


SIGDIAL 2019



**20th Annual Meeting of the
Special Interest Group on Discourse and
Dialogue**



Proceedings of the Conference

11-13 September 2019
Stockholm, Sweden

In cooperation with:

Association for Computational Linguistics (ACL)

International Speech Communication Association (ISCA)

Association for the Advancement of Artificial Intelligence (AAAI)

We thank our sponsors:

Honda Research Institute

Amazon Alexa

Rasa Technologies

Toshiba Research Europe

Interactions

Spotify

Educational Testing Service (ETS)

KTH Royal Institute of Technology

Microsoft Research

Apple

Monash University

Platinum



Gold



Silver



Measuring the Power of Learning.™



MONASH
INFORMATION
TECHNOLOGY

Bronze

TOSHIBA

In cooperation with



©2019 The Association for Computational Linguistics

Association for Computational Linguistics (ACL)
209 N. Eighth Street
Stroudsburg, PA 18360
USA
Tel: +1-570-476-8006
Fax: +1-570-476-0860
acl@aclweb.org

ISBN 978-1-950737-61-1

Introduction

We are excited to welcome you to SIGDIAL 2019, the 20th Annual Meeting of the Special Interest Group on Discourse and Dialogue. This year the conference is being held in Stockholm, Sweden, on September 11-13, with the Satellite Event YRRSDS 2019 (Young Researchers' Roundtable on Spoken Dialog Systems), and in close temporal proximity with Interspeech 2019, held in Gratz, Austria, and SemDial 2019, held in London, UK.

The SIGDIAL conference is a premier publication venue for research in discourse and dialogue. This year, the program includes three keynote talks, six oral presentation sessions, three poster sessions including six demonstrations, a panel entitled “The Future of Dialogue Research” organized by Phil Cohen, and a special session entitled “Implications of Deep Learning for Dialogue Modeling” organized by Nigel Ward, Yun-Nung (Vivian) Chen, Tatsuya Kawahara and Gabriel Skantze.

We received a record 146 submissions this year, about one third more than the submissions received in 2018. The 146 submissions comprised 93 long papers, 43 short papers and 10 demo descriptions. All submissions received at least three reviews. When making our selections for the program, we carefully considered the reviews and the comments made during the discussions among reviewers. The members of the Program Committee did an excellent job in reviewing the submitted papers, and we thank them for their essential role in selecting the accepted papers and helping produce a high quality program for the conference. In line with the SIGDIAL tradition, our aim has been to create a balanced program that accommodates as many favourably rated papers as possible. We accepted 51 papers: 33 long papers—three of which were converted to short papers, 13 short papers, and five demo descriptions. These numbers give an overall acceptance rate of 35%, with the following rates for the different types of papers: 35% for long papers, 30% for short papers and 50% for demo descriptions. It is worth noting that the acceptance rate for long papers was significantly lower than that of previous years – a result of the unusually large number of submissions.

Each of the three conference days features one keynote and one poster session, with the remaining time given to oral presentations, demos, the panel and the special session. The oral presentations comprise 16 of the long papers and three long papers selected for the special session. The three poster sessions feature the remaining long papers, all the short papers and two work-in-progress special session papers. In terms of content, about a quarter of the accepted papers discuss datasets and evaluation issues, and approximately half employ deep learning to address problems in discourse and dialogue—a trend also exhibited in recent Language Technology conferences. Finally, this SIGDIAL features an invited demo that showcases research conducted in the department of Robotics, Perception and Learning at KTH, the host institution.

A conference of this scale requires advice, help and enthusiastic participation of many parties, and we have a big ‘thank you’ to say to all of them. Regarding the program, we thank our three keynote speakers, Dan Bohus (Microsoft Research, Redmond, Washington, US), Mirella Lapata (University of Edinburgh, UK) and Helen Meng (Chinese University of Hong Kong, China) for their inspiring talks on situated interaction, learning neural natural language interfaces, and dialogue research application to healthcare, e-commerce and education. We also thank the organizer of the panel on the Future of Dialogue Research, and the organizers of the special session on Implications of Deep Learning for Dialogue Modeling. We are grateful for their smooth and efficient coordination with the main conference. In addition, we thank Alex Papangelis, Mentoring Chair for SIGDIAL 2019, for his dedicated work on the mentoring process. The goal of mentoring is to assist authors of papers that contain important ideas but require significant stylistic modifications. In total, seven of the accepted papers received mentoring, and we thank our mentoring team for their excellent support of the authors.

We extend special thanks to our Local Chair, Gabriel Skantze, and his team, including the student

volunteers who provide on-site assistance. SIGDIAL 2019 would not have been possible without their effort in arranging the conference venue, handling registration, making banquet arrangements, numerous preparations for the conference, and last but not least, Gabriel's personal contributions, which exceeded those of a local organizer.

Mikio Nakano, our Sponsorship Chair, has conducted the massive task of recruiting and liaising with our conference sponsors, many of whom continue to contribute year after year. Sponsorships support valuable aspects of the program, such as lunches and the conference banquet. We thank Mikio for his dedicated work and his assistance with conference planning. We gratefully acknowledge the support of our sponsors: (Platinum level) Honda Research Institute, Interactions and Microsoft Research; (Gold level) Amazon Alexa, Apple, Rasa Technologies and Spotify; (Silver level) Educational Testing Service (ETS) and Monash University; and (Bronze level) Toshiba Research Europe. We also thank the KTH Royal Institute of Technology for its generous sponsorship as host.

Koichiro Yoshino, our publicity chair, was tireless in the design and maintenance of the SIGDIAL 2019 website, cheerfully coping with multiple and constant changes; and Stefan Ultes, our publication chair, capped the long organizational process by putting together these high quality conference proceedings. We thank the SIGdial board, both current and emeritus officers, Gabriel Skantze, Mikio Nakano, Vikram Ramanarayanan, Ethan Selfridge, Kallirroi Georgila, Jason Williams and Amanda Stent, for their advice and support from beginning to end.

We once again thank our program committee members for committing their time to help us select an excellent technical program. Finally, we thank all the authors who submitted to the conference and all the conference participants for making SIGDIAL 2019 a success and for growing the research areas of discourse and dialogue with their fine work.

Satoshi Nakamura, General Chair

Milica Gašić and Ingrid Zukerman, Program Co-Chairs

Committee Listings

General Chair:

Satoshi Nakamura, Nara Institute of Science and Technology, Japan

Program Chairs:

Milica Gašić, Heinrich Heine University Düsseldorf, Germany

Ingrid Zukerman, Monash University, Australia

Local Chair:

Gabriel Skantze, KTH Royal Institute of Technology, Sweden

Sponsorship Chair:

Mikio Nakano, Honda Research Institute, Japan

Mentoring Chair:

Alexandros Papangelis, Uber AI, United States

Publication Chair:

Stefan Ultes, Daimler AG, Germany

Publicity Chair:

Koichiro Yoshino, Nara Institute of Science and Technology, Japan

SIGdial Officers:

President: Gabriel Skantze, KTH Royal Institute of Technology, Sweden

Vice President: Mikio Nakano, Honda Research Institute, Japan

Secretary: Vikram Ramanarayanan, Educational Testing Service (ETS) Research, United States

Treasurer: Ethan Selfridge, Interactions, United States

President Emeritus: Jason Williams, Apple, United States

Vice President Emeritus: Kallirroi Georgila, University of Southern California, United States

Program Committee:

Masahiro Araki, Kyoto Institute of Technology, Japan

Ron Artstein, USC Institute for Creative Technologies, United States

Timo Baumann, University of Hamburg, Germany

Frederic Bechet, Aix Marseille Université - LIS/CNRS, France

Steve Beet, Aculab plc, United Kingdom

Jose Miguel Benedi, Universitat Politècnica de València, Spain

Luciana Benotti, Universidad Nacional de Córdoba, Argentina

Nate Blaylock, Nuance Communications, United States

Dan Bohus, Microsoft Research, United States

Johan Boye, KTH, Sweden

Paweł Budzianowski, University of Cambridge, United Kingdom

Hendrik Buschmeier, Bielefeld University, Germany

Andrew Caines, University of Cambridge, United Kingdom
Giuseppe Carenini, University of British Columbia, Canada
Christophe Cerisara, Université de Lorraine, CNRS, LORIA, France
Joyce Chai, Michigan State University, United States
Senthil Chandramohan, Microsoft, United States
Lin Chen, Head of AI, Cambia Health Solutions, United States
Philip Cohen, Monash University, Australia
Mark Core, University of Southern California, United States
Heriberto Cuayahuitl, University of Lincoln, United Kingdom
Vera Demberg, Saarland University, Germany
Nina Dethlefs, University of Hull, United Kingdom
David DeVault, University of Southern California, United States
Barbara Di Eugenio, University of Illinois at Chicago, United States
Giuseppe Di Fabbrizio, VUI, Inc., United States
Maxine Eskenazi, Carnegie Mellon University, United States
Keelan Evanini, Educational Testing Service, United States
Mauro Falcone, Fondazione Ugo Bordoni, Italy
Raquel Fernandez, ILLC, University of Amsterdam, Netherlands
Kallirroi Georgila, University of Southern California, United States
Alborz Geramifard, Facebook, United States
Jonathan Ginzburg, Université Paris-Diderot (Paris 7), France
Ivan Habernal, Technische Universität Darmstadt, Germany
Dilek Hakkani-Tur, Amazon Alexa AI, United States
Helen Hastie, Heriot-Watt University, United Kingdom
Michael Heck, Heinrich Heine University, Germany
Ryuichiro Higashinaka, NTT Media Intelligence Labs., Japan
Takuya Hiraoka, NEC Central Research Laboratories, Japan
Keikichi Hirose, University of Tokyo, Japan
David M. Howcroft, Heriot-Watt University, United Kingdom
Michimasa Inaba, University of Electro-Communications, Japan
Koji Inoue, Kyoto University, Japan
David Janiszek, Université Paris Descartes, France
Kristiina Jokinen, AIRC, AIST, Japan
Arne Jonsson, Linkping University, Sweden
Pamela Jordan, University of Pittsburgh, United States
Filip Jurcicek, Apple Inc., United Kingdom
Tatsuya Kawahara, Kyoto University, Japan
Simon Keizer, Emotech North, United Kingdom
Casey Kennington, Boise State University, United States
Chandra Khatri, Uber AI, United States
Norihide Kitaoka, Toyohashi University of Technology, Japan
Kazunori Komatani, Osaka University, Japan
Stefan Kopp, Bielefeld University, Germany
Fabrice Lefevre, Avignon University, France
James Lester, North Carolina State University, United States
Junyi Jessy Li, University of Texas at Austin, United States
Pierre Lison, Norwegian Computing Centre, Norway
Diane Litman, University of Pittsburgh, United States
Bing Liu, Facebook, United States

Eduardo Lleida Solano, University of Zaragoza, Spain
José Lopes, Heriot Watt University, United Kingdom
Ramon Lopez-Cozar, University of Granada, Spain
Annie Louis, University of Edinburgh, United Kingdom
Nurul Lubis, Heinrich Heine University, Germany
Matthew Marge, Army Research Laboratory, United States
Helen Meng, Chinese University of Hong Kong, China
Teruhisa Misu, Honda Research Institute USA, United States
Satoshi Nakamura, Nara Institute of Science and Technology, Japan
Mikio Nakano, Honda Research Institute, Japan
Shashi Narayan, Google, United Kingdom
Douglas O'Shaughnessy, INRS-EMT (Univ. of Quebec), Canada
Mari Ostendorf, University of Washington, United States
Alexandros Papangelis, Uber AI, United States
Volha Petukhova, Saarland University, Germany
Paul Piwek, The Open University, United Kingdom
Andrei Popescu-Belis, HEIG-VD / HES-SO, Switzerland
Rashmi Prasad, Interactions Corporation, United States
Matthew Purver, Queen Mary University of London, United Kingdom
Vikram Ramanarayanan, Educational Testing Service R&D, United States
Antoine Raux, Apple, United States
Ehud Reiter, University of Aberdeen, United Kingdom
Norbert Reithinger, DFKI GmbH, Germany
Giuseppe Riccardi, University of Trento, Italy
Antonio Roque, Tufts University, United States
Carolyn Rose, Carnegie Mellon University, United States
Sophie Rosset, LIMSI, CNRS, Universit Paris-Saclay, France
Sakriani Sakti, Nara Institute of Science and Technology (NAIST) / RIKEN AIP, Japan
Ruhi Sarikaya, Amazon, United States
Niko Schenk, Goethe University Frankfurt am Main, Germany
David Schlangen, Bielefeld University, Germany
Ethan Selfridge, Interactions Corp, United States
Gabriel Skantze, KTH Speech Music and Hearing, Sweden
Manfred Stede, University of Potsdam, Germany
Georg Stemmer, Intel Corp., Germany
Matthew Stone, Rutgers University, United States
Svetlana Stoyanchev, Interactions Corporation, United States
Kristina Striegnitz, Union College, United States
Hiroaki Sugiyama, NTT Communication Science Labs., Japan
Antnio Teixeira, DETI/IEETA, University of Aveiro, Portugal
Joel Tetreault, Grammarly, United States
Sofia Thunberg, Linköping University, Sweden
Takenobu Tokunaga, Tokyo Institute of Technology, Japan
David Traum, University of Southern California, United States
Bo-Hsiang Tseng, University of Cambridge, United Kingdom
Gokhan Tur, Uber, United States
Stefan Ultes, Daimler AG, Germany
David Vandyke, Apple, United Kingdom
Marilyn Walker, University of California Santa Cruz, United States

Hsin-Min Wang, Academia Sinica, Taiwan
Nigel Ward, University of Texas at El Paso, United States
Jason D Williams, Microsoft Research, United States
Yen-chen Wu, University of Cambridge, United Kingdom
Koichiro Yoshino, Nara Institute of Science and Technology, Japan
Steve Young, Cambridge University, United Kingdom
Kai Yu, Shanghai Jiao Tong University, China
Zhou Yu, University of California, Davis, United States

Secondary Reviewers: Alexandre Denis, Sahar Ghannay, Churn-Jung Liao, Yi Ma, Sonja Stange,
Hiroki Tanaka, Martin Villalba

Mentors:

Dimitrios Alikaniotis, Grammarly, United States
Timo Baumann, University of Hamburg, Germany
Hendrik Buschmeier, Bielefeld University, Germany
Ivan Habernal, Technische Universität Darmstadt, Germany
Kornel Laskowski, Carnegie Mellon University, United States
Stefan Ultes, Daimler AG, Germany
David Vandyke, Apple, United Kingdom

Invited Speakers:

Dan Bohus, Microsoft Research, United States
Mirella Lapata, University of Edinburgh, United Kingdom
Helen Meng, Chinese University of Hong Kong, China

Table of Contents

<i>Deep Reinforcement Learning For Modeling Chit-Chat Dialog With Discrete Attributes</i> Chinnadhurai Sankar and Sujith Ravi	1
<i>Improving Interaction Quality Estimation with BiLSTMs and the Impact on Dialogue Policy Learning</i> Stefan Ultes	11
<i>Lifelong and Interactive Learning of Factual Knowledge in Dialogues</i> Sahisnu Mazumder, Bing Liu, Shuai Wang and Nianzu Ma	21
<i>Few-Shot Dialogue Generation Without Annotated Data: A Transfer Learning Approach</i> Igor Shalyminov, Sungjin Lee, Arash Eshghi and Oliver Lemon	32
<i>SIM: A Slot-Independent Neural Model for Dialogue State Tracking</i> Chenguang Zhu, Michael Zeng and Xuedong Huang	40
<i>Simple, Fast, Accurate Intent Classification and Slot Labeling for Goal-Oriented Dialogue Systems</i> Arshit Gupta, John Hewitt and Katrin Kirchhoff	46
<i>Time Masking: Leveraging Temporal Information in Spoken Dialogue Systems</i> Rylan Conway and Mathias Lambert	56
<i>To Combine or Not To Combine? A Rainbow Deep Reinforcement Learning Agent for Dialog Policies</i> Dirk Vath and Ngoc Thang Vu	62
<i>Contextualized Representations for Low-resource Utterance Tagging</i> Bhargavi Paranjape and Graham Neubig	68
<i>Capturing Dialogue State Variable Dependencies with an Energy-based Neural Dialogue State Tracker</i> Anh Duong Trinh, Robert J. Ross and John D. Kelleher	75
<i>Leveraging Non-Conversational Tasks for Low Resource Slot Filling: Does it help?</i> Samuel Louvan and Bernardo Magnini	85
<i>Collaborative Multi-Agent Dialogue Model Training Via Reinforcement Learning</i> Alexandros Papangelis, Yi-Chia Wang, Piero Molino and Gokhan Tur	92
<i>Scoring Interactional Aspects of Human-Machine Dialog for Language Learning and Assessment using Text Features</i> Vikram Ramanarayanan, Matthew Mulholland and Yao Qian	103
<i>Spoken Conversational Search for General Knowledge</i> Lina M. Rojas Barahona, Pascal Bellec, Benoit Besset, Martinho Dossantos, Johannes Heinecke, munshi asadullah, Olivier Leblouch, Jeanyves. Lancien, Geraldine Damnati, Emmanuel Mory and Frederic Herledan	110
<i>Graph2Bots, Unsupervised Assistance for Designing Chatbots</i> Jean-Leon Bouraoui, Sonia Le Meitour, Romain Carbou, Lina M. Rojas Barahona and Vincent Lemaire	114
<i>On a Chatbot Conducting Dialogue-in-Dialogue</i> Boris Galitsky, Dmitry Ilvovsky and Elizaveta Goncharova	118
<i>DeepCopy: Grounded Response Generation with Hierarchical Pointer Networks</i> Semih Yavuz, Abhinav Rastogi, Guan-Lin Chao and Dilek Hakkani-Tur	122

<i>Towards End-to-End Learning for Efficient Dialogue Agent by Modeling Looking-ahead Ability</i> Zhuoxuan Jiang, Xian-Ling Mao, Ziming Huang, Jie Ma and Shaochun Li	133
<i>Unsupervised Dialogue Spectrum Generation for Log Dialogue Ranking</i> Xinnuo Xu, Yizhe Zhang, Lars Liden and Sungjin Lee	143
<i>Tree-Structured Semantic Encoder with Knowledge Sharing for Domain Adaptation in Natural Language Generation</i> Bo-Hsiang Tseng, Paweł Budzianowski, Yen-chen Wu and Milica Gasic	155
<i>Structured Fusion Networks for Dialog</i> Shikib Mehri, Tejas Srinivasan and Maxine Eskenazi	165
<i>Flexibly-Structured Model for Task-Oriented Dialogues</i> Lei Shu, Piero Molino, Mahdi Namazifar, Hu Xu, Bing Liu, Huaixiu Zheng and Gokhan Tur ..	178
<i>FriendsQA: Open-Domain Question Answering on TV Show Transcripts</i> Zhengzhe Yang and Jinho D. Choi	188
<i>Foundations of Collaborative Task-Oriented Dialogue: What’s in a Slot?</i> Philip Cohen	198
<i>Speaker-adapted neural-network-based fusion for multimodal reference resolution</i> Diana Kleingarn, Nima Nabizadeh, Martin Heckmann and Dorothea Kolossa	210
<i>Learning Question-Guided Video Representation for Multi-Turn Video Question Answering</i> Guan-Lin Chao, Abhinav Rastogi, Semih Yavuz, Dilek Hakkani-Tur, Jindong Chen and Ian Lane	215
<i>Zero-shot transfer for implicit discourse relation classification</i> Murathan Kurfalı and Robert Östling	226
<i>A Quantitative Analysis of Patients’ Narratives of Heart Failure</i> Sabita Acharya, Barbara Di Eugenio, Andrew Boyd, Richard Cameron, Karen Dunn Lopez, Pamela Martyn-Nemeth, Debaleena Chattopadhyay, Pantea Habibi, Carolyn Dickens, Haleh Vatani and Amer Ardati	232
<i>TDDiscourse: A Dataset for Discourse-Level Temporal Ordering of Events</i> Aakanksha Naik, Luke Breitfeller and Carolyn Rose	239
<i>Real Life Application of a Question Answering System Using BERT Language Model</i> Francesca Alloatti, Luigi Di Caro and Gianpiero Sportelli	250
<i>Hierarchical Multi-Task Natural Language Understanding for Cross-domain Conversational AI: HER-MIT NLU</i> Andrea Vanzo, Emanuele Bastianelli and Oliver Lemon	254
<i>Dialog State Tracking: A Neural Reading Comprehension Approach</i> Shuyang Gao, Abhishek Sethi, Sanchit Agarwal, Tagyoung Chung and Dilek Hakkani-Tur	264
<i>Cross-Corpus Data Augmentation for Acoustic Addressee Detection</i> Oleg Akhtiamov, Ingo Siegert, Alexey Karpov and Wolfgang Minker	274
<i>A Scalable Method for Quantifying the Role of Pitch in Conversational Turn-Taking</i> Kornel Laskowski, Marcin Wlodarczak and Mattias Heldner	284

<i>A Large-Scale User Study of an Alexa Prize Chatbot: Effect of TTS Dynamism on Perceived Quality of Social Dialog</i>	
Michelle Cohn, Chun-Yen Chen and Zhou Yu	293
<i>Influence of Time and Risk on Response Acceptability in a Simple Spoken Dialogue System</i>	
Andisheh Partovi and Ingrid Zukerman	307
<i>Characterizing the Response Space of Questions: a Corpus Study for English and Polish</i>	
Jonathan Ginzburg, Zulipiye Yusupujang, Chuyuan Li, Kexin Ren and Paweł Łupkowski	320
<i>From Explainability to Explanation: Using a Dialogue Setting to Elicit Annotations with Justifications</i>	
Nazia Attari, Martin Heckmann and David Schlangen	331
<i>Prediction of User Emotion and Dialogue Success Using Audio Spectrograms and Convolutional Neural Networks</i>	
Athanasios Lykartsis and Margarita Kotti	336
<i>Modelling Adaptive Presentations in Human-Robot Interaction using Behaviour Trees</i>	
Nils Axelsson and Gabriel Skantze	345
<i>Coached Conversational Preference Elicitation: A Case Study in Understanding Movie Preferences</i>	
Filip Radlinski, Krisztian Balog, Bill Byrne and Karthik Krishnamoorthi	353
<i>A Crowd-based Evaluation of Abuse Response Strategies in Conversational Agents</i>	
Amanda Cercas Curry and Verena Rieser	361
<i>A Dynamic Strategy Coach for Effective Negotiation</i>	
Yiheng Zhou, He He, Alan W Black and Yulia Tsvetkov	367
<i>Investigating Evaluation of Open-Domain Dialogue Systems With Human Generated Multiple References</i>	
Prakhar Gupta, Shikib Mehri, Tiancheng Zhao, Amy Pavel, Maxine Eskenazi and Jeffrey Bigham	379
<i>User Evaluation of a Multi-dimensional Statistical Dialogue System</i>	
Simon Keizer, Ondřej Dušek, Xingkun Liu and Verena Rieser	392
<i>Dialogue Act Classification in Team Communication for Robot Assisted Disaster Response</i>	
Tatiana Anikina and Ivana Kruijff-Korbayova	399
<i>Multi-Task Learning of System Dialogue Act Selection for Supervised Pretraining of Goal-Oriented Dialogue Policies</i>	
Sarah McLeod, Ivana Kruijff-Korbayova and Bernd Kiefer	411
<i>B. Rex: a dialogue agent for book recommendations</i>	
Mitchell Abrams, Luke Gessler and Matthew Marge	418
<i>SpaceRefNet: a neural approach to spatial reference resolution in a real city environment</i>	
Dmytro Kalpakchi and Johan Boye	422
<i>Which aspects of discourse relations are hard to learn? Primitive decomposition for discourse relation classification</i>	
Charlotte Roze, Chloé Braud and Philippe Muller	432
<i>Discourse Relation Prediction: Revisiting Word Pairs with Convolutional Networks</i>	
Siddharth Varia, Christopher Hidey and Tuhin Chakrabarty	442

Conference Program

11 September 2019

09:00–09:20 *Welcome*

09:20–10:20 *Keynote 1 - Learning Natural Language Interfaces with Neural Models*
Mirella Lapata

10:20–10:50 *Coffee Break*

10:50–12:05 **Session 1 - Policy and Knowledge**

Deep Reinforcement Learning For Modeling Chit-Chat Dialog With Discrete Attributes

Chinnadhurai Sankar and Sujith Ravi

Improving Interaction Quality Estimation with BiLSTMs and the Impact on Dialogue Policy Learning

Stefan Ultes

Lifelong and Interactive Learning of Factual Knowledge in Dialogues

Sahisnu Mazumder, Bing Liu, Shuai Wang and Nianzu Ma

12:05–13:20 *Lunch*

11 September 2019 (continued)

13:20–15:10 Poster and Demos 1

Few-Shot Dialogue Generation Without Annotated Data: A Transfer Learning Approach

Igor Shalyminov, Sungjin Lee, Arash Eshghi and Oliver Lemon

SIM: A Slot-Independent Neural Model for Dialogue State Tracking

Chenguang Zhu, Michael Zeng and Xuedong Huang

Simple, Fast, Accurate Intent Classification and Slot Labeling for Goal-Oriented Dialogue Systems

Arshit Gupta, John Hewitt and Katrin Kirchhoff

Time Masking: Leveraging Temporal Information in Spoken Dialogue Systems

Rylan Conway and Mathias Lambert

To Combine or Not To Combine? A Rainbow Deep Reinforcement Learning Agent for Dialog Policies

Dirk V ath and Ngoc Thang Vu

Contextualized Representations for Low-resource Utterance Tagging

Bhargavi Paranjape and Graham Neubig

Capturing Dialogue State Variable Dependencies with an Energy-based Neural Dialogue State Tracker

Anh Duong Trinh, Robert J. Ross and John D. Kelleher

Leveraging Non-Conversational Tasks for Low Resource Slot Filling: Does it help?

Samuel Louvan and Bernardo Magnini

Collaborative Multi-Agent Dialogue Model Training Via Reinforcement Learning

Alexandros Papangelis, Yi-Chia Wang, Piero Molino and Gokhan Tur

Scoring Interactional Aspects of Human-Machine Dialog for Language Learning and Assessment using Text Features

Vikram Ramanarayanan, Matthew Mulholland and Yao Qian

Spoken Conversational Search for General Knowledge

Lina M. Rojas Barahona, Pascal Bellec, Benoit Besset, Martinho Dossantos, Johannes Heinecke, munshi asadullah, Olivier Leblouch, Jeanyves Lancien, Geraldine Damnati, Emmanuel Mory and Frederic Herledan

11 September 2019 (continued)

Graph2Bots, Unsupervised Assistance for Designing Chatbots

Jean-Leon Bouraoui, Sonia Le Meitour, Romain Carbou, Lina M. Rojas Barahona and Vincent Lemaire

On a Chatbot Conducting Dialogue-in-Dialogue

Boris Galitsky, Dmitry Ilvovsky and Elizaveta Goncharova

15:10–15:40 *Coffee Break*

15:40–16:55 **Session 2 (Special Session) - Implications of Deep Learning for Dialogue Modeling**

DeepCopy: Grounded Response Generation with Hierarchical Pointer Networks

Semih Yavuz, Abhinav Rastogi, Guan-Lin Chao and Dilek Hakkani-Tur

Towards End-to-End Learning for Efficient Dialogue Agent by Modeling Looking-ahead Ability

Zhuoxuan Jiang, Xian-Ling Mao, Ziming Huang, Jie Ma and Shaochun Li

Unsupervised Dialogue Spectrum Generation for Log Dialogue Ranking

Xinnuo Xu, Yizhe Zhang, Lars Liden and Sungjin Lee

16:55–17:55 *Panel: The Future of Dialogue Research*

Organizer: Phil Cohen

Vikram Ramanarayanan, Educational Testing Service (ETS) Research, USA

Sujith Savi, Google, USA

Gabriel Skantze, KTH Royal Institute of Technology, Sweden

18:15–19:45 *Reception*

12 September 2019

09:00–10:00 *Keynote 2 - Situated Interaction*
Dan Bohus

10:00–10:30 *Coffee Break*

10:30–11:45 **Session 3 - Generation and End-to-end Dialogue Systems**

Tree-Structured Semantic Encoder with Knowledge Sharing for Domain Adaptation in Natural Language Generation

Bo-Hsiang Tseng, Paweł Budzianowski, Yen-chen Wu and Milica Gasic

Structured Fusion Networks for Dialog

Shikib Mehri, Tejas Srinivasan and Maxine Eskenazi

Flexibly-Structured Model for Task-Oriented Dialogues

Lei Shu, Piero Molino, Mahdi Namazifar, Hu Xu, Bing Liu, Huaixiu Zheng and Gokhan Tur

11:45–13:00 *Lunch*

13:00–14:15 **Poster and Demos 2**

FriendsQA: Open-Domain Question Answering on TV Show Transcripts

Zhengzhe Yang and Jinho D. Choi

Foundations of Collaborative Task-Oriented Dialogue: What's in a Slot?

Philip Cohen

Speaker-adapted neural-network-based fusion for multimodal reference resolution

Diana Kleingarn, Nima Nabizadeh, Martin Heckmann and Dorothea Kolossa

Learning Question-Guided Video Representation for Multi-Turn Video Question Answering

Guan-Lin Chao, Abhinav Rastogi, Semih Yavuz, Dilek Hakkani-Tur, Jindong Chen and Ian Lane

12 September 2019 (continued)

Zero-shot transfer for implicit discourse relation classification

Murathan Kurfalı and Robert Östling

A Quantitative Analysis of Patients' Narratives of Heart Failure

Sabita Acharya, Barbara Di Eugenio, Andrew Boyd, Richard Cameron, Karen Dunn Lopez, Pamela Martyn-Nemeth, Debaleena Chattopadhyay, Pantea Habibi, Carolyn Dickens, Haleh Vatani and Amer Ardati

TDDiscourse: A Dataset for Discourse-Level Temporal Ordering of Events

Aakanksha Naik, Luke Breitfeller and Carolyn Rose

Real-time Generation of Unambiguous Spatial Referring Expressions

Fethiye Irmak Dogan, Sinan Kalkan and Iolanda Leite

Real Life Application of a Question Answering System Using BERT Language Model

Francesca Alloatti, Luigi Di Caro and Gianpiero Sportelli

14:15–15:05 Session 4 - Understanding and Dialogue State Tracking

Hierarchical Multi-Task Natural Language Understanding for Cross-domain Conversational AI: HERMIT NLU

Andrea Vanzo, Emanuele Bastianelli and Oliver Lemon

Dialog State Tracking: A Neural Reading Comprehension Approach

Shuyang Gao, Abhishek Sethi, Sanchit Agarwal, Tagyoung Chung and Dilek Hakkani-Tur

15:05–15:35 Coffee Break

12 September 2019 (continued)

15:35–16:25 Session 5 - Acoustics

Cross-Corpus Data Augmentation for Acoustic Addressee Detection
Oleg Akhramov, Ingo Siegert, Alexey Karpov and Wolfgang Minker

A Scalable Method for Quantifying the Role of Pitch in Conversational Turn-Taking
Kornel Laskowski, Marcin Włodarczyk and Mattias Heldner

16:25–17:10 Sponsor Session

18:30–21:00 *Banquet at Vasa Museum*

13 September 2019

09:00–10:00 *The Many Facets of Dialog*
Helen Meng

10:00–10:30 *Coffee Break*

10:30–11:45 **Session 6 - Evaluation and Data**

A Large-Scale User Study of an Alexa Prize Chatbot: Effect of TTS Dynamism on Perceived Quality of Social Dialog

Michelle Cohn, Chun-Yen Chen and Zhou Yu

Influence of Time and Risk on Response Acceptability in a Simple Spoken Dialogue System

Andisheh Partovi and Ingrid Zukerman

Characterizing the Response Space of Questions: a Corpus Study for English and Polish

Jonathan Ginzburg, Zulipiye Yusupujiang, Chuyuan Li, Kexin Ren and Paweł Łupkowski

11:45–13:00 *Lunch*

13 September 2019 (continued)

13:00–14:50 Poster and Demos 3

From Explainability to Explanation: Using a Dialogue Setting to Elicit Annotations with Justifications

Nazia Attari, Martin Heckmann and David Schlangen

Prediction of User Emotion and Dialogue Success Using Audio Spectrograms and Convolutional Neural Networks

Athanasios Lykartsis and Margarita Kotti

Modelling Adaptive Presentations in Human-Robot Interaction using Behaviour Trees

Nils Axelsson and Gabriel Skantze

Coached Conversational Preference Elicitation: A Case Study in Understanding Movie Preferences

Filip Radlinski, Krisztian Balog, Bill Byrne and Karthik Krishnamoorthi

A Crowd-based Evaluation of Abuse Response Strategies in Conversational Agents

Amanda Cercas Curry and Verena Rieser

A Dynamic Strategy Coach for Effective Negotiation

Yiheng Zhou, He He, Alan W Black and Yulia Tsvetkov

Investigating Evaluation of Open-Domain Dialogue Systems With Human Generated Multiple References

Prakhar Gupta, Shikib Mehri, Tiancheng Zhao, Amy Pavel, Maxine Eskenazi and Jeffrey Bigham

User Evaluation of a Multi-dimensional Statistical Dialogue System

Simon Keizer, Ondřej Dušek, Xingkun Liu and Verena Rieser

Dialogue Act Classification in Team Communication for Robot Assisted Disaster Response

Tatiana Anikina and Ivana Kruijff-Korbayova

Multi-Task Learning of System Dialogue Act Selection for Supervised Pretraining of Goal-Oriented Dialogue Policies

Sarah McLeod, Ivana Kruijff-Korbayova and Bernd Kiefer

B. Rex: a dialogue agent for book recommendations

Mitchell Abrams, Luke Gessler and Matthew Marge

13 September 2019 (continued)

14:20–14:50 *Coffee Break (during Poster and Demos 3)*

14:50–16:05 **Session 7 - Discourse**

SpaceRefNet: a neural approach to spatial reference resolution in a real city environment

Dmytro Kalpakchi and Johan Boye

Which aspects of discourse relations are hard to learn? Primitive decomposition for discourse relation classification

Charlotte Roze, Chloé Braud and Philippe Muller

Discourse Relation Prediction: Revisiting Word Pairs with Convolutional Networks

Siddharth Varia, Christopher Hidey and Tuhin Chakrabarty

16:05–16:20 *Short Coffee Break*

16:20–17:20 *Business meeting, Awards and Closing*

Keynote Abstracts

Keynote 1 - Learning Natural Language Interfaces with Neural Models

Mirella Lapata

Professor, School of Informatics, University of Edinburgh, UK

Abstract

In Spike Jonze's futuristic film "Her", Theodore, a lonely writer, forms a strong emotional bond with Samantha, an operating system designed to meet his every need. Samantha can carry on seamless conversations with Theodore, exhibits a perfect command of language, and is able to take on complex tasks. She filters his emails for importance, allowing him to deal with information overload, she proactively arranges the publication of Theodore's letters, and is able to give advice using common sense and reasoning skills.

In this talk I will present an overview of recent progress on learning natural language interfaces which might not be as clever as Samantha but nevertheless allow users to interact with various devices and services using every day language. I will address the structured prediction problem of mapping natural language utterances onto machine-interpretable representations and outline the various challenges it poses. For example, the fact that the translation of natural language to formal language is highly non-isomorphic, data for model training is scarce, and natural language can express the same information need in many different ways. I will describe a general modeling framework based on neural networks which tackles these challenges and improves the robustness of natural language interfaces.

Biography

Mirella Lapata is professor of natural language processing in the School of Informatics at the University of Edinburgh. Her research focuses on getting computers to understand, reason with, and generate natural language. She is the first recipient (2009) of the British Computer Society and Information Retrieval Specialist Group (BCS/IRSG) Karen Sparck Jones award and a Fellow of the Royal Society of Edinburgh. She has also received best paper awards in leading NLP conferences and has served on the editorial boards of the Journal of Artificial Intelligence Research, the Transactions of the ACL, and Computational Linguistics. She was president of SIGDAT (the group that organizes EMNLP) in 2018.

Keynote 2 - Situated Interaction

Dan Bohus *Senior Principal Researcher, Perception and Interaction Group, Microsoft Research, Redmond, Washington, US*

Abstract

Physically situated dialog is a complex, multimodal affair that goes well beyond the spoken word. When interacting with each other, people incrementally coordinate their actions to simultaneously resolve several different problems: they manage engagement, coordinate on taking turns, recognize intentions, and establish and maintain common ground as a basis for contributing to the conversation. A wide array of non-verbal signals are brought to bear. Proximity and body pose, attention and gaze, head nods and hand gestures, prosody and facial expressions, all play important roles in the intricate, mixed-initiative, fluidly coordinated process we call interaction. And just like a couple of decades ago advances in speech recognition opened up the field of spoken dialog systems, today advances in vision and other perceptual technologies are again opening up new horizons – we are starting to be able to build machines that can understand these social signals and the physical world around them, and begin to participate in physically situated interactions and collaborations with people.

In this talk, using a number of research vignettes from my work, I will draw attention to some of the challenges and opportunities that lie ahead of us in this exciting space. In particular, I will discuss issues with managing engagement and turn-taking in multiparty open-world settings, and more generally highlight the importance of timing and fine-grained coordination in situated interaction. Finally, I will conclude by describing a framework that promises to simplify the development of physically situated interactive systems and enable more research and faster progress in this area.

Biography

Dan Bohus is a Senior Principal Researcher in the Perception and Interaction Group at Microsoft Research. His work centers on the study and development of computational models for physically situated spoken language interaction and collaboration. The long term question that shapes his research agenda is how can we enable interactive systems to reason more deeply about their surroundings and seamlessly participate in open-world, multiparty dialog and collaboration with people? Prior to joining Microsoft Research, Dan obtained his Ph.D. from Carnegie Mellon University.

Keynote 3 - The Many Facets of Dialog

Helen Meng *SProfessor, Department of Systems Engineering and Engineering Management, Chinese University of Hong Kong, China*

Abstract

Dialog is a most fascinating form of human communication. The back-and-forth exchanges convey the speaker's message to the listener, and the listener can derive information about the speaker's thoughts, intent, well-being, emotions and much more. This talk presents an overview of dialog research that concerns our group at The Chinese University of Hong Kong. In the domain of education and learning, we have been recording in-class student group discussions in the flipped-classroom setting of a freshman elite mathematics course. We investigate features in the weekly, within-group dialogs that may relate to class performance and learning efficacy. In the domain of e-commerce, we are developing dialog models based on approximately 20 million conversation turns, to support a virtual shopping assistant in customer inquiries and orders, logistics tracking, etc. In the domain of health and wellbeing, we are capturing and analysing dialogs between health professionals (or their virtual equivalent) and subjects in cognitive screening tests. We also conduct research in both semantic interpretation and dialog state tracking, as well as affective design of virtual conversational assistants. For the former, we have developed a Convex Polytopic Model for extracting a knowledge representation from user inputs in dialog turns by generating a compact convex polytope to enclose all the data points projected to a latent semantic space. The polytope vertices represent extracted semantic concepts. Each user input can then be "interpreted" as a sequence of polytope vertices which represent the user's goals and dialog states. For the latter, we have developed a multimodal, multi-task, deep learning framework to infer the user's emotive state and emotive state change simultaneously. This enables virtual conversational assistants to understand the emotive state in the user's input and to generate an appropriate emotive system response in the dialog turn, which will further influence the user's emotive state in the subsequent dialog turn. Such an affective design will be able to enhance user experience in conversational dialogs with intelligent virtual assistants.

Biography

Helen Meng is Patrick Huen Wing Ming Professor of Systems Engineering and Engineering Management at The Chinese University of Hong Kong (CUHK). She is the Founding Director of the CUHK Ministry of Education (MoE)-Microsoft Key Laboratory for Human-Centric Computing and Interface Technologies (since 2005), Tsinghua-CUHK Joint Research Center for Media Sciences, Technologies and Systems (since 2006), and Co-Director of the Stanley Ho Big Data Decision Analytics Research Center (since 2013). Previously, she served as CUHK Faculty of Engineering's Associate Dean (Research), Chairman of the Department of Systems Engineering and Engineering Management, Editor-in-Chief of the IEEE Transactions on Audio, Speech and Language Processing, Member of the IEEE Signal Processing Society Board of Governors, ISCA Board Member and presently Member of the ISCA International Advisory Council. She was elected APSIPA's inaugural Distinguished Lecturer 2012-2013 and ISCA Distinguished Lecturer 2015-2016. Her awards include the Ministry of Education Higher Education Outstanding Scientific Research Output Award 2009, Hong Kong Computer Society's inaugural Outstanding ICT Woman Professional Award 2015, Microsoft Research Outstanding Collaborator Award 2016 (1 in 32 worldwide), IEEE ICME 2016 Best Paper Award, IBM Faculty Award 2016, HKPWE Outstanding Women Professionals and Entrepreneurs Award 2017 (1 in 20 since 1999), Hong Kong ICT Silver Award 2018 in Smart Inclusion, and the CogInfoComm2018 Best Paper Award. Helen received all her degrees from MIT. Her research interests include big data decision analytics, and artificial intelligence especially for speech and language technologies to support multilingual and multimodal human-computer interaction. Helen has given invited / keynote presentations including INTERSPEECH 2018 Plenary Talk, World Economic Forum Global Future Council 2018, Taihe Workshop on Building Stakeholder Networks on AI Ethics and Governance 2019 and the World Peace Forum 2019. She has served in numerous Government appointments, including

Chairlady of the Research Grants Council's Assessment Panel for Competitive Research Funding Schemes for the Local Self-financing Degree Sector, Chairlady of the Working Party on Manpower Survey of the Information/Innovation Technology Sector (since 2013), as well as Steering Committee Member of Hong Kong's Electronic Health Record (eHR) Sharing. Helen is a Fellow of HKCS, HKIE, IEEE and ISCA.

Deep Reinforcement Learning For Modeling Chit-Chat Dialog With Discrete Attributes

Chinnadhurai Sankar *

Mila, Université de Montréal
chinnadhurai@gmail.com

Sujith Ravi

Google Research
sravi@google.com

Abstract

Open domain dialog systems face the challenge of being repetitive and producing generic responses. In this paper, we demonstrate that by conditioning the response generation on *interpretable* discrete dialog attributes and *composed* attributes, it helps improve the model perplexity and results in diverse and interesting non-redundant responses. We propose to formulate the dialog attribute prediction as a reinforcement learning (RL) problem and use policy gradients methods to optimize utterance generation using long-term rewards. Unlike existing RL approaches which formulate the token prediction as a policy, our method reduces the complexity of the policy optimization by limiting the action space to dialog attributes, thereby making the policy optimization more practical and sample efficient. We demonstrate this with experimental and human evaluations.

1 Introduction

Following the success of neural machine translation systems (Bahdanau et al., 2015; Sutskever et al., 2014; Cho et al., 2014), there has been a growing interest in adapting the encoder-decoder models to model open-domain conversations (Sordani et al., 2015; Serban et al., 2016a,b; Vinyals and Le, 2015). This is done by framing the next utterance generation as a machine translation problem by treating the dialog history as the source sequence and the next utterance as the target sequence. Then the models are trained end-to-end with Maximum Likelihood (MLE) objective without any hand crafted structures like slot-value pairs, dialog manager, etc used in conventional dialog modeling (Lagus and Kuusisto, 2002). Such data driven approaches are worth pursuing in the context of open-domain conversations since the next utterance distribution in open-domain conversations

exhibit high entropy which makes it impractical to manually craft good features.

While the encoder-decoder approaches are promising, lack of specificity has been one of the many challenges (Wei et al., 2017) in modelling non-goal oriented dialogs. Recent encoder-decoder based models usually tend to generate generic or dull responses like “*I don’t know.*”. One of the main causes are the implicit imbalances present in the dialog datasets that tend to potentially handicap the models into generating uninteresting responses.

Imbalances in a dialog dataset can be broadly divided into two categories: *many-to-one* and *one-to-many*. *Many-to-one* imbalance occurs when the dataset contain very similar responses to several different dialog contexts. In such scenarios, decoder learns to ignore the context (considering it as noise) and behaves like a regular language model. Such a decoder would not generalize to new contexts and will end up predicting generic responses for all contexts. In the *one-to-many* case, the dataset may exhibit a different type of imbalance where a certain type of generic response may be present in abundance compared to other plausible interesting responses for the same dialog context (Wei et al., 2017). When trained with a maximum-likelihood (MLE) objective, generative models usually tend to place more probability mass around the most commonly observed responses for a given context. So, we end up observing little variance in the generated responses in such cases. While these two imbalances are problematic for training a dialog model, they are also inherent characteristics of a dialog dataset which cannot be removed.

Several approaches have been proposed in the literature to address the generic response generation issue. Li et al. (2016) propose to modify the loss function to increase the diversity in the generated responses. Multi-resolution RNN (Serban et al., 2017) addresses the above issue by additionally

*Work done during internship at Google

conditioning with entity information in the previous utterances. Alternatively, Song et al. (2016) uses external knowledge from a retrieval model to condition the response generation. Latent variable models inspired by Conditional Variational Autoencoders (CVAEs) are explored in (Shen et al., 2017; Zhao et al., 2017). While models with continuous latent variables tend to be uninterpretable, discrete latent variable models exhibit high variance during inference. Shen et al. (2017) append discrete attributes such as sentiment to the latent representation to generate next utterance.

1.1 Contributions

New Conditional Dialog Generation Model. Drawing insights from (Shen et al., 2017; Zhou et al., 2017), we propose a *conditional utterance generation model* in which the next utterance is conditioned on the dialog attributes corresponding to the next utterance. To do this, we first predict the higher level dialog attributes corresponding to the next response. Then we generate the next utterance conditioned on the dialog context and predicted attributes. Dialog attribute of an utterance refers to discrete features or aspects associated with the utterance. Example attributes include dialog-acts, sentiment, emotion, speaker id, speaker personality or other user defined discrete features of an utterance. While previous research works lack the framework to learn to predict the attributes of the next utterance and mainly view the next utterance’s attribute as a control variable in their models, our method learns to predict the attributes in an end-to-end manner. This alleviates the need to have utterances annotated with attributes during inference.

RL for Dialog Attribute Selection. Further, it also enables us to formulate the dialog attribute selection as a reinforcement learning (RL) problem and optimize the policy initialized by the supervised training using REINFORCE (Williams, 1992). While the Supervised pre-training helps the model to generate utterances coherent with the dialog history, the RL formulation encourages the model to generate utterances optimized for long term rewards like diversity, user-satisfaction scores etc. This way of optimizing the policy over the discrete dialog attribute space is more practical as the action space is low dimensional instead of the entire vocabulary (as common in policies which involve predicting the next token to generate).

By using REINFORCE (Williams, 1992) to further optimize the dialog attribute selection process, We then show improvements in specificity of the generated responses both qualitatively (based on human evaluations) and quantitatively (with respect to the *diversity* measures). The diversity scores, *distinct-1* and *distinct-2* are computed as the fraction of uni-grams and bi-grams in the generated responses as described in (Li et al., 2016).

Improvements on Dialog datasets demonstrated through quantitative & qualitative Evaluations: Additionally, we annotate an existing open domain dialog dataset using dialog attribute classifiers trained with tagged datasets like Switchboard (Godfrey et al., 1992; Jurafsky et al., 1997), Frames (Schulz et al., 2017) and demonstrate both quantitative (in terms of token perplexity/embedding metrics (Rus and Lintean, 2012; Mitchell and Lapata, 2008)) and qualitative improvements (based on human evaluations) in generating interesting responses. In this work, we show results with two types of dialog attributes - sentiment and dialog-acts. It is worth investigating this approach as we need not invest much in training classifiers for very high accuracy and we show empirically that annotations from classifiers with low accuracy are able to boost token perplexity. We conjecture that the irregularities in the auto-annotated dialog attributes induce a regularization effect while training deep neural networks analogous to the dropout mechanism. Also, annotating utterances with many types of dialog attributes could increase the regularization effect and potentially tip the utterance generation in the favor of certain low frequency but interesting responses.

In this work, we are mainly interested in exploring the impact of the jointly modelling extra discrete dialog attributes along with dialog history for next utterance generation and their contribution to addressing the generic response problem. Although our approach is flexible enough to include latent variables additionally, we mainly focus on the contribution of dialog attributes to address the “generic” response issue in this work.

2 Attribute Conditional HRED

In this paper, we extend the *HRED* (Serban et al., 2016a) model (elaborated in the Appendix section) by jointly modelling the utterances with the dialog attributes of each utterance. *HRED* is an encoder-decoder model consisting of a token-level RNN

encoder and an utterance-level RNN encoder to summarize the dialog context followed by a token-level RNN decoder to generate the next utterance. The joint probability can be factorized into dialog attributes prediction, followed by next utterance generation conditioned on the predicted dialog attributes as shown in equation 1 .

$$P(U_m, DA_{1:K} | U_{1:m-1}) = \prod_{i=1}^K P(DA_i | U_{1:m-1}) * P(U_m | U_{1:m-1}, DA_{1:K}) \quad (1)$$

where $DA_{1:K}$ denote K different dialog attributes corresponding to the utterance U_m . U_m is the m_{th} utterance, $U_{1:m-1}$ are the past utterances. For instance, if we condition on three dialog attributes - *sentiment*, *dialog-acts* and *emotion*, we would have $K = 3$. Further, we assume that the dialog attributes are conditionally independent given the dialog context. More simply, we predict the attributes of the next utterance and then, condition on the previous context & the predicted attributes to generate the next utterance.




Figure 1: Dialog attribute classification: We predict the dialog attribute of the next utterance based on the previous context and attributes corresponding to the previous utterances. Please note that we depict only a single attribute for convenience

2.1 Dialog Attribute Prediction

We predict the dialog attribute of the next utterance conditioned on the context vector i.e. summary of the previous utterances and the dialog attributes of the previous utterances. We first pass the attributes of all the previous utterances through an RNN. We combine only the last hidden state of this RNN with the context vector (represents the summary of all the previous utterances) to predict the dialog attribute of the next utterance as shown in Figure 1.

If the dialog dataset is not annotated with the dialog attributes, we build a classifier (with a manually tagged dataset) to annotate the dialog attributes.

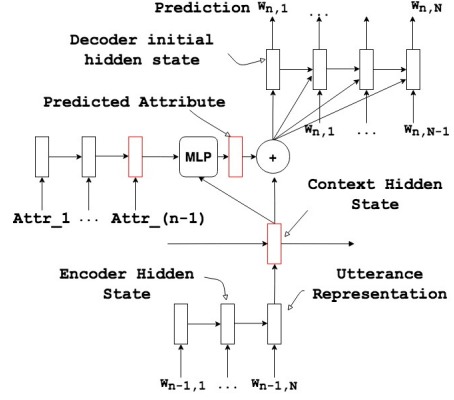


Figure 2: Attribute Conditional HRED : Token generation is additionally conditioned on the predicted dialog attributes. The dialog attribute’s embedding is concatenated with the context vector.

This classifier is a simple MLP. We empirically show that this classifier need not have high accuracy to improve the dialog modeling. We hypothesize that few misclassified attributes could potentially provide a regularization effect similar to the dropout mechanism (Srivastava et al., 2014).

2.2 Conditional Response Generation

After the dialog attributes prediction, we generate the next utterance conditioned on the dialog context and the predicted attributes as shown in Figure 2. Token generation of the next utterance is modelled as in equation 2. The context and attributes are combined by concatenating their corresponding hidden states.

$$h_{dec,m,n} = f_{dec}(h_{dec,m,n-1}, w_{m,n-1}, \mathbf{c}_m) \quad (2)$$

where $h_{dec,m,n}$ is the recurrent hidden state of the decoder after seeing $n - 1$ words in the m -th utterance, f_{dec} is the token level response decoder, and

$$\mathbf{c}_m = [s_{m-1}; da_m^1; da_m^2; \dots; da_m^K] \quad (3)$$

where s_{m-1} is the summary of previous $m - 1$ utterances (recurrent hidden state of the utterance-level encoder), and $da_m^1, da_m^2, \dots, da_m^K$ are the K dialog attribute embeddings corresponding to the m -th utterance.

During inference, we first predict the dialog attributes of the dialog context. We then predict the dialog attribute of the next utterance conditioned on the predicted attribute and the hierarchical utterance representations. We combine the predicted attribute’s embedding vector with the context representation to generate the next utterance. Looking from another perspective, we could formulate

the conditional utterance generation problem as a multi-task problem where we jointly learn to predict the dialog attributes and tokens of the next utterance.

2.3 RL for Dialog Attribute Prediction

Often the MLE objective does not capture the true goal of the conversation and lacks the framework which can take developer-defined rewards into account for modelling such goals. Also, the MLE-based seq2seq models fail to model long term influence of the utterances on the dialog flow causing coherency issues. This calls for a Reinforcement Learning (RL) based framework which has the ability to optimize policies for maximizing long term rewards. At the core, the MLE objective tries to increase the conditional utterance probabilities and influences the model to place higher probabilities over the commonly occurring utterances. On the other hand, RL based methods circumvent this issue by shifting the optimization problem to maximizing long term rewards which could promote diversity, coherency, etc.

Previous approaches [Li et al. \(2016\)](#); [Kottur et al. \(2017\)](#); [Lewis et al. \(2017\)](#) propose to model the token prediction of the next utterance as a reinforcement learning problem and optimize the models to maximize hand-crafted rewards for improving diversity, coherency, and ease of answering. Their approaches involves pre-training the encoder-decoder models with supervised training and then refining the utterance generation further with RL using the hand-engineered rewards. Their state space consists of the dialog context representation (encoder hidden states). Their action space at a given time step includes all possible words that the decoder can generate (which is very large).

While this approach is appealing, policy gradient methods are known to suffer from high variance when using large action spaces. This makes training extremely unstable and requires significant engineering efforts to train successfully.

Another potential drawback with directly acting over the vocabulary space is that the RL optimization procedure tends to strip away the linguistic / natural language aspects learned during the supervised pre-training step, as observed in ([Kottur et al., 2017](#); [Lewis et al., 2017](#)). Since the primary focus of the RL objective function is to improve the final reward (which may not emphasize on the linguistic aspects of the generated responses, for

e.g., diversity scores), the optimization algorithm could lead the decoder into generating unnatural responses. We propose to avoid both the issues by reducing the action space to a higher level abstraction space i.e. the dialog attributes. Our action space comprises the discrete dialog attributes and the state space is the dialog context. Intuitively, this enables the RL policy to view the dialog attributes as control variables for improving dialog flow and modelling long term influence. For instance, if the input response was “*how old are you?*”, an RL policy optimized to maximize conversation length and engagement could choose to set one of the next utterance attributes as a question-type to generate a response like “*why do you ask?*” instead of a straightforward answer, to keep the conversation engaging. Thus, we believe that this approach enables the model to predict such rare but interesting utterances to which the MLE objective fails to give attention.

Our policy network comprises of the encoders and the attribute prediction network. Given the previous utterances $U_{1:m-1}$, the policy network first encodes them by using the encoders. Then this encoded representation is passed to the attribute prediction network. The output of the attribute prediction network is the action. While there are many ways to design the reward function, we adopt the *ease-of-answering* reward introduced by [Li et al. \(2016\)](#) - negative log-likelihood of a set of manually constructed dull utterances (usually the most commonly occurring phrases in the dataset) in response to the next generated utterance. Let \mathbb{S} be the set of dull utterances. With the sampled dialog-acts, $DA_{1:K}$ from the policy network, we generate the next utterance U_m using the decoder. Then we add this generated utterance to the context and predict the probability of seeing one of the dull utterances in the $m + 1$ -th step. This is used to compute the reward as follows:

$$R = \frac{1}{|\mathbb{S}|} \sum_{s \in \mathbb{S}} \frac{1}{N_s} \log P(s|U_{1:m}), \quad (4)$$

where N_s is the number of tokens in the dull utterance s . The normalization avoids the reward function attending to only the longer dull responses. We use REINFORCE ([Williams, 1992](#)) to optimize our policy, $P_{RL}(DA_{1:K}|U_{1:m-1})$. The expected reward is given by equation 5.

$$J(\theta) = \mathbb{E}[R(U_{1:m-1}, DA_{1:K})] \quad (5)$$

The gradient is estimated as in equation 6.

$$\nabla J(\theta_{RL}) = (R - b) \nabla \log P_{RL}(DA_{1:K} | U_{1:m-1}), \quad (6)$$

where b is the reward baseline (computed as the running average of the rewards during training). We initialize the policy with the supervised training and add an L2-loss to penalize the network weights from moving away from the supervised network weights.

3 Training Setup

Datasets: We first start with the Reddit-discourse dataset (Zhang et al., 2017) for training dialog attribute classifiers and modelling utterance generation.

Reddit: The Reddit discourse dataset (Zhang et al., 2017) is manually pre-annotated with dialog-acts via crowd sourcing. The dialog-acts comprise of *answer, question, humor, agreement, disagreement, appreciation, negative reaction, elaboration, announcement*. It comprises conversations from around 9000 randomly sampled Reddit threads with over 100000 comments and an average of 12 turns per thread.

Open-Subtitles: Additionally, we show results with the unannotated Open-Subtitles dataset (Tiedemann, 2009) (we randomly sample up to 2 million dialogs for training and validation). We tag the dataset with dialog attributes using pre-trained classifiers.

We experiment with two types of dialog attributes in this paper - *sentiment and dialog-acts*. We annotate the utterances with sentiment tags - *positive, negative, neutral* using the Stanford CoreNLP tool (Manning et al., 2014). We adopt the dialog-acts from two annotated dialog corpus - Switchboard (Godfrey et al., 1992) and Frames (Schulz et al., 2017).

Switchboard: Switchboard corpus (Godfrey et al., 1992) is a collection of 1155 chit-chat style telephonic conversations based on 70 topics. Jurafsky et al. (1997) revised the original tags to 42 dialog-acts. In our experiments, we restrict dialog-acts to the top-10 most frequently annotated tags in the corpus - *Statement-non-opinion, Acknowledge, Statement-opinion, Agree/Accept, Abandoned or Turn-Exit, Appreciation, Yes-No-Question, Non-verbal, Yes answers, Conventional-closing*. We consider the top-10 frequently annotated tags as a simple solution to avoid the class imbalance issue (the *Statement-non-opinion* act is tagged 72824

times, while *Thanking* is tagged only 67 times) for training the dialog attribute classifiers.

Frames: Frames (Schulz et al., 2017) is a task oriented dialog corpus collected in the *Wizard-of-Oz* fashion. It comprises of 1369 human-human dialogues with an average of 15 turns per dialog. The wizards had access to a database of hotels and flights information and had to converse with users to help finalize vacation plans. The dataset has 20 different types of dialog-acts annotations. Like the Switchboard corpus, we adopt the top 10 frequently occurring acts in the dataset for our experiments - *inform, offer, request, suggest, switch-frame, no result, thank you, sorry, greeting, affirm*.

Model Details: We use two-layer GRUs (Chung et al., 2014) for both encoder and decoders with hidden sizes of 512. We restrict the vocabulary for both the datasets to top 25000 frequency occurring tokens. The dialog attribute classifier for dialog attributes is a simple 2-layer MLP with layer sizes of 256, and 10 respectively. We use the rectified linear unit (ReLU) as the non-linear activation function for the MLPs and use dropout rate of 0.3 for the token embeddings, hidden-hidden transition matrices of the encoder and decoder GRUs.

Training Details: We ran our experiments in Nvidia Tesla-K80 GPUs and optimized using the ADAM optimizer with the default hyperparameters used in (Merity et al., 2017, 2018). All models are trained with batch size 128 and a learning rate 0.0001.

4 Experimental Results

In this section, we present the experimental results along with qualitative analysis.

In Section 4.1, we discuss the dialog attribute classification results for different model architectures trained on the Reddit, Switchboard and Frames datasets.

In Section 4.2, we first demonstrate quantitative improvements (token perplexity/embedding based metrics) for the Attribute conditional HRED model with the manually annotated Reddit dataset. Further, we discuss the model perplexity improvements along with sample conversations and human evaluation results on the Open-Subtitles dataset. We annotate it with sentiment and dialog-acts (from Switchboard/Frames datasets) using pre-trained classifiers described in Section 4.1.

Finally, in Section 4.3, we analyze the quality of the generated responses after RL fine-tuning us-

ing *diversity* scores (*distinct-1*, *distinct-2*), sample conversations and human evaluation results for diversity and relevance.

4.1 Dialog Attribute Prediction

In this section, we present the experiments with the model architectures for the dialog attribute prediction - dialog-acts from Reddit, Switchboard and Frames datasets. First, we demonstrate the performance of the dialog-acts classifiers on the Reddit dataset as shown in Table 1.

Model	Acc(%)
$F(U_t)$	57
$F(DA_{t-1,t-2})$	54
$F(U_t, DA_{t-1,t-2})$	68

Table 1: Dialog-acts prediction accuracy in Reddit validation set.

The model $F(U_t)$ refers to the architecture which predicts the dialog-acts based on current utterance U_t alone. The tokens in the current utterance U_t are fed through a two-layer GRU and the final hidden state is used to predict the dialog-acts. The model $F(DA_{t-1,t-2})$ predicts the current utterance’s dialog-acts DA_t based on the dialog-acts corresponding to the previous two utterances. We consider the dialog-acts prediction problem as a sequence modelling problem where we feed the dialog-acts into a single-layer GRU and predict the current dialog-acts conditioned on the previous dialog-acts. We settled on conditioning on the dialog-acts corresponding to the previous two utterances alone as we didn’t observe any boost in the classifier performance from the older dialog-acts. As seen in Table 1, conditioning additionally on the dialog attributes helps improve classifier performance.

Next, we train classifiers to predict dialog-acts of utterances of the Switchboard and Frames corpus. In our experiments, the number of act types is 11 - the top 10 most frequently occurring acts in the corpus and "others" category covering the rest of the tags.

As seen from Table 2, classifier performance is not really high and yet, contribute to improvements in perplexity for the conditional Seq2Seq models (discussed in Section 4.2). While we aim for better classifier performance, it is important to note here that the primary objective of such dialog attribute classifiers is to tag unannotated open-domain dia-

Corpus	Num Acts	Acc(%)
Reddit	9	68.1
Switchboard	11	67.9
Frames	11	71.1

Table 2: Dialog-acts prediction accuracy for classifiers trained on validation set of different datasets.

log datasets. As future work, we will study how the classification errors influence response generation.

4.2 Utterance Evaluation

Following (Serban et al., 2016a), we use token perplexity and embedding based metrics (average, greedy and extrema) (Mitchell and Lapata, 2008; Rus and Lintean, 2012) for utterance evaluation.

Metric	LM	Seq2Seq	Seq2Seq+Attr
Perplexity	176	170	163
Greedy	-	0.47	0.54
Extrema	-	0.37	0.47
Average	-	0.67	0.62

Table 3: Perplexity and Embedding Metrics for the Reddit validation set.

Reddit: First, we evaluate Seq2Seq models trained on the manually annotated Reddit corpus as shown in Table 3. *Seq2Seq+Attr* refers to our model where we condition on the dialog-acts additionally. Please note that we use the notation "Attr" here to maintain generality as it may refer to other dialog attributes like sentiment later in this section. For both the baseline and conditional Seq2Seq models, we consider a dialog context involving the previous two turns as we did not observe significant performance improvement with three or more turns. We use a 2-layer GRU language model as a baseline for comparison. As seen from Table 3, *Seq2Seq+Attr* fares well both in terms of perplexity and embedding metrics. Higher perplexity observed in the Reddit corpus could be due to the presence of several topics in the dataset (exhibits high entropy) and fewer dialogs compared to other open domain dialog datasets.

Open-Subtitles: With promising results on the manually tagged Reddit corpus, we now evaluate our attribute conditional HRED model on the unannotated Open-Subtitles dataset. We tag the Open-Subtitles dataset with the sentiment tags using the Stanford Core-NLP tool (Manning et al., 2014) and

Model	Attributes	Num Dialogs(in Millions)			
		0.2 M	0.5 M	1 M	2 M
Seq2seq	-	101.63	80.05	74.78	67.28
Seq2seq	Sentiment	98.61	79.15	72.23	66.11
Seq2seq	Switchboard	97.03	77.81	71.51	64.21
Seq2seq	Frames	96.61	77.41	72.01	65.33
Seq2seq	Sentiment, Switchboard	96.67	78.01	72.17	66.01
Seq2seq	Sentiment, Frames	96.32	77.61	72.15	66.13
Seq2seq	Switchboard, Frames	94.80	77.40	71.18	65.01

Table 4: Validation Perplexity for the Open-Subtitles dataset.

dialog-acts from Frames & Switchboard corpus using the pre-trained classifiers described in Section 4.1.

In Table 4, we compare the model perplexity when trained on varying dialog corpus size. In most of the cases, we observe that the conditioning with acts from both the frames and switchboard yields the lowest perplexity. We observe that the perplexity improvement is substantial for smaller datasets which is also corroborated from the experiments with the Reddit dataset.

Human Evaluation: Following the human evaluation setting in (Li et al., 2016), we randomly sample 200 input message and the generated outputs from the *Seq2Seq+Attr* & *Seq2Seq* models. We present each of them to 3 judges and ask them to decide which of the two outputs is 1) relevant and 2) diverse or interesting. Ties are permitted. Results for human evaluation are shown in Table 8. We observe that *Seq2Seq+Attr* performs better than the *Seq2Seq* model both in terms of *diversity* and *relevance*.

Seq2Seq+Attr vs Seq2Seq			
Metric	Wins(%)	Losses(%)	Ties(%)
<i>Diversity</i>	42	24.16	33.84
<i>Relevance</i>	40.16	36.83	23.01

Table 5: Human Evaluation results: *Seq2Seq+Attr* vs *Seq2Seq*

Please note that the *Seq2Seq+Attr* model performs better in terms of diversity compared to the relevancy. This is in line with our expectations, as the purpose of dialog attribute annotations is to help the model focus better on less-frequent responses.

Additionally, we present a few sample conversations in Table 6, where we observe that the *Seq2Seq+Attr* model generates more interesting responses.

Input:	i wish i was home watching tv.
Seq2Seq:	i dont know what i was thinking about
Seq2Seq+Attr:	i cant wait to see it.
Input:	He used from his charity to settle legal problems.
Seq2Seq:	i have no idea what youre talking about
Seq2Seq+Attr:	i dont think he is going to be a president.
Input:	tell us how you really feel
Seq2Seq:	i dont understand why
Seq2Seq+Attr:	lmao i could hella picture your reaction

Table 6: Sample conversations

4.3 RL For Dialog Attribute Prediction

For the RL fine-tuning, we report the diversity scores of the generated responses with the models trained on the Open-Subtitles dataset in Table 7. The diversity scores, *distinct-1* and *distinct-2* are computed as the fraction of uni-grams and bi-grams in the generated responses following the previous work by Li et al. (2015).

Model	<i>distinct-1</i>	<i>distinct-2</i>
Seq2Seq	0.004	0.013
Seq2Seq+Attr	0.005	0.018
RL	0.011	0.033

Table 7: Diversity scores on the Open-Subtitles validation set after RL fine-tuning .

We use the model conditioned on acts from both Switchboard and Frames for the *Seq2Seq+Attr* and *RL* cases. The action space for the policy in this case, covers the 10 acts from Switchboard and Frames each. We choose a collection of commonly occurring phrases in the Open-Subtitles dataset as the set of dull responses, \mathbb{S} for the reward computation in equation 4. We observe that the RL fine-tuning improves over the conditional seq2seq in terms of the diversity scores.

Human Evaluation: As described in Section 4.2, we present each of the 200 randomly sampled input-response pairs of the *Seq2Seq + Attr* and *RL* models to 3 judges and ask to them rate each sample for *diversity* and *relevance*. From Table 8, we can see that the *RL* model significantly performs better both in terms of *diversity* and *relevance*.

Qualitative Analysis: In Table 9, we present the percentage of the commonly occurring generic

RL vs Seq2Seq+Attr			
Metric	Wins(%)	Losses(%)	Ties(%)
<i>Diversity</i>	54.66	28.50	16.84
<i>Relevance</i>	43.33	26.62	30.05

Table 8: Human Evaluation results: *RL* vs *Seq2Seq+Attr*

responses from the Open-Subtitles dataset in the validation set samples corresponding to the *RL* and *Seq2Seq + Attr* models. We observe very low percentages of such generic responses in the samples after RL fine-tuning. It is interesting to note that RL model has successfully learned to minimize the generation of other dull responses like *I would love to be*, *I would love to see*, *I dont want to* apart from expected the dull responses, \mathbb{S} (used in the reward computation). At the same time, RL model has scored higher in terms of the *Relevancy* metric, as seen in Table 8 which indicates that the RL fine-tuning actually explores interesting responses whilst avoiding the generic responses.

Generic Responses	RL(%)	Seq2Seq + Attr(%)
thank you so much	7.56	7.32
i dont understand why	0.0	15.64
i would love to see	0.66	5.65
i dont know how	0.0	13.97
i dont want to	1.66	3.99
i dont know why	0.0	3.66
i would love to be	0.99	2.21
i have no idea	4.31	3.33

Table 9: Percentage of generic responses after RL fine-tuning.

Additionally, we present a few sample conversations in Table 10, where we observe that the *RL* model generates more diverse and relevant responses.

5 Related Work

There are several works focusing on dialog-acts classification and clustering based analysis (Reithinger and Klesen, 1997; Liu, 2006; Khanpour et al., 2016; Ang et al., 2005; Crook et al., 2009; Stolcke et al., 2000; Ezen-Can and Boyer, 2013). Shen et al. (2017) additionally add sentiment feature to the latent variables in the VAE setting for utterance generation. In our work, we use dialog attributes from different sources - Switchboard and Frames corpus to model utterance generation in a more realistic setting. As for the RL setting, existing research efforts include (Li et al., 2016;

Input:	i'm honestly a bit confused why no one has brought me or my books any cake
Seq2Seq+Attr:	i dont think i dont think anything
RL:	i cant wait to see you in the city.
Input:	ive been playing spaceship with my year old niece for the past few days
Seq2Seq+Attr:	i dont even know what i was talking about.
RL:	this is the best thing ive ever seen.
Input:	it makes me so happy that you like them
Seq2Seq+Attr:	i dont know what i was thinking about it
RL:	i was just thinking about the same thing

Table 10: Sample conversations

Dhingra et al., 2016; Jaques et al., 2016) which formulate the token prediction as a RL policy in Seq2Seq models. However, searching over a huge vocabulary space typically involves training with huge number of samples and careful fine-tuning of the policy optimization algorithms. Additionally, as discussed in Section 2.3, it requires precautionary measures to prevent the RL algorithm from removing the linguistic aspects of the generated utterances. In another related research work, Serban et al. (2017) use dialog-acts as one among their hand crafted features to select responses from an ensemble of dialog systems. They use dialog-acts in their RL policy, however their action space comprises of responses from an ensemble of dialog models. They include dialog-acts in their features for their distributed state representation.

6 Conclusion

In this work, we address the dialog utterance generation problem by jointly modeling previous dialog context and discrete dialog attributes. We analyze both quantitatively (model perplexity and other embedding based metrics) and qualitatively (human evaluation, sample conversations) to validate that *composed* dialog attributes help generate interesting responses. Further, we formulate the dialog attribute prediction problem as a reinforcement learning problem. We fine tune the attribute selection policy network trained with supervised learning using REINFORCE and demonstrate improvements in diversity scores compared to the Seq2Seq model. In the future, we plan to extend the model for additional dialog attributes like emotion, speaker persona etc. and evaluate the controllability aspect of the responses based on the dialog attributes.

References

- Jeremy Ang, Yang Liu, and Elizabeth Shriberg. 2005. Automatic dialog act segmentation and classification in multiparty meetings. In *ICASSP (1)*, pages 1061–1064.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings Of The International Conference on Representation Learning (ICLR 2015)*.
- K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. 2014. [Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation](#). *ArXiv e-prints*.
- J. Chung, C. Gulcehre, K. Cho, and Y. Bengio. 2014. [Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling](#). *ArXiv e-prints*.
- Nigel Crook, Ramón Granell, and Stephen G. Pulman. 2009. Unsupervised classification of dialogue acts using a dirichlet process mixture model. In *Proceedings of the SIGDIAL 2009 Conference, The 10th Annual Meeting of the Special Interest Group on Discourse and Dialogue, 11-12 September 2009, London, UK*, pages 341–348.
- B. Dhingra, L. Li, X. Li, J. Gao, Y.-N. Chen, F. Ahmed, and L. Deng. 2016. [Towards End-to-End Reinforcement Learning of Dialogue Agents for Information Access](#). *ArXiv e-prints*.
- Aysu Ezen-Can and Kristy Elizabeth Boyer. 2013. Unsupervised classification of student dialogue acts with query-likelihood clustering. In *Proceedings of the 6th International Conference on Educational Data Mining, Memphis, Tennessee, USA, July 6-9, 2013*, pages 20–27.
- John J. Godfrey, Edward C. Holliman, and Jane McDaniel. 1992. [Switchboard: Telephone speech corpus for research and development](#). In *Proceedings of the 1992 IEEE International Conference on Acoustics, Speech and Signal Processing - Volume 1, ICASSP'92*, pages 517–520, Washington, DC, USA. IEEE Computer Society.
- N. Jaques, S. Gu, D. Bahdanau, J. M. Hernández-Lobato, R. E. Turner, and D. Eck. 2016. [Sequence Tutor: Conservative Fine-Tuning of Sequence Generation Models with KL-control](#). *ArXiv e-prints*.
- D. Jurafsky, R. Bates, N. Coccaro, R. Martin, M. Meteer, K. Ries, E. Shriberg, A. Stolcke, P. Taylor, and C. Van Ess-Dykema. 1997. [Automatic detection of discourse structure for speech recognition and understanding](#). In *1997 IEEE Workshop on Automatic Speech Recognition and Understanding Proceedings*, pages 88–95.
- Hamed Khanpour, Nishitha Guntakandla, and Rodney D. Nielsen. 2016. Dialogue act classification in domain-independent conversations using a deep recurrent neural network. In *COLING 2016, 26th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers, December 11-16, 2016, Osaka, Japan*, pages 2012–2021.
- S. Kottur, J. M. F. Moura, S. Lee, and D. Batra. 2017. [Natural Language Does Not Emerge ‘Naturally’ in Multi-Agent Dialog](#). *ArXiv e-prints*.
- Krista Lagus and Jukka Kuusisto. 2002. Topic identification in natural language dialogues using neural networks. In *Proceedings of the SIGDIAL 2002 Workshop, The 3rd Annual Meeting of the Special Interest Group on Discourse and Dialogue, Thursday, July 11, 2002 to Friday, July 12, 2002, Philadelphia, PA, USA*, pages 95–102.
- M. Lewis, D. Yarats, Y. N. Dauphin, D. Parikh, and D. Batra. 2017. [Deal or No Deal? End-to-End Learning for Negotiation Dialogues](#). *ArXiv e-prints*.
- J. Li, M. Galley, C. Brockett, J. Gao, and B. Dolan. 2015. [A Diversity-Promoting Objective Function for Neural Conversation Models](#). *ArXiv e-prints*.
- J. Li, M. Galley, C. Brockett, J. Gao, and B. Dolan. 2016. A diversity-promoting objective function for neural conversation models. In *The North American Chapter of the Association for Computational Linguistics (NAACL)*, pages 110–119.
- J. Li, W. Monroe, A. Ritter, M. Galley, J. Gao, and D. Jurafsky. 2016. [Deep Reinforcement Learning for Dialogue Generation](#). *ArXiv e-prints*.
- Yang Liu. 2006. Using SVM and error-correcting codes for multiclass dialog act classification in meeting corpus. In *INTERSPEECH 2006 - ICSLP, Ninth International Conference on Spoken Language Processing, Pittsburgh, PA, USA, September 17-21, 2006*.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. [The Stanford CoreNLP natural language processing toolkit](#). In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60.
- S. Merity, N. Shirish Keskar, and R. Socher. 2017. [Regularizing and Optimizing LSTM Language Models](#). *ArXiv e-prints*.
- S. Merity, N. Shirish Keskar, and R. Socher. 2018. [An Analysis of Neural Language Modeling at Multiple Scales](#). *ArXiv e-prints*.
- Jeff Mitchell and Mirella Lapata. 2008. Vector-based models of semantic composition. In *ACL*, pages 236–244.
- Norbert Reithinger and Martin Klesen. 1997. Dialogue act classification using language models. In *EuroSpeech*.

- Vasile Rus and Mihai Lintean. 2012. A comparison of greedy and optimal assessment of natural language student input using word-to-word similarity metrics. In *Proceedings of the Seventh Workshop on Building Educational Applications Using NLP*, pages 157–162. Association for Computational Linguistics.
- Hannes Schulz, Jeremie Zumer, Layla El Asri, and Shikhar Sharma. 2017. [A frame tracking model for memory-enhanced dialogue systems](#). *CoRR*, abs/1706.01690.
- I. V. Serban, C. Sankar, M. Germain, S. Zhang, Z. Lin, S. Subramanian, T. Kim, M. Pieper, S. Chandar, N. R. Ke, S. Rajeshwar, A. de Brebisson, J. M. R. Sotelo, D. Suhubdy, V. Michalski, A. Nguyen, J. Pineau, and Y. Bengio. 2017. [A Deep Reinforcement Learning Chatbot](#). *ArXiv e-prints*.
- Iulian Vlad Serban, Tim Klinger, Gerald Tesauro, Kartik Talamadupula, Bowen Zhou, Yoshua Bengio, and Aaron C. Courville. 2017. Multiresolution recurrent neural networks: An application to dialogue response generation. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 3288–3294.
- Iulian Vlad Serban, Alessandro Sordoni, Yoshua Bengio, Aaron C. Courville, and Joelle Pineau. 2016a. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proceedings of AAAI*.
- Iulian Vlad Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron C. Courville, and Yoshua Bengio. 2016b. A hierarchical latent variable encoder-decoder model for generating dialogues. *CoRR*, abs/1605.06069.
- X. Shen, H. Su, Y. Li, W. Li, S. Niu, Y. Zhao, A. Aizawa, and G. Long. 2017. [A Conditional Variational Framework for Dialog Generation](#). *ArXiv e-prints*.
- Y. Song, R. Yan, X. Li, D. Zhao, and M. Zhang. 2016. [Two are Better than One: An Ensemble of Retrieval and Generation-Based Dialog Systems](#). *ArXiv e-prints*.
- Alessandro Sordoni, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and Bill Dolan. 2015. A neural network approach to context-sensitive generation of conversational responses. *arXiv preprint arXiv:1506.06714*.
- Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of machine learning research*, 15(1):1929–1958.
- Andreas Stolcke, Klaus Ries, Noah Coccaro, Elizabeth Shriberg, Rebecca Bates, Daniel Jurafsky, Paul Taylor, Rachel Martin, Carol Van Ess-Dykema, and Marie Meteer. 2000. Dialogue act modeling for automatic tagging and recognition of conversational speech. *Computational linguistics*, 26(3):339–373.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 3104–3112.
- Jörg Tiedemann. 2009. News from OPUS - A collection of multilingual parallel corpora with tools and interfaces. In N. Nicolov, K. Bontcheva, G. Angelova, and R. Mitkov, editors, *Recent Advances in Natural Language Processing*, volume V, pages 237–248. John Benjamins, Amsterdam/Philadelphia, Borovets, Bulgaria.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- B. Wei, S. Lu, L. Mou, H. Zhou, P. Poupart, G. Li, and Z. Jin. 2017. [Why Do Neural Dialog Systems Generate Short and Meaningless Replies? A Comparison between Dialog and Translation](#). *ArXiv e-prints*.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*.
- Amy X. Zhang, Bryan Culbertson, and Praveen Paritosh. 2017. Characterizing online discussion using coarse discourse sequences. In *Proceedings of the 11th International AAAI Conference on Weblogs and Social Media, ICWSM '17*.
- T. Zhao, R. Zhao, and M. Eskenazi. 2017. [Learning Discourse-level Diversity for Neural Dialog Models using Conditional Variational Autoencoders](#). *ArXiv e-prints*.
- H. Zhou, M. Huang, T. Zhang, X. Zhu, and B. Liu. 2017. [Emotional Chatting Machine: Emotional Conversation Generation with Internal and External Memory](#). *ArXiv e-prints*.

Improving Interaction Quality Estimation with BiLSTMs and the Impact on Dialogue Policy Learning

Stefan Ultes

Daimler AG

Sindelfingen, Germany

stefan.ultes@daimler.com

Abstract

Learning suitable and well-performing dialogue behaviour in statistical spoken dialogue systems has been in the focus of research for many years. While most work which is based on reinforcement learning employs an objective measure like task success for modelling the reward signal, we use a reward based on user satisfaction estimation. We propose a novel estimator and show that it outperforms all previous estimators while learning temporal dependencies implicitly. Furthermore, we apply this novel user satisfaction estimation model live in simulated experiments where the satisfaction estimation model is trained on one domain and applied in many other domains which cover a similar task. We show that applying this model results in higher estimated satisfaction, similar task success rates and a higher robustness to noise.

1 Introduction

One prominent way of modelling the decision-making component of a spoken dialogue system (SDS) is to use (partially observable) Markov decision processes ((PO)MDPs) (Lemon and Pietquin, 2012; Young et al., 2013). There, reinforcement learning (RL) (Sutton and Barto, 1998) is applied to find the optimal system behaviour represented by the policy π . Task-oriented dialogue systems model the reward r , used to guide the learning process, traditionally with task success as the principal reward component (Gašić and Young, 2014; Lemon and Pietquin, 2007; Daubigney et al., 2012; Levin and Pieraccini, 1997; Young et al., 2013; Su et al., 2015, 2016).

An alternative approach proposes user satisfaction as the main reward component (Ultes et al., 2017a). However, the applied statistical user satisfaction estimator heavily relies on handcrafted temporal features. Furthermore, the impact of the

estimation performance on the resulting dialogue policy remains unclear.

In this work, we propose a novel LSTM-based user satisfaction reward estimator that is able to learn the temporal dependencies implicitly and compare the performance of the resulting dialogue policy with the initially used estimator.

Optimising the dialogue behaviour to increase user satisfaction instead of task success has multiple advantages:

1. The user satisfaction is more domain-independent as it can be linked to interaction phenomena independent of the underlying task (Ultes et al., 2017a).
2. User satisfaction is favourable over task success as it represents more accurately the user’s view and thus whether the user is likely to use the system again in the future. Task success has only been used as it has shown to correlate well with user satisfaction (Williams and Young, 2004).

Based on previous work by Ultes et al. (2017a), the interaction quality (IQ)—a less subjective version of user satisfaction¹—will be used for estimating the reward. The estimation model is thus based on domain-independent, interaction-related features which do not have any information available about the goal of the dialogue. This allows the reward estimator to be applicable for learning in unseen domains.

The originally applied IQ estimator heavily relies on handcrafted temporal features. In this work, we will present a deep learning-based IQ estimator that utilises the capabilities of recurrent neural networks to get rid of all handcrafted fea-

¹The relation of US and IQ has been closely investigated in (Schmitt and Ultes, 2015; Ultes et al., 2013).

tures that encode temporal effects. By that, these temporal dependencies may be learned instead.

The applied RL framework is shown in Figure 1. Within this setup, both IQ estimators are used for learning dialogue policies in several domains to analyse their impact on general dialogue performance metrics.

The remainder of the paper is organised as follows: in Section 2, related work is presented focusing on dialogue learning and the type of reward that is applied. In Section 3, the interaction quality is presented and how it is used in the reward model. The deep learning-based interaction quality estimator proposed in this work is then described in detail in Section 4 followed by the experiments and results both of the estimator itself and the resulting dialogue policies in Section 5.

2 Relevant Related Work

Most of previous work on dialogue policy learning focuses on employing task success as the main reward signal (Gašić and Young, 2014; Gašić et al., 2014; Lemon and Pietquin, 2007; Daubigney et al., 2012; Levin and Pieraccini, 1997; Young et al., 2013; Su et al., 2015, 2016). However, task success is usually only computable for pre-defined tasks e.g., through interactions with simulated or recruited users, where the underlying goal is known in advance. To overcome this, the required information can be requested directly from users at the end of each dialogue (Gašić et al., 2013). However, this can be intrusive, and users may not always cooperate.

An alternative is to use a task success estimator (El Asri et al., 2014b; Su et al., 2015, 2016). With the right choice of features, these can also be applied to new and unseen domains (Vandyke et al., 2015). However, these models still attempt to estimate completion of the underlying task, whereas our model evaluates the overall user experience.

In this paper, we show that an interaction quality reward estimator trained on dialogues from a bus information system will result in well-performing dialogues both in terms of success rate and user satisfaction on five other domains, while only using interaction-related, domain-independent information, i.e., not knowing anything about the task of the domain.

Others have previously introduced user satisfaction into the reward (Walker et al., 1998;

Walker, 2000; Rieser and Lemon, 2008b,a) by using the PARADISE framework (Walker et al., 1997). However, PARADISE relies on the existence of explicit task success information which is usually hard to obtain.

Furthermore, to derive user ratings within that framework, users have to answer a questionnaire which is usually not feasible in real world settings. To overcome this, PARADISE has been used in conjunction with expert judges instead (El Asri et al., 2012, 2013) to enable unintrusive acquisition of dialogues. However, the problem of mapping the results of the questionnaire to a scalar reward value still exists.

Therefore, we use interaction quality (Section 3) in this work because it uses scalar values applied by experts and only uses task-independent features that are easy to derive.

3 Interaction Quality Reward Estimation

In this work, the reward estimator is based on the interaction quality (IQ) (Schmitt and Ultes, 2015) for learning information-seeking dialogue policies. IQ represents a less subjective variant of user satisfaction: instead of being acquired from users directly, experts annotate pre-recorded dialogues to avoid the large variance that is often encountered when users rate their dialogues directly (Schmitt and Ultes, 2015).

IQ is defined on a five-point scale from five (satisfied) down to one (extremely unsatisfied). To derive a reward from this value, the equation

$$R_{IQ} = T \cdot (-1) + (iq - 1) \cdot 5 \quad (1)$$

is used where R_{IQ} describes the final reward. It is applied to the final turn of the dialogue of length T with a final IQ value of iq . A per-turn penalty of -1 is added to the dialogue outcome. This results in a reward range of 19 down to $-T$ which is consistent with related work (Gašić and Young, 2014; Vandyke et al., 2015; Su et al., 2016, e.g.) in which binary task success (TS) was used to define the reward as:

$$R_{TS} = T \cdot (-1) + \mathbb{1}_{TS} \cdot 20, \quad (2)$$

where $\mathbb{1}_{TS} = 1$ only if the dialogue was successful, $\mathbb{1}_{TS} = 0$ otherwise. R_{TS} will be used as a baseline.

The problem of estimating IQ has been cast as a classification problem where the target classes

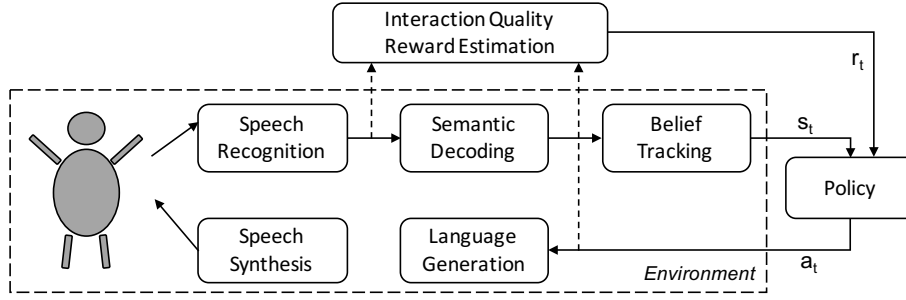


Figure 1: The RL framework integrating an interaction quality reward estimator as proposed by [Ultes et al. \(2017a\)](#). The policy learns to take action a_t at time t while being in state s_t and receiving reward r_t .




Figure 2: Modelling of temporal information in the interaction parameters used as input to the IQ estimator.

are the distinct IQ values. The input consists of domain-independent variables called interaction parameters. These parameters incorporate information from the automatic speech recognition (ASR) output and the preceding system action. Most previous approaches used this information, which is available at every turn, to compute temporal features by taking sums, means or counts from the turn-based information for a window of the last 3 system-user-exchanges² and the complete dialogue (see Fig. 2). The baseline IQ estimation approach as applied by [Ultes et al. \(2017a\)](#) (and originating from [Ultes et al. \(2015\)](#)) used a feature set of 16 parameters as shown in Table 1 with a support vector machine (SVM) ([Vapnik, 1995](#); [Chang and Lin, 2011](#)).

The LEGO corpus ([Schmitt et al., 2012](#)) provides data for training and testing and consists of 200 dialogues (4,885 turns) from the Let’s Go bus information system ([Raux et al., 2006](#)). There, users with real needs were able to call the system to get information about the bus schedule. Each turn of these 200 dialogues has been annotated with IQ (representing the quality of the dialogue up to the current turn) by three experts. The final IQ label has been assigned using the median of the three individual labels.

Previous work has used the LEGO corpus with

²a system turn followed by a user turn

Table 1: The parameters used for IQ estimation extracted on the exchange level from each user input plus counts, sums and rates for the whole dialogue (#,% ,Mean) and for a window of the last 3 turns ({.}).

	Parameter	Description
Exchange level	ASRRognitionStatus	ASR status: <i>success, no match, no input</i>
	ASRConfidence	confidence of top ASR results
	RePrompt?	is the system question the same as in the previous turn?
	ActivityType	general type of system action: <i>statement, question</i>
	Confirmation?	is system action confirm?
Dialogue level	MeanASRConfidence	mean ASR confidence if ASR is success
	#Exchanges	number of exchanges (turns)
	#ASRSuccess	count of ASR status is success
	%ASRSuccess	rate of ASR status is success
	#ASRRejections	count of ASR status is reject
%ASRRejections	rate of ASR status is reject	
Window level	{Mean}ASRConfidence	mean ASR confidence if ASR is success
	{#}ASRSuccess	count of ASR is success
	{#}ASRRejections	count of ASR status is reject
	{#}RePrompts	count of times RePrompt? is true
	{#}SystemQuestions	count of ActivityType is question

a full IQ feature set (which includes additional partly domain-related information) achieving an unweighted average recall³ (UAR) of 0.55 using ordinal regression ([El Asri et al., 2014a](#)), 0.53 using a two-level SVM approach ([Ultes and Minker, 2013](#)), and 0.51 using a hybrid-HMM ([Ultes and Minker, 2014](#)). Human performance on the same task is 0.69 UAR ([Schmitt and Ultes, 2015](#)). A deep learning approach using only non-temporal features achieved an UAR of 0.55 ([Rach et al., 2017](#)).

³UAR is the arithmetic average of all class-wise recalls.

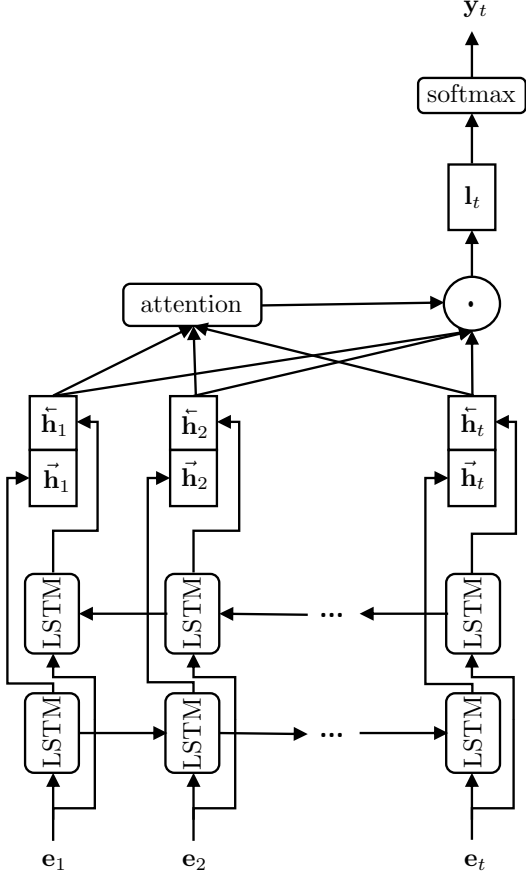


Figure 3: The architecture of the proposed BiLSTM model with self attention. For each time t , the exchange level parameter of all exchanges e_i of the sub-dialogue $i \in \{1 \dots t\}$ are encoded to their respective hidden representation \mathbf{h}_i and are considered and weighted with the self attention mechanism to finally estimate the IQ value \mathbf{y}_t at time t .

4 LSTM-based Interaction Quality Estimation

The proposed IQ estimation model will be used as a reward estimator as depicted in Figure 1. With parameters that are collected from the dialogue system modules for each time step t , the reward estimator derives the reward r_t that is used for learning the dialogue policy π .

The architecture of our proposed IQ estimation model is shown in Figure 3. It is based on the idea that the temporal information that has previously been explicitly encoded with the window and dialogue interaction parameter levels may be learned instead by using recurrent neural networks. Thus, only the exchange level parameters \mathbf{e}_t are considered (see Table 1). Long Short-Term Memory (LSTM) cells are at the core of the model and have originally been proposed by Hochreiter

and Schmidhuber (1997) as a recurrent variant that remedies the vanishing gradient problem (Bengio et al., 1994).

As shown in Figure 3, the exchange level parameters form the input vector \mathbf{e}_t for each time step or turn t to a bi-directional LSTM (Graves et al., 2013) layer. The input vector \mathbf{e}_t encodes the nominal parameters ASRRecognitionStatus, Activity-Type, and Confirmation? as 1-hot representations. In the BiLSTM layer, two hidden states are computed: $\vec{\mathbf{h}}_t$ constitutes the forward pass through the current sub-dialogue and $\overleftarrow{\mathbf{h}}_t$ the backwards pass:

$$\vec{\mathbf{h}}_t = \text{LSTM}(\mathbf{e}_t, \vec{\mathbf{h}}_{t-1}) \quad (3)$$

$$\overleftarrow{\mathbf{h}}_t = \text{LSTM}(\mathbf{e}_t, \overleftarrow{\mathbf{h}}_{t+1}) \quad (4)$$

The final hidden layer is then computed by concatenating both hidden states:

$$\mathbf{h}_t = [\vec{\mathbf{h}}_t, \overleftarrow{\mathbf{h}}_t] \quad (5)$$

Even though information from all time steps may contribute to the final IQ value, not all time steps may be equally important. Thus, an attention mechanism (Vaswani et al., 2017) is used that evaluates the importance of each time step t' for estimating the IQ value at time t by calculating a weight vector $\alpha_{t,t'}$.

$$\mathbf{g}_{t,t'} = \tanh(\mathbf{h}_t^T \mathbf{W}_t + \mathbf{h}_{t'}^T \mathbf{W}_{t'} + \mathbf{b}_t) \quad (6)$$

$$\alpha_{t,t'} = \text{softmax}(\sigma(\mathbf{W}_a \mathbf{g}_{t,t'} + \mathbf{b}_a)) \quad (7)$$

$$\mathbf{l}_t = \sum_{t'} \alpha_{t,t'} \mathbf{h}_{t'} \quad (8)$$

Zheng et al. (2018) describe this as follows: “The attention-focused hidden state representation \mathbf{l}_t of an [exchange] at time step t is given by the weighted summation of the hidden state representation $\mathbf{h}_{t'}$ of all [exchanges] at time steps t' , and their similarity $\alpha_{t,t'}$ to the hidden state representation \mathbf{h}_t of the current [exchange]. Essentially, \mathbf{l}_t dictates how much to attend to an [exchange] at any time step conditioned on their neighbourhood context.”

To calculate the final estimate \mathbf{y}_t of the current IQ value at time t , a softmax layer is introduced:

$$\mathbf{y}_t = \text{softmax}(\mathbf{l}_t) \quad (9)$$

For estimating the interaction quality using a BiLSTM, the proposed architecture frames the task as a classification problem where each sequence is labelled with one IQ value. Thus, for

Table 2: Performance of the proposed LSTM-based variants with the traditional cross-validation setup. Due to overlapping sub-dialogues in the train and test sets, the performance of the LSTM-based models achieve unrealistically high performance.

	UAR	κ	ρ	eA	Ep.
LSTM	0.78	0.85	0.91	0.99	101
BiLSTM	0.78	0.85	0.92	0.99	100
LSTM+att	0.74	0.82	0.91	0.99	101
BiLSTM+att	0.75	0.83	0.91	0.99	93
Rach et al. (2017)	0.55	0.68	0.83	0.94	-
Ultes et al. (2015)	0.55	-	-	0.89	-

each time step t , the IQ value needs to be estimated for the corresponding sub-dialogue consisting of all exchanges from the beginning up to t . Framing the problem like this is necessary to allow for the application of a BiLSTM-approach and still be able to only use information that would be present at the current time step t in an ongoing dialogue interaction.

To analyse the influence of the BiLSTM, a model with a single forward-LSTM layer is also investigated where

$$\mathbf{h}_t = \vec{\mathbf{h}}_t. \quad (10)$$

Similarly, a model without attention is also analysed where

$$\mathbf{l}_t = \mathbf{h}_t. \quad (11)$$

5 Experiments and Results

The proposed BiLSTM IQ estimator is both trained and evaluated on the LEGO corpus and applied within the IQ reward estimation framework (Fig. 1) on several domains within a simulated environment.

5.1 Interaction Quality Estimation

To evaluate the proposed BiLSTM model with attention (BiLSTM+att), it is compared with three of its own variants: a BiLSTM without attention (BiLSTM) as well as a single forward-LSTM layer with attention (LSTM+att) and without attention (LSTM). Additional baselines are defined by Rach et al. (2017) who already proposed an LSTM-based architecture that only uses non-temporal features, and the SVM-based estimation model as originally used for reward estimation by Ultes et al. (2015).

The deep neural net models have been implemented with Keras (Chollet et al., 2015) using

Table 3: Performance of the proposed LSTM-based variants with the dialogue-wise cross-validation setup. The models by Rach et al. (2017) and Ultes et al. (2015) have been re-implemented. The BiLSTM with attention mechanism performs best in all evaluation metrics.

	UAR	κ	ρ	eA	Ep.
LSTM	0.51	0.63	0.78	0.93	8
BiLSTM	0.53	0.63	0.78	0.93	8
LSTM+att	0.52	0.63	0.79	0.92	40
BiLSTM+att	0.54	0.65	0.81	0.94	40
Rach et al. (2017)	0.45	0.58	0.79	0.88	82
Ultes et al. (2015)	0.44	0.53	0.69	0.86	-

the self-attention implementation as provided by Zheng et al. (2018)⁴. All models were trained against cross-entropy loss using RmsProp (Tieleman and Hinton, 2012) optimisation with a learning rate of 0.001 and a mini-batch size of 16.

As evaluation measures, the unweighted average recall (UAR)—the arithmetic average of all class-wise recalls—, a linearly weighted version of Cohen’s κ , and Spearman’s ρ are used. As missing the correct estimated IQ value by only one has little impact for modelling the reward, a measure we call the extended accuracy (eA) is used where neighbouring values are taken into account as well.

All experiments were conducted with the LEGO corpus (Schmitt et al., 2012) in a 10-fold cross-validation setup for a total of 100 epochs per fold. The results are presented in Table 2. Due to the way the task is framed (one label for each sub-dialogue), memorising effects may be observed with the traditional cross-validation setup that has been used in previous work. Hence, the results in Table 2 show very high performance, which is likely to further increase with ongoing training. However, the corresponding models are likely to generalise poorly.

To alleviate this, a dialogue-wise cross-validation setup has been employed also consisting of 10 folds of disjoint sets of dialogues. By that, it can be guaranteed that there are no overlapping sub-dialogues in the training and test sets. All results of these experiments are presented in Table 3 with the absolute improvement of the two main measures UAR and eA over the SVM-based approach of Ultes et al. (2015) visualised in Figure 4.

⁴Code freely available at <https://github.com/CyberZHG/keras-self-attention>




Figure 4: Absolute improvement of the IQ estimation models over the originally employed model by (Ultes et al., 2017a) for IQ-based reward estimation with the dialogue-wise cross-validation setup. UAR and eA take values from 0 to 1

The proposed BiLSTM+att model outperforms existing models and the baselines in all four performance measures by achieving an UAR of 0.54 and an eA of 0.94 after 40 epochs. Furthermore, both the BiLSTM and the attention mechanism by themselves improve the performance in terms of UAR. Based on this findings, the BiLSTM+att model is selected as reward estimator for the experiments in the dialogue policy learning setup as shown in Figure 1.

5.2 Dialogue Policy Learning

To analyse the impact of the IQ reward estimator on the resulting dialogue policy, experiments are conducted comparing three different reward models. The two baselines are in accordance to Ultes et al. (2017a): having the objective task success as principal reward component (R_{TS}) and having the interaction quality estimated by a support vector machine as principal reward component (R_{IQ}^s). TS can be computed by comparing the outcome of each dialogue with the pre-defined goal. Of course, this is only possible in simulation and when evaluating with paid subjects. This goal information is not available to the IQ estimators, nor is it required. Both baselines are compared to our proposed BiLSTM model to estimate the interaction quality used as principal reward component (R_{IQ}^{bi}).

For learning the dialogue behaviour, a policy model based on the GP-SARSA algorithm (Gašić and Young, 2014) is used. This is a value-based method that uses a Gaussian process to approximate the state-value function. As it takes into account the uncertainty of the approximation, it

Table 4: Statistics of the domains the IQ reward estimator is trained on (LetsGo) and applied to (rest).

Domain	Code	# constraints	# DB items
LetsGo		4	-
CamRestaurants	CR	3	110
CamHotels	CH	5	33
SFRestaurants	SR	6	271
SFHotels	SH	6	182
Laptops	L	6	126

is very sample efficient and may even be used to learn a policy directly through real human interaction (Gašić et al., 2013).

The decisions of the policy are based on a summary space representation of the dialogue state tracker. In this work, the focus tracker (Henderson et al., 2014)—an effective rule-based tracker—is used. For each dialogue decision, the policy chooses exactly one summary action out of a set of summary actions which are based on general dialogue acts like *request*, *confirm* or *inform*. The exact number of system actions varies for the domains and ranges from 16 to 25.

To measure the dialogue performance, the task success rate (TSR) and the average interaction quality (AIQ) are measured: the TSR represents the ratio of dialogues for which the system was able to provide the correct result. AIQ is calculated based on the estimated IQ values of the respective model (AIQ^{bi} for the BiLSTM and AIQ^s for the SVM) at the end of each dialogue. As there are two IQ estimators, a distinction is made between AIQ^s and AIQ^{bi} . Additionally, the average dialogue length (ADL) is reported.

For the simulation experiments, the performance of the trained policies on five different domains was evaluated: Cambridge Hotels and Restaurants, San Francisco Hotels and Restaurants, and Laptops. The complexity of each domain is shown in Table 4 and compared to the LetsGo domain (the domain the estimators have been trained on).

The dialogues were created using the publicly available spoken dialogue system toolkit PyDial (Ultes et al., 2017b)⁵ which contains an implementation of the agenda-based user simulator (Schatzmann and Young, 2009) with an additional error model. The error model simulates the required semantic error rate (SER) caused in the real system by the noisy speech channel. For each

⁵Code freely available at <http://www.pydial.org>




Figure 5: Results of the simulated experiments for all domains showing task success rate (TSR) only. Each value is computed after 100 evaluation / 1,000 training dialogues averaged over three trials. Numerical results with significance indicators are shown in Table 5.

domain, all three reward models are compared on three SERs: 0%, 15%, and 30%. More specifically, the applied evaluation environments are based on Env. 1, Env. 3, and Env. 6, respectively, as defined by Casanueva et al. (2017). Hence, for each domain and for each SER, policies have been trained using 1,000 dialogues followed by an evaluation step of 100 dialogues. The task success rates in Figure 5 with exact numbers shown in Table 5 were computed based on the evaluation step averaged over three train/evaluation cycles with different random seeds.

As already known from the experiments conducted by Ultes et al. (2017a), the results of the SVM IQ reward estimator show similar results in terms of TSR for R_{IQ}^s and R_{TS} in all domains for an SER of 0%. This finding is even stronger when comparing R_{IQ}^{bi} and R_{TS} . These high TSRs are achieved while having the dialogues of both IQ-based models result in higher AIQ values compared to R_{TS} throughout the experiments. Of course, only the IQ-based model is aware of the IQ concept and indeed is trained to optimise it.

For higher SERs, the TSRs lightly degrade for the IQ-based reward estimators. However, there seems to be a tendency that the TSR for R_{IQ}^{bi} is more robust against noise compared to R_{IQ}^s while still resulting in better AIQ values.

Finally, even though the differences are mostly not significant, there is also a tendency for R_{IQ}^{bi} to result in shorter dialogues compared to both R_{IQ}^s and R_{TS} .

6 Discussion

One of the major questions of this work addresses the impact of an IQ reward estimator on the resulting dialogues where the IQ estimator achieves better performance than previous ones. Analysing the results of the dialogue policy learning experiment leads to the conclusion that the policy learned with R_{IQ}^{bi} performs similar or better than R_{IQ}^s through out all experiments while still achieving better average user satisfaction compared to R_{TS} . Especially for noisy environments, the improvement is relevant.

The BiLSTM clearly performs better on the LEGO corpus while learning the temporal dependencies instead of using handcrafted ones. However, it entails the risk that these learned temporal dependencies are too specific to the original data so that the model does not generalise well anymore. This would mean that it would be less suitable to be applied to dialogue policy learning for different domains. Luckily, the experiments clearly show that this is not the case.

Obviously, the experiments have only been conducted in a simulated environment and not verified in a user study with real humans. However, the general framework of applying an IQ reward estimator for learning a dialogue policy has already been successfully validated with real user experiments by Ultes et al. (2017a) and it seems rather unlikely that the changes we induce by changing the reward estimator lead to a fundamentally different result.

Table 5: Results of the simulated experiments for all domains showing task success rate (TSR), average interaction quality estimated with the SVM (AIQ^s) and the BiLSTM (AIQ^{bi}), and average dialogue length (ADL) in number of turns. Each value is computed after 100 evaluation / 1,000 training dialogues averaged over three trials with different random seeds. ^{1,2,3} marks statistically significant difference compared to R_{TS} , to R_{IQ}^s , and to AIQ^{bi} , respectively ($p < 0.05$, T-test for TSR and ADL, Mann-Whitney-U test for AIQ).

Domain	SER	TSR			AIQ^s		AIQ^{bi}		ADL		
		R_{TS}	R_{IQ}^s	R_{IQ}^{bi}	R_{TS}	R_{IQ}^s	R_{TS}	R_{IQ}^{bi}	R_{TS}	R_{IQ}^s	R_{IQ}^{bi}
CR	0%	1.00 ^{2,3}	0.99 ¹	0.99 ¹	3.64 ²	3.90 ¹	3.68 ³	3.83 ¹	4.68	4.88	4.59
	15%	0.97	0.94	0.96	3.35 ²	3.65 ¹	3.45 ³	3.63 ¹	5.85 ³	5.33	5.10 ¹
	30%	0.94	0.92	0.90	3.15 ²	3.34 ¹	3.22	3.30	6.34	6.30	6.25
CH	0%	0.98	0.99	0.99	3.26 ²	3.62 ¹	3.33	3.44	5.71	5.61	5.40
	15%	0.96 ²	0.89 ^{1,3}	0.93 ²	2.90	2.88	3.14	3.14	6.28 ²	7.26 ^{1,3}	6.31 ²
	30%	0.86	0.88	0.87	2.38 ²	2.79 ¹	2.79 ³	3.02 ¹	7.94 ³	7.31	6.99 ¹
SR	0%	0.98	0.97	0.98	3.04 ²	3.53 ¹	3.13 ³	3.37 ¹	6.26	6.03	5.80
	15%	0.90 ³	0.88	0.84 ¹	2.40 ²	3.00 ¹	2.85 ³	3.01 ¹	7.99	7.55	7.33
	30%	0.71	0.77	0.78	2.03 ²	2.52 ¹	2.46 ³	2.78 ¹	9.77 ³	9.41	8.50 ¹
SH	0%	0.97	0.99	0.98	3.15 ²	3.52 ¹	3.17 ³	3.36 ¹	5.99 ²	5.50 ¹	5.76
	15%	0.88	0.88	0.89	2.63 ²	2.94 ¹	2.77 ³	3.17 ¹	7.98 ³	7.59 ³	6.63 ^{1,2}
	30%	0.83 ²	0.76 ¹	0.80	2.50	2.63	2.70 ³	2.87 ¹	8.38	9.21	8.37
L	0%	0.98	0.99	0.99	3.26 ²	3.61 ¹	3.28	3.41	5.78	5.44	5.60
	15%	0.89	0.88	0.92	2.58 ²	2.97 ¹	2.92 ³	3.17 ¹	7.19	7.34	6.73
	30%	0.80	0.74	0.77	2.43	2.57	2.79	2.92	8.22 ²	9.32 ^{1,3}	7.97 ²
All	0%	0.98	0.98	0.98	3.23 ²	3.65 ¹	3.31	3.48	5.76	5.50	5.47
	15%	0.92	0.89	0.91	2.76 ²	3.10 ¹	3.02 ^{2,0}	3.20 ¹	7.13	7.06	6.52
	30%	0.83	0.81	0.82	2.49	2.80	2.78	2.97	8.20 ²	8.23 ^{1,3}	7.66 ²

7 Conclusion

In this work we proposed a novel model for interaction quality estimation based on BiLSTMs with attention mechanism that clearly outperformed the baseline while learning all temporal dependencies implicitly. Furthermore, we analysed the impact of the performance increase on learned policies that use this interaction quality estimator as the principal reward component. The dialogues of the proposed interaction quality estimator show a slightly higher robustness towards noise and shorter dialogues while still yielding good performance in terms of both of task success rate and (estimated) user satisfaction. This has been demonstrated by training the reward estimator on a bus information domain and applying it to learn dialogue policies in five different domains (Cambridge restaurants and hotels, San Francisco restaurants and hotels, Laptops) in a simulated experiment.

For future work, we aim at extending the interaction quality estimator by incorporating domain-independent linguistic data to further improve the estimation performance. Furthermore, the effects of using a user satisfaction-based reward estimator needs to be applied to more complex tasks.

References

- Yoshua Bengio, Patrice Simard, Paolo Frasconi, et al. 1994. Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on neural networks*, 5(2):157–166.
- Iñigo Casanueva, Paweł Budzianowski, Pei-Hao Su, Nikola Mrkšić, Tsung-Hsien Wen, Stefan Ultes, Lina Rojas-Barahona, Steve Young, and Milica Gašić. 2017. A benchmarking environment for reinforcement learning based task oriented dialogue management. In *Deep Reinforcement Learning Symposium, 31st Conference on Neural Information Processing Systems (NIPS)*.
- Chih-Chung Chang and Chih-Jen Lin. 2011. LIB-SVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- François Chollet et al. 2015. Keras. <https://keras.io>.
- Lucie Daubigny, Matthieu Geist, and Olivier Pietquin. 2012. Off-policy Learning in Large-scale POMDP-based Dialogue Systems. In *Proceedings of the 37th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2012)*, pages 4989–4992, Kyoto (Japan). IEEE.
- Layla El Asri, Hatim Khouzaimi, Romain Laroche, and Olivier Pietquin. 2014a. Ordinal regression for in-

- teraction quality prediction. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3245–3249. IEEE.
- Layla El Asri, Romain Laroche, and Olivier Pietquin. 2012. Reward Function Learning for Dialogue Management. In *Proceedings of the 6th Starting AI Researchers’ Symposium (STAIRS)*, pages 95–106. IOS Press.
- Layla El Asri, Romain Laroche, and Olivier Pietquin. 2013. Reward shaping for statistical optimisation of dialogue management. In *Statistical Language and Speech Processing*, pages 93–101. Springer.
- Layla El Asri, Romain Laroche, and Olivier Pietquin. 2014b. Task completion transfer learning for reward inference. *Proc of MLIS*.
- Milica Gašić, Catherine Breslin, Matthew Henderson, Dongho Kim, Martin Szummer, Blaise Thomson, Pirros Tsiakoulis, and Steve J. Young. 2013. On-line policy optimisation of Bayesian spoken dialogue systems via human interaction. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8367–8371. IEEE.
- Milica Gašić, Dongho Kim, Pirros Tsiakoulis, Catherine Breslin, Matthew Henderson, Martin Szummer, Blaise Thomson, and Steve J. Young. 2014. Incremental on-line adaptation of POMDP-based dialogue managers to extended domains. In *Proceedings of the 15th International Conference on Spoken Language Processing (INTERSPEECH)*, pages 140–144. ISCA.
- Milica Gašić and Steve J. Young. 2014. Gaussian processes for POMDP-based dialogue manager optimization. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 22(1):28–40.
- Alex Graves, Navdeep Jaitly, and Abdel-rahman Mohamed. 2013. Hybrid speech recognition with deep bidirectional lstm. In *2013 IEEE workshop on automatic speech recognition and understanding*, pages 273–278. IEEE.
- Matthew Henderson, Blaise Thomson, and Jason Williams. 2014. The second dialog state tracking challenge. In *15th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, volume 263.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Oliver Lemon and Olivier Pietquin. 2007. Machine learning for spoken dialogue systems. In *European Conference on Speech Communication and Technologies (Interspeech’07)*, pages 2685–2688.
- Oliver Lemon and Olivier Pietquin. 2012. *Data-Driven Methods for Adaptive Spoken Dialogue Systems*. Springer New York.
- Esther Levin and Roberto Pieraccini. 1997. A stochastic model of computer-human interaction for learning dialogue strategies. In *Eurospeech*, volume 97, pages 1883–1886.
- Niklas Rach, Wolfgang Minker, and Stefan Ultes. 2017. Interaction quality estimation using long short-term memories. In *Proceedings of the 18th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 164–169. Association for Computational Linguistics.
- Antoine Raux, Dan Bohus, Brian Langner, Alan W. Black, and Maxine Eskenazi. 2006. Doing research on a deployed spoken dialogue system: One year of let’s go! experience. In *Proc. of the International Conference on Speech and Language Processing (ICSLP)*.
- Verena Rieser and Oliver Lemon. 2008a. Automatic learning and evaluation of user-centered objective functions for dialogue system optimisation. In *Proceedings of the Sixth International Language Resources and Evaluation (LREC’08)*, pages 2356–2361, Marrakech, Morocco. European Language Resources Association (ELRA). [Http://www.lrec-conf.org/proceedings/lrec2008/](http://www.lrec-conf.org/proceedings/lrec2008/).
- Verena Rieser and Oliver Lemon. 2008b. Learning effective multimodal dialogue strategies from wizard-of-oz data: Bootstrapping and evaluation. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 638–646. ACL.
- Jost Schatzmann and Steve J. Young. 2009. The hidden agenda user simulation model. *Audio, Speech, and Language Processing, IEEE Transactions on*, 17(4):733–747.
- Alexander Schmitt and Stefan Ultes. 2015. Interaction quality: Assessing the quality of ongoing spoken dialog interaction by experts—and how it relates to user satisfaction. *Speech Communication*, 74:12–36.
- Alexander Schmitt, Stefan Ultes, and Wolfgang Minker. 2012. A parameterized and annotated spoken dialog corpus of the cmu let’s go bus information system. In *International Conference on Language Resources and Evaluation (LREC)*, pages 3369–337.
- Pei-Hao Su, M. Gašić, N. Mrkšić, L. Rojas-Barahona, Stefan Ultes, D. Vandyke, T. H. Wen, and S. Young. 2016. On-line active reward learning for policy optimisation in spoken dialogue systems. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 2431–2441. Association for Computational Linguistics.
- Pei-Hao Su, David Vandyke, Milica Gašić, Dongho Kim, Nikola Mrkšić, Tsung-Hsien Wen, and Steve J. Young. 2015. Learning from real users: Rating dialogue success with neural networks for reinforcement learning in spoken dialogue systems. In *Interspeech*, pages 2007–2011. ISCA.

- Richard S. Sutton and Andrew G. Barto. 1998. *Reinforcement Learning: An Introduction*, 1st edition. MIT Press, Cambridge, MA, USA.
- T. Tieleman and G. Hinton. 2012. Lecture 6.5—RmsProp: Divide the gradient by a running average of its recent magnitude. COURSEERA: Neural Networks for Machine Learning.
- Stefan Ultes, Paweł Budzianowski, Iñigo Casanueva, Nikola Mrkšić, Lina Rojas-Barahona, Pei-Hao Su, Tsung-Hsien Wen, Milica Gašić, and Steve Young. 2017a. Domain-independent user satisfaction reward estimation for dialogue policy learning. In *Interspeech*, pages 1721–1725. ISCA.
- Stefan Ultes, Matthias Kraus, Alexander Schmitt, and Wolfgang Minker. 2015. Quality-adaptive spoken dialogue initiative selection and implications on reward modelling. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 374–383. ACL.
- Stefan Ultes and Wolfgang Minker. 2013. Improving interaction quality recognition using error correction. In *Proceedings of the 14th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 122–126. Association for Computational Linguistics.
- Stefan Ultes and Wolfgang Minker. 2014. Interaction Quality Estimation in Spoken Dialogue Systems Using Hybrid-HMMs. In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 208–217. Association for Computational Linguistics.
- Stefan Ultes, Lina M. Rojas-Barahona, Pei-Hao Su, David Vandyke, Dongho Kim, Iñigo Casanueva, Paweł Budzianowski, Nikola Mrkšić, Tsung-Hsien Wen, Milica Gašić, and Steve J. Young. 2017b. Pydial: A multi-domain statistical dialogue system toolkit. In *ACL Demo*. Association of Computational Linguistics.
- Stefan Ultes, Alexander Schmitt, and Wolfgang Minker. 2013. On quality ratings for spoken dialogue systems – experts vs. users. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 569–578. Association for Computational Linguistics.
- David Vandyke, Pei-Hao Su, Milica Gašić, Nikola Mrkšić, Tsung-Hsien Wen, and Steve Young. 2015. Multi-domain dialogue success classifiers for policy training. In *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, pages 763–770. IEEE.
- Vladimir N. Vapnik. 1995. *The nature of statistical learning theory*. Springer-Verlag New York, Inc., New York, NY, USA.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Marilyn Walker. 2000. An application of reinforcement learning to dialogue strategy selection in a spoken dialogue system for email. *Journal of Artificial Intelligence Research*, 12:387–416.
- Marilyn Walker, Jeanne C Fromer, and Shrikanth S. Narayanan. 1998. Learning optimal dialogue strategies: A case study of a spoken dialogue agent for email. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics-Volume 2*, pages 1345–1351. Association for Computational Linguistics.
- Marilyn Walker, Diane J. Litman, Candace A. Kamm, and Alicia Abella. 1997. PARADISE: a framework for evaluating spoken dialogue agents. In *Proceedings of the eighth conference on European chapter of the Association for Computational Linguistics (EACL)*, pages 271–280, Morristown, NJ, USA. Association for Computational Linguistics.
- Jason D. Williams and Steve J. Young. 2004. Characterizing task-oriented dialog using a simulated asr channel. In *Proceedings of the 8th International Conference on Spoken Language Processing (Interspeech 2004)*, pages 185–188.
- Steve J. Young, Milica Gašić, Blaise Thomson, and Jason D. Williams. 2013. POMDP-based statistical spoken dialog systems: A review. *Proceedings of the IEEE*, 101(5):1160–1179.
- Guineng Zheng, Subhabrata Mukherjee, Xin Luna Dong, and Feifei Li. 2018. Opentag: Open attribute value extraction from product profiles. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1049–1058. ACM.

Lifelong and Interactive Learning of Factual Knowledge in Dialogues

Sahisnu Mazumder, Bing Liu, Shuai Wang, Nianzu Ma

Department of Computer Science, University of Illinois at Chicago, USA

sahisnumazumder@gmail.com, liub@uic.edu

shuaiwanghk@gmail.com, jingyima005@gmail.com

Abstract

Dialogue systems are increasingly using *knowledge bases* (KBs) storing real-world facts to help generate quality responses. However, as the KBs are inherently incomplete and remain fixed during conversation, it limits dialogue systems' ability to answer questions and to handle questions involving entities or relations that are not in the KB. In this paper, we make an attempt to propose an engine for *Continuous and Interactive Learning of Knowledge* (CILK) for dialogue systems to give them the ability to continuously and interactively learn and infer new knowledge during conversations. With more knowledge accumulated over time, they will be able to learn better and answer more questions. Our empirical evaluation shows that CILK is promising.

1 Introduction

Dialogue systems, including question-answering (QA) systems are now commonly used in practice. Early such systems were built mainly based on rules and information retrieval techniques (Banchs and Li, 2012; Ameixa et al., 2014; Lowe et al., 2015; Serban et al., 2015). Recent deep learning models (Vinyals and Le, 2015; Xing et al., 2017; Li et al., 2017c) learn from large corpora. However, since they do not use explicit knowledge bases (KBs), they often suffer from generic and dull responses (Xing et al., 2017; Young et al., 2018). KBs have been used to deal with the problem (Ghazvininejad et al., 2018; Le et al., 2016; Young et al., 2018; Long et al., 2017; Zhou et al., 2018). Many task-oriented dialogue systems (Eric and Manning, 2017; Madotto et al., 2018) also use KBs to support information-seeking conversations.

One major shortcoming of existing systems that use KBs is that the KBs are fixed once the dialogue systems are deployed. However, it is almost impossible for the initial KBs to contain all possible

knowledge that the user may ask, not to mention that new knowledge appears constantly. It is thus highly desirable for dialogue systems to learn by themselves while in use, i.e., *learning on the job in lifelong learning* (Chen and Liu, 2018). Clearly, the system can (1) extract more knowledge from the Web or other sources, and (2) learn directly from users during conversations. This paper focuses on the latter and makes an attempt to propose an engine for *Continuous and Interactive Learning of Knowledge* (CILK) to give the dialogue system the ability to acquire/learn new knowledge from the user during conversation. Specifically, it focuses on learning new knowledge interactively from the user when the system is unable to answer a user's WH-question. The acquired new knowledge makes the system better able to answer future user questions, and no longer be limited by the fixed knowledge provided by the human developers.

The type of knowledge that the CILK engine focuses on is the facts that can be expressed as triples, (h, r, t) , which means that the *head entity* h and the *tail entity* t can be linked by the *relation* r . An example of a fact is $(Boston, LocatedInCountry, USA)$, meaning that *Boston is located in USA*. This paper only develops the core engine. It does not study other dialogue functions like response generation, semantic parsing, fact extraction from user utterances, entity linking, etc., which have been studied extensively before and are assumed to be available for use. Thus, this paper works only with structured queries $(h, r, ?)$, e.g., $(Boston, LocatedInCountry, ?)$ meaning "In what Country is Boston located?" or $(?, r, t)$, e.g., $(?, PresidentOf, USA)$ meaning "Who is the President of USA?" It assumes that a semantic parser is available that can convert natural language queries from users into query triples. Similarly, it assumes an information extraction tool like OpenIE (Angeli et al., 2015) is employed to extract facts as triples (h, r, t) from

user’s utterances during conversation. Building a full-fledged dialogue system that can also learn during conversation is a huge undertaking and is out of the scope of this paper. We thus only investigate the core knowledge learning engine here. We also assume that the user has good intentions (i.e., user answers questions with 100% conformity about the veracity of his/her facts)¹; but is not omniscient (opposed to the teacher-student learning setup).

Problem Definition: Given a user query / question $(h, r, ?)$ [or $(?, r, t)$], where r and h (or t) may not be in the KB (i.e., unknown), our goal is two-fold: (i) *answering the user query or rejecting the query to remain unanswered* in the case when the correct answer is believed to not exist in the KB and (ii) *learning / acquiring some knowledge (facts) from the user to help the answering task*. We only focus on the setting where the query cannot be answered *directly* with the current KB and need inference over existing facts, as considering structured query, it’s trivial to retrieve the answer if the answer triple is already in KB. We further distinguish two types of queries: **(1) closed-world queries**, where h (or t) and r are known to the KB, and **(2) open-world queries**, where either one or both h (or t) and r are unknown to the KB.

It is easy to see that the problem is essentially a *lifelong learning* problem (Chen and Liu, 2018), where each query to be processed is a task and the knowledge gained is retained in the KB. To process a new query/task, the knowledge learned and accumulated from the past queries can be leveraged.

For each new open-world query, the proposed approach works in two steps:

Step 1 - Interact with the user: It converts open-world queries (2) to closed-world queries (1) by asking the user questions related to h (or t) and r to make them known to the KB (added to KB). The reason for the conversion will be clear below. The user answers, called *supporting facts* (SFs), are the new knowledge to be added to KB. This step is also called *interactive knowledge learning*. Note, closed-world queries (1) do not need this step.

Step 2 - Infer the query answer: It solves closed-world queries (1) by inferring the query answer. The main idea is to use each entity e in the KB to form a candidate triple (h, r, e) (or (e, r, t)),

¹We envision that the proposed engine is incorporated into a dialogue system in a multi-user environment. The system can perform cross-verification with other users by asking them whether the knowledge (facts) from a user is correct.

USER: *(Boston, LocatedInCountry, ?)* “In what Country is Boston located?” [Query]
CILK: I do not know what “located in Country” means? Can you provide me an example? [Ask for Clue]
USER: *(London, LocatedInCountry, UK)*. “London is located in UK.” [SF1]
CILK: Got it. Can you tell me a fact about “Boston”? [Ask for Entity Fact]
USER: *(Harvard University, UniversityLocatedIn, Boston)*. “Harvard university is located in Boston.” [SF2]
CILK: *(Boston, LocatedInCountry, USA)* “Boston is located in USA.” [Answer]

Figure 1: An example of interactive learning and inference. Note that CILK only works with triples. Each triple above is assumed to be extracted from the sentence after it. *Ask for Clue* and *Ask for Entity Fact* are interaction query types, discussed in Sec. 3. SF denotes supporting fact.

which is then scored. The entity e with the highest score is predicted as the answer of the query.

Scoring each candidate is modeled as a *knowledge base completion* (KBC) problem (Lao and Cohen, 2010; Bordes et al., 2011). KBC aims to infer new facts (knowledge) from existing facts in a KB and is defined as a *link prediction* problem: Given a query triple, $(e, r, ?)$ [or $(?, r, e)$], it predicts a tail entity t_{true} [head entity h_{true}] which makes the query triple true and thus should be added to the KB. KBC makes the *closed-world* assumption that h , r and t are all *known* to exist in the KB (Lao et al., 2011; Bordes et al., 2011, 2013; Nickel et al., 2015). This is not suitable for knowledge learning in conversations because in a conversation, the user can ask or say anything, which may contain entities and relations that are not in the KB. CILK removes the closed-world assumption and allows all h (or t) and/or r to be *unknown* (not in the KB). Step 1 above basically asks the user questions to make h (or t) and/or r known to the KB. Then, an existing KBC model as a query inference model can be applied to retrieve an answer entity from KB.

Figure 1 shows an example. CILK acquires supporting facts SF1 and SF2 to accomplish the goal of *knowledge learning* and utilizes these pieces of knowledge along with existing KB facts to answer the user query (i.e., to infer over the query relation “*LocatedInCountry*”). CILK aims to achieve these two sub-goals. The new knowledge (SFs) is added to the KB for future use². We evaluate CILK using two real-world KBs: *Nell* and *WordNet* and obtain promising results.

²The inferred query answer is not added to the KB as it may be incorrect. But it can be added in a multi-user environment through cross-verification (see footnote 1 and Sec. 4).

2 Related Work

To the best of our knowledge, no existing system can perform the proposed task. We reported a preliminary research in (Mazumder et al., 2018).

CILK is related to interactive language learning (Wang et al., 2016, 2017), which is mainly about language grounding, not about knowledge learning. Li et al. (2017a,b) and Zhang et al. (2017) train chatbots using human teachers who can ask and answer the chatbot questions. Ono et al. (2017), Otsuka et al. (2013), Ono et al. (2016) and Komatani et al. (2016) allow a system to ask the user whether its prediction of category of a term is correct or not. Compared to these works, CILK performs interactive knowledge learning and inference (over existing and acquired knowledge) while conversing with users after the dialogue system has been deployed (i.e., *learning on the job* (Chen and Liu, 2018)) without any teacher supervision or help.

NELL (Mitchell et al., 2015) updates its KB using facts extracted from the Web (complementary to our work). We do not do Web fact extraction.

KB completion (KBC) has been studied in recent years (Lao et al., 2011; Bordes et al., 2011, 2015; Mazumder and Liu, 2017). But they mainly handle facts with known entities and relations. Neelakantan et al. (2015) work on fixed unknown relations with known embeddings, but does not allow unknown entities. Xiong et al. (2018) also deal with queries involving unknown relations, but known entities in the KB. Shi and Wenginger (2018) handles unknown entities by exploiting an external text corpus. None of the KBC methods perform conversational knowledge learning like CILK.

3 Proposed Technique

As discussed in Sec. 1, given a query $(e, r, ?)$ [or $(?, r, e)$]³ from the user, CILK interacts with the user to acquire supporting facts to answer the query. Such an interactive knowledge learning and inference task is realized by the cooperation of three primary components of CILK: **Knowledge base** (KB) \mathcal{K} , **Interaction Module** \mathcal{I} and **Inference Model** \mathcal{M} . The interaction module \mathcal{I} decides whether to ask or not and formulates questions to ask the user for supporting facts. The acquired supporting facts are added to the KB \mathcal{K} and used in training the Inference Model \mathcal{M} which then performs inference over the query (i.e., answers the query).

³Either e or r or both may not exist in the KB

In the following subsections, we formalize the interactive knowledge learning problem (Sec. 3.1), describe the Inference Model \mathcal{M} (Sec. 3.2) and discuss how CILK interacts and processes a query from the user (Sec. 3.3).

3.1 Problem Formulation

CILK’s KB \mathcal{K} is a triple store $\{(h, r, t)\} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$, where \mathcal{E} is the entity set and \mathcal{R} is the relation set. Let q be a query of the form $(e, r, ?)$ [or $(?, r, e)$] issued to CILK, where e is termed as *query entity* and r as the *query relation*. If $e \notin \mathcal{E}$ and/or $r \notin \mathcal{R}$ (we also say $e, r \notin \mathcal{K}$), we call q an *open-world* query. Otherwise, q is referred to as a *closed-world* query, i.e., both e and r exist in \mathcal{K} . Given \mathcal{K} and a query q , the query inference task is defined as follows: If q is of the form $(e, r, ?)$, the goal is to predict a tail entity $t_{true} \in \mathcal{E}$ such that (e, r, t_{true}) holds. We call such q a *tail query*. If q is of the form $(?, r, e)$, the goal is to predict a head entity $h_{true} \in \mathcal{E}$ such that (h_{true}, r, e) holds. We call such q a *head query*. In the open-world setting, it’s quite possible that the answer entity t_{true} (for a tail query) or h_{true} (for a head query) does not exist in the KB (in \mathcal{E}). In such cases, the inference model \mathcal{M} cannot find the true answer. We thus further extend the goal of query inference task to either finding answer entity t_{true} (h_{true}) for q or rejecting q to indicate that the answer does not exist in \mathcal{E} .

Given an open-world (head / tail) query q from user u , CILK interacts with u to acquire a set of supporting facts (SFs) [i.e., a set of *clue* triples C_r involving query relation r and/or a set of *entity fact* triples F_e involving query entity e] for learning r and e (discussed in Sec 3.3). In Figure 1, (*London, LocatedInCountry, UK*) is a clue of query relation “*LocatedInCountry*” and (*Harvard University, UniversityLocatedIn, Boston*) is an entity fact involving query entity “*Boston*”. In this interaction process, CILK decides and asks questions to the user for knowledge acquisition in multiple dialogue turns (see Figure 1). This is **step 1** as discussed in Sec. 1 and will be further discussed in Sec. 3.3.

Once SFs are gathered, it uses $(\mathcal{K} \cup C_r \cup F_e)$ to infer q , which is **step 2** in Sec. 1 and will be detailed in Sec. 3.2. We refer to the whole interaction process involving multi-turn knowledge acquisition followed by the query inference step as a *dialogue session*. In summary, CILK is assumed to operate in multiple dialogue sessions with different users and acquire knowledge in each session and thereby, continuously learns new knowledge over time.

3.2 Inference Model

Given a query q , the Inference Model \mathcal{M} attempts to infer q by predicting the answer entity from \mathcal{E} . In particular, it selects each entity $e_i \in \mathcal{E}$ and forms $|\mathcal{E}|$ number of candidate triples $\{d_1, \dots, d_{|\mathcal{E}|}\}$, where d_i is of the form (e, r, e_i) for a tail query [or (e_i, r, e) for a head query] and then score each d_i to quantify the relevancy of e_i of being an answer to q . The top ranked entity e_i is returned as the predicted answer of q . We deal with the case of query rejection by \mathcal{M} later.

We use the neural knowledge base embedding (KBE) approach (Bordes et al., 2011, 2013; Yang et al., 2014) to design \mathcal{M} . Given a KB represented as a triple store, a neural KBE method learns to encode relational information in the KB using low-dimensional representations (embeddings) of entities and relations and uses the learned representations to predict the correctness of unseen triples. In particular, the goal is to learn representations for entities and relations such that valid triples receive high scores (or low energies) and invalid triples receive low scores (or high energies) defined by a scoring function $S(\cdot)$. The embeddings can be learned via a neural network. In a typical (linear) KBE model, given a triple (h, r, t) , input entity h, t and relation r correspond to high-dimensional vectors (either ‘‘one-hot’’ index vector or ‘‘n-hot’’ feature vector) $\mathbf{x}_h, \mathbf{x}_t$ and \mathbf{x}_r respectively, which are then projected into low dimensional vectors $\mathbf{v}_h, \mathbf{v}_t$ and \mathbf{v}_r using an entity embedding matrix W_E and relation embedding matrix W_R as given by $\mathbf{v}_h = W_E \mathbf{x}_h, \mathbf{v}_r = W_R \mathbf{x}_r$ and $\mathbf{v}_t = W_E \mathbf{x}_t$. The scoring function $S(\cdot)$ is then used to compute a validity score $S(h, r, t)$ of the triple.

Any KBE model can be used for learning \mathcal{M} . For evaluation, we adopt DistMult (Yang et al., 2014) for its state-of-the art performance over many other KBE models (Kadlec et al., 2017). The scoring function of DistMult is defined as follows:

$$S(h, r, t) = \mathbf{v}_h^T \text{diag}(\mathbf{v}_r) \mathbf{v}_t = \sum_{i=1}^N \mathbf{v}_h[i] \mathbf{v}_r[i] \mathbf{v}_t[i] \quad (1)$$

where $\text{diag}(\mathbf{v}_r)$ is the diagonal matrix in \mathbf{v}_r .

The parameters of \mathcal{M} , i.e., W_E and W_R , are learned by minimizing a margin-based ranking objective \mathcal{L} , which encourages the scores of positive triples to be higher than those of negative triples:

$$\mathcal{L} = \sum_{d \in D^+} \sum_{d' \in D^-} \max\{S(d') - S(d) + 1, 0\} \quad (2)$$

where, D^+ is a set of triples observed in \mathcal{K} , treated as positive triples. D^- is a set of negative triples

obtained by corrupting either head entity or tail entity of each +ve triple (h, r, t) in D^+ by replacing it with a randomly chosen entity h' and t' respectively from \mathcal{K} such that the corrupted triples $(h', r, t), (h, r, t') \notin \mathcal{K}$. Note, \mathcal{M} is trained continuously by sampling a set of +ve triples and correspondingly constructing a set of -ve triples as the KB expands with acquired supporting facts to improve its inference capability over new queries (involving new query relations and entities). Thus, the embedding matrices W_E and W_R also grow linearly over time.

Rejection in KB Inference. For a query with no answer entity existing in \mathcal{K} , CILK attempts to reject the query from being answered. To decide whether to reject the query or not, CILK maintains a **threshold buffer** \mathcal{T} that stores entity and relation specific prediction thresholds and updates it continuously over time, as described below.

Besides the dataset for training \mathcal{M} , CILK also creates a validation dataset D_{vd} , consisting of a set of validation query tuples of the form (q, E^+, E^-) . Here, q is either a head or tail query involving query entity e and relation r , $E^+ = \{e_1^+, \dots, e_p^+\}$ is the set of p positive (true answer) entities in \mathcal{K} and $E^- = \{e_1^-, \dots, e_n^-\}$ is the set of n negative entities randomly sampled from \mathcal{K} such that $E^+ \cap E^- = \emptyset$.

Let $D_{vd}^e = \{(q, E^+, E^-) \mid (q, E^+, E^-) \in D_{vd}, e \in q\}$ be the validation query tuple set involving entity e and $D_{vd}^r = \{(q, E^+, E^-) \mid (q, E^+, E^-) \in D_{vd}, r \in q\}$ be the validation query tuple set involving relation r . Then, we compute $\mathcal{T}[z]$, (i.e., prediction threshold for z , where z is either e or r) as the average of the mean scores of triples involving +ve entities and mean scores of triples involving -ve entities, computed over all q in D_{vd}^z , given by-

$$\mathcal{T}[z] = \frac{1}{2|D_{vd}^z|} \sum_{(q, E^+, E^-) \in D_{vd}^z} \mu_E^+ + \mu_E^- \quad (3)$$

where $\mu_E^+ = \frac{1}{|E^+|} \sum_{e_i^+ \in E^+} S(q, e_i^+)$ and $\mu_E^- = \frac{1}{|E^-|} \sum_{e_i^- \in E^-} S(q, e_i^-)$. Here, $S(q, e_i^+) = S(e, r, e_i^+)$ if q is a tail query and $S(e_i^+, r, e)$ if q is a head query. $S(q, e_i^-)$ can be explained in a similar way.

Given a head or tail query q involving query entity e and relation r , we compute the prediction threshold μ_q for q as $\mu_q = \max\{\mathcal{T}[e], \mathcal{T}[r], 0\}$.

Inference Decision Making. If $\tilde{e} \in \mathcal{E}$ is the predicted answer entity by \mathcal{M} for query q and $S(q, \tilde{e}) > \mu_q$, CILK responds to user with answer \tilde{e} . Otherwise, q gets rejected.

Algorithm 1 CILK Knowledge Learning and Inference

Input: query $q_j = (e, r, ?)$ or $(?, r, e)$ issued by user at session- j ; \mathcal{K}_j : CILK's KB at session- j ; \mathcal{P}_j : Performance Buffer at session- j ; \mathcal{T}_j : Threshold Buffer at session- j ; \mathcal{M}_j : trained Inference Model at session- j ; α : probability of treating an acquired supporting fact as training triple; ρ : % of entities or relations in \mathcal{K}_j that belong to the diffident set.
Output: \tilde{e} : predicted entity as answer of query q_j in session- j .

```
1: if  $r \notin \mathcal{K}_j$  or IsDiffident( $r, \mathcal{P}_j, \rho$ ) then
2:    $C_r \leftarrow \text{AskUserforCLUE}(r)$  {acquire supporting facts to learn  $r$ 's embedding}
3: end if
4: if  $e \notin \mathcal{K}_j$  or IsDiffident( $e, \mathcal{P}_j, \rho$ ) then
5:    $F_e \leftarrow \text{AskUserforEntityFacts}(e)$  {Acquire supporting facts to learn  $e$ 's embedding}
6: end if
7: if  $C_r \neq \emptyset$  then
8:    $\mathcal{K}_{(j+\frac{1}{2})} \leftarrow$  Add clue triples from  $C_r$  into  $\mathcal{K}_j$  and randomly mark  $\alpha\%$  of  $C_r$  as training triples and  $(1-\alpha)\%$  as validation triples respectively in  $\mathcal{K}_j$ .
9: end if
10: if  $F_e \neq \emptyset$  then
11:    $\mathcal{K}_{(j+\frac{1}{2})} \leftarrow$  Add fact triples from  $F_e$  into  $\mathcal{K}_{(j+\frac{1}{2})}$  and randomly mark  $\alpha\%$  of these triples as training triples and  $(1-\alpha)\%$  as validation triples.
12: end if
13:  $D_{tr}^r, D_{vd}^r \leftarrow \text{SampleTripleSet}(\mathcal{K}_{(j+\frac{1}{2})}, r)$ 
14:  $D_{tr}^e, D_{vd}^e \leftarrow \text{SampleTripleSet}(\mathcal{K}_{(j+\frac{1}{2})}, e)$ 
15:  $\mathcal{M}_{j+1} \leftarrow \text{TrainInfModel}(\mathcal{M}_j, D_{tr}^r \cup D_{tr}^e)$ 
16:  $\mathcal{P}_{j+1}, \mathcal{T}_{j+1} \leftarrow \text{UpdatePerfandThreshBuffer}(\mathcal{M}_{j+1}, (D_{vd}^r \cup D_{vd}^e), \mathcal{P}_j, \mathcal{T}_j)$ 
17:  $\tilde{e} \leftarrow \text{PredictAnswerEntity}(\mathcal{M}_{j+1}, q_j, \mathcal{T}_{j+1})$ 
```

3.3 Working of CILK

Given a query q involving unknown query entity e and/or relation r , CILK has to ask the user to provide supporting facts to learn embeddings of e and r in order to infer q . However, the user in a given session can only provide very few supporting facts, which may not be sufficient for learning good embeddings of e and r . Moreover, to accumulate a sufficiently good validation dataset for learning $\mathcal{T}[e]$ and $\mathcal{T}[r]$, CILK needs to gather more triples from users involving e and r . But, asking for SFs for any entity and/or relation can be annoying to the user and also, is unnecessary if CILK has already learned good embeddings of that entity and/or relation (i.e., CILK has performed well in predicting true answer entity for queries involving that entity and/or relation in past dialogue sessions with other users). Thus, it is more reasonable to ask for SFs for the *known* entities and/or relations for which CILK is not confident about performing inference accurately, *besides the unknown ones*.

To minimize the rate of user interaction and justify the knowledge acquisition process, CILK uses a **performance buffer** \mathcal{P} to store the performance

statistics of CILK in past dialogue sessions. We use Mean Reciprocal Rank (MRR) to measure the performance of \mathcal{M} (discussed in Sec. 4.1). In particular, $\mathcal{P}[e]$ and $\mathcal{P}[r]$ denote the avg. MRR achieved by \mathcal{M} while answering queries involving e and r respectively, evaluated on the validation dataset D_{vd} . At the end of each dialogue session, CILK detects the set of bottom $\rho\%$ query relations and entities in \mathcal{P} based on MRR scores evaluated on the validation dataset. We call these sets the **diffident** relation and entity sets respectively for the next dialogue session. If the query relation and/or entity issued in the next session belongs to the *diffident* relation or entity set, CILK asks the user for supporting facts⁴. Otherwise, it proceeds with inference, answering or rejecting the query.

Algorithm 1 shows the interactive knowledge learning and inference process of CILK on a query $q_j = (e, r, ?)$ or $(?, r, e)$ in a given dialogue session- j . Let $\mathcal{K}_j, \mathcal{P}_j, \mathcal{T}_j$ and \mathcal{M}_j be the current version of KB, performance buffer, threshold buffer and inference model of CILK at the point when session- j starts. Then, the interactive knowledge learning and inference proceeds as follows:

- If $r \notin \mathcal{K}_j$ or r is diffident in \mathcal{P}_j , the interaction module \mathcal{I} of CILK asks the user to provide clue(s) C_r involving r [Line 1-3]. Similarly, if $e \notin \mathcal{K}_j$ or e is diffident in \mathcal{P}_j , \mathcal{I} asks the user to provide entity fact(s) F_e involving e [Line 4-6].

- If the user provides C_r and/or F_e , \mathcal{I} augments \mathcal{K}_j with triples from C_r and F_e respectively and \mathcal{K}_j expands to \mathcal{K}_{j+1} [Line 7-12]. In this process, $\alpha\%$ of the triples in C_r and F_e are randomly marked as *training* triples and rest $(1-\alpha)\%$ are marked as *validation* triples while storing them in \mathcal{K}_j .

- Next, a set of training triples D_{tr}^r, D_{tr}^e and a set of validation triples D_{vd}^r, D_{vd}^e are sampled randomly from \mathcal{K}_{j+1} involving r and e respectively [Line 13-14] for training and evaluating \mathcal{M}_j . While sampling, we set the ratio of number of training triples to that of validation triples as α to maintain a fixed training and validation set distribution. The size for $(D_{tr}^r \cup D_{tr}^e)$ is set at most N_{tr} (tuned based on real-time training requirements).

- Next, \mathcal{M}_j is trained with $(D_{tr}^r \cup D_{tr}^e)$ and gets updated to \mathcal{M}_{j+1} [Line 15]. Note that, training \mathcal{M}_j with $(D_{tr}^r \cup D_{tr}^e)$ encourages \mathcal{M}_j to learn the embeddings of both r and e before inferring q_j .

⁴Note, if (unknown) e or r appears the first time in a user query, then it cannot be in the diffident set. But the system has to ask the user question by default.

Table 1: Dataset statistics [*kwn* = known, *unk* = unknown]

KB Statistics	WordNet	Nell
# Relations ($\mathcal{K}_{org} / \mathcal{K}_b$)	18 / 12	150 / 142
# Entities ($\mathcal{K}_{org} / \mathcal{K}_b$)	13, 595 / 13, 150	11, 443 / 10, 547
# Triples ($\mathcal{K}_{org} / \mathcal{K}_b$)	53, 573 / 33, 159	66, 529 / 51,252
# Test relations (<i>kwn</i> / <i>unk</i>)	18 (12 / 6)	25 (17 / 8)
# initial Train / initial valid / test (or query) triples (D_q)	29846 / 3323 / 1180	46056 / 5196 / 1250
Test (or query) triples (D_q) statistics [(<i>e</i>, <i>r</i>, ?) or (?, <i>r</i>, <i>e</i>)]		
% triples with only <i>e unk</i>	8.05	19.36
% triples with only <i>r unk</i>	30.25	21.84
% triples both <i>e</i> and <i>r unk</i>	5.25	10.16

Then, we evaluate \mathcal{M}_{j+1} with $(D_{vd}^r \cup D_{vd}^e)$ in order to update the performance buffer \mathcal{P}_j into \mathcal{P}_{j+1} and threshold buffer \mathcal{T}_j into \mathcal{T}_{j+1} [Line 16]. Finally, \mathcal{M}_{j+1} is invoked by CILK to either infer q_j for predicting an answer entity \tilde{e} from \mathcal{K}_{j+1} [Line 17] or reject q_j to indicated that the true answer does not exist in \mathcal{K}_{j+1} . Note, CILK trains \mathcal{M}_j and infers q [Line 13-17] only if $e, q \in \mathcal{K}_{j+1}$.

4 Experiments

As indicated earlier, the proposed CILK system is best used in a *multi-user* environment, so it naturally observes many more query triples (hence, accumulates more facts) from different users over time. Presently CILK fulfills its knowledge learning requirement by only adding the *supporting facts* into the KB. The predicted query triples are not added as they are unverified knowledge. However, in practice, CILK can store these predicted triples in the KB as well after checking their correctness through *cross-verification* while conversing with other users in some future related conversations by smartly asking them. Note that CILK may not verify its prediction with the same user who asked the question/query q because he/she may not know the answer(s) for q . However, there is no problem that it acquires the correct answer(s) of q when it asks q to some other user u' in a future related conversation and u' answers q . At this point, CILK can incorporate q into its KB and also, train itself using triple q . We do not address the issue here.

4.1 Evaluation Setup

Evaluation of CILK with real users in a crowd-source based setup would be very difficult to conduct and prohibitively time-consuming (and expensive) as it needs a large number of real-time and continuous user interaction. Thus, we design a simulated interactive environment for the evaluation.

We create a **simulated user** (a program) to interact with CILK, where the simulated user issues a query to CILK and CILK answers the query. The

(simulated) user has (1) a **knowledge base** (\mathcal{K}_u) for answering questions from CILK, and (2) an **query dataset** (D_q) from which the user issues queries to CILK.⁵ Here, D_q consists of a set of structured query triples q of the form $(e, r, ?)$ and $(?, r, e)$ readable by CILK. In practice, the user only issues queries to CILK, but cannot evaluate the performance of the system unless the user knows the answer. To evaluate the performance of CILK on D_q in the simulated setting, we also collect the answer set for each query $q \in D_q$ (discussed shortly).

As CILK is supposed to perform continuous online knowledge acquisition and learning, we evaluate its performance on the streaming query dataset. We assume that, CILK has been deployed with an initial knowledge base (\mathcal{K}_b) and the inference model \mathcal{M} has been trained over all triples in \mathcal{K}_b for a given number of epochs N_{init} . We call \mathcal{K}_b the **base KB** of CILK which serves as its knowledge base at the time point (t_{eval}) when our evaluation starts. And the training process of \mathcal{M} using triples in \mathcal{K}_b is referred to as the **initial training phase** of CILK onwards. In the initial training phase, we randomly split \mathcal{K}_b triples into a set of training triples D_{tr} and a set of validation triples D_{vd} with 9:1 ratio (we use $\alpha = 0.9$) and train \mathcal{M} with D_{tr} . D_{vd} is used to tune model hyper-parameters and populate initial performance and threshold buffers \mathcal{P} and \mathcal{T} respectively. D_{tr} , D_{vd} , \mathcal{P} , and \mathcal{T} get updated continuously after t_{eval} in the **online training and evaluation phase** (with new acquired triples) during interaction with the simulated user.

The relations and entities in \mathcal{K}_b are regarded as *known* relations and *known* entities to CILK till t_{eval} . Thus, the initial inference model \mathcal{M} is trained and validated with triples involving only *known* relations and *known* entities (in \mathcal{K}_b). During the online training and evaluation phase, CILK faces queries (from D_q) involving both *known* and *unknown* relations and entities. More specifically, if a relation (entity) appearing in a query $q \in D_q$ exists in \mathcal{K}_b , we consider that query relation (entity) as *known* query relation (entity). Otherwise, it is referred to as *unknown* query relation (entity).

We create simulated user’s KB \mathcal{K}_u , base KB (\mathcal{K}_b) and query dataset D_q from two standard KB datasets: (1) **WordNet** (Bordes et al., 2013) and (2) **Nell** (Gardner et al., 2014). From each KB dataset,

⁵Using \mathcal{K}_u and D_q , we can create **simulated dialogues** as well. Utterances in a dialogue can be created using a language template for each triple. Likewise, extraction of triples from utterances can be done using templates as well.

Table 2: Comparison of predictive performance of various versions of CILK. For each KB dataset, we compare the first four (Threshold) variants denoted as “X-BTr” and last three (dataset sampling strategy) variants denoted as “MaxTh-X” and marked the highest H@1 and H@10 values (among each of the groups of four and three) in bold. Thus, some columns have at max. two values marked bold (due to the two comparison groups). *MaxTh-BTr* in the table is the version of CILK proposed in Sec. 3.

	Rel - K / Ent - K			Rel - K / Ent -UNK			Rel - UNK / Ent - K			Rel - UNK / Ent -UNK			Overall		
	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10
WordNet															
EntTh-BTr	0.46	34.57	57.23	0.04	3.50	4.38	0.20	16.21	25.80	0.07	4.83	8.06	0.33	25.03	40.89
RelTh-BTr	0.45	12.71	16.32	0.04	7.89	7.89	0.21	12.30	16.51	0.07	9.67	9.67	0.33	12.09	15.39
MinTh-BTr	0.45	33.81	57.99	0.03	2.63	3.50	0.22	15.93	28.05	0.07	4.84	8.06	0.33	24.43	41.91
MaxTh-BTr	0.45	34.72	56.87	0.04	5.26	6.14	0.20	15.92	25.79	0.07	6.45	9.67	0.33	25.27	40.95
MaxTh-EntTr	0.42	26.07	42.74	0.26	19.29	22.80	0.19	11.79	15.17	0.23	17.74	20.96	0.33	20.77	31.60
MaxTh-RelTr	0.45	34.48	55.93	0.003	2.63	3.51	0.13	11.25	18.01	0.11	8.06	16.13	0.30	23.46	38.09
Nell															
EntTh-BTr	0.37	26.80	47.28	0.06	4.47	7.22	0.15	9.58	19.97	0.04	1.64	7.36	0.22	16.18	29.78
RelTh-BTr	0.37	17.01	25.05	0.06	3.78	4.13	0.16	8.72	17.67	0.03	3.28	4.92	0.23	11.35	17.49
MinTh-BTr	0.37	26.63	47.30	0.06	5.33	8.60	0.15	10.24	23.21	0.03	1.64	5.72	0.23	16.41	30.57
MaxTh-BTr	0.37	27.57	47.58	0.06	4.30	7.57	0.16	10.69	19.61	0.03	4.92	8.20	0.23	17.16	30.03
MaxTh-EntTr	0.34	21.82	42.65	0.13	3.95	7.91	0.22	16.48	20.56	0.06	4.06	4.06	0.24	15.46	27.44
MaxTh-RelTr	0.37	26.60	47.07	0.04	3.44	5.85	0.20	12.18	17.67	0.06	3.28	10.67	0.23	16.67	29.29

we first build a fairly large triple store and use it as the original KB (\mathcal{K}_{org}) and then, create \mathcal{K}_u of user, base KB (\mathcal{K}_b) of CILK and D_q from \mathcal{K}_{org} , as discussed below (Table 1 shows the results).

Simulated User, Base KB Creation and Query Dataset Generation. In Nell, we found 150 relations with ≥ 300 triples, and we randomly selected 25 relations for D_q . We shuffle the list of 25 relations, select 34% of them as *unknown* relations and consider the rest (66%) as *known* relations.

For each *known* relation r , we randomly shuffle the list of distinct triples for r , choose (maximum) 250 triples and randomly select 20% as test and add a randomly chosen subset of the rest of the triples along with the leftovers (not in the list of 250), into \mathcal{K}_b and the other subset are added to \mathcal{K}_u (to provide supporting facts involving poorly learned *known* relations and/or *entities*, if asked [see Sec 3.3]).

For each *unknown* relation r , we remove all triples of r from \mathcal{K}_{org} , randomly choose 20% triples among them and reserve them as query triples for unknown r . Rest 80% triples of *unknown* r are added to \mathcal{K}_u (for providing clues). In this process, we also make sure that the query instances involving *unknown* r are excluded from \mathcal{K}_u . Thus, the user cannot provide the query triple itself as a clue to CILK (during inference) and also, to simulate the case that the user does not know the answer of its issued query. Note, if the user cannot provide a clue for an unknown query relation or a fact for an unknown query entity (not likely), CILK will not be able to correctly answer the query.

At this point, D_q consists of query triples involving both *known* and *unknown* relations, but all *known* entities. To create queries in D_q having unknown entities, we randomly choose 20% of the

Table 3: Performance of CILK Threshold variants on Rejection and prediction decisions. Here, AE (\neg AE) means true answer entity exists (does not exist) in KB. “Pr(pred|AE)” means the probability of predicting an answer, given the true answer exists in KB. “Pr(Reject| \neg AE)” means probability of rejecting the query, given true answer does not exist in KB.

	WordNet		Nell	
	Pr(pred AE)	Pr(Reject \neg AE)	Pr(pred AE)	Pr(Reject \neg AE)
EntTh-BTr	0.85	0.24	0.82	0.15
RelTh-BTr	0.20	0.92	0.26	0.72
MinTh-BTr	0.90	0.18	0.86	0.10
MaxTh-BTr	0.83	0.33	0.72	0.31

entities in D_q triples, remove all triples involving those entities from \mathcal{K}_{org} and add them to \mathcal{K}_u . Now, \mathcal{K}_{org} gets reduced to \mathcal{K}_b (base KB). Next, for each query triple $(h, r, t) \in D_q$, we convert the triple into a head query $q=(?, r, t)$ [or a tail query $q=(h, r, ?)$] by randomly deleting the head or tail entity. We also collect the answer set for each $q \in D_q$ based on observed triples in \mathcal{K}_{org} for CILK evaluation. Note, the generated query triples (with answer entity) in D_q are not directly in \mathcal{K}_b or \mathcal{K}_u .

The WordNet dataset being small, we use all its 18 relations for creating $D_q, \mathcal{K}_u, \mathcal{K}_b$ following Nell. As mentioned earlier, the triples in \mathcal{K}_b are randomly split into 90% training and 10% validation datasets for simulating *initial training phase* of CILK.

Hyper-parameter Settings. Embedding dimensions of entity and relations are empirically set as 250 for WordNet and Nell, initial training epochs N_{init} as 100 for WordNet (140 for Nell), training batch size 128, N_{tr} as 500, $|D_{vd}^r \cup D_{vd}^e|$ as 50, $\alpha = 0.9$, $\rho = 20\%$, random seed as 1000, 4 negative triples generated per positive triple, online training epoch as 5 (2) for each closed (open) world query processing, and learning rate 0.001 for both KB datasets. L2-regularization parameter set as 0.001. Adam optimizer is used for optimization.

Table 4: Overall Performance of **MaxTh-BTr** (CILK), varying the maximum number of clues (#C) and entity facts (#EF) acquired from user per dialogue session (if asked by the interaction module \mathcal{I}).

(#C, #EF)	WordNet			Nell		
	MRR	H@1	H@10	MRR	H@1	H@10
(1, 1)	0.30	22.09	37.83	0.23	16.89	31.14
(1, 2)	0.32	23.00	39.25	0.25	18.11	31.30
(1, 3)	0.33	25.27	40.95	0.23	17.16	30.03
(1, 3)-U	0.31	23.52	38.15	0.21	15.77	28.64
(2, 2)	0.32	23.43	39.05	0.23	16.82	30.33

Compared Models. Since there is no existing work that solves our proposed problem, we compare various versions of CILK, constructed based on different types of prediction threshold μ_q for query rejection (Sec. 3.2) and various online training $D_{tr} = (D_{tr}^r \cup D_{tr}^e)$ and validation dataset $D_{vad} = (D_{vad}^r \cup D_{vad}^e)$ sampling strategies [see Line 13-14 of Algorithm 1] as discussed below:

- **CILK variants based on prediction threshold types**, namely *EntTh-BTr*, *RelTh-BTr*, *MinTh-BTr* and *MaxTh-BTr* (see Table 2). For *EntTh-BTr*, we use $\mu_q = \max\{\mathcal{T}[e], 0\}$, for *RelTh-BTr*, we use $\mu_q = \max\{\mathcal{T}[r], 0\}$, for *MinTh-BTr*, we use $\mu_q = \max\{\min\{\mathcal{T}[e], \mathcal{T}[r]\}, 0\}$ and *MaxTh-BTr* uses $\mu_q = \max\{\mathcal{T}[e], \mathcal{T}[r], 0\}$ as proposed in Sec 3.2. Here, “*BTr*” indicates that the CILK variant samples triples involving both query entity and relation from KB to build D_{tr} and D_{vad} .

- **CILK variants based on dataset sampling strategies:** *MaxTh-BTr* (as explained above), *MaxTh-EntTr* and *MaxTh-RelTr* (see Table 2). Given the query entity e and query relation r , *MaxTh-EntTr* only samples triples involving e and *MaxTh-RelTr* samples only triples involving r to build D_{tr} and D_{vad} . Note, if the sampled dataset D_{tr} (D_{vad}) is \emptyset , CILK skips online training (validation) steps for that session.

Evaluation Metrics. We use two common KBE evaluation metrics: *mean reciprocal rank* (MRR) and *Hits@k* ($H@k$). MRR is the average inverse rank of the top ranked true answer entity for all queries (Bordes et al., 2013). $Hits@k$ is the proportion of test queries for which the true answer entity has appeared in top- k (ranked) predictions. Higher MRR and $Hits@k$ indicate better performance.

4.2 Results and Analysis

For evaluation on a given KB (WordNet or Nell), we randomly generate a chronological ordering of all query instances in D_q , which are fed to the trained CILK (after the initial training phase is over) in a streaming fashion, and then evaluate

Table 5: Performance of **MaxTh-BTr** (CILK) on test queries observed over time, given the model has made a prediction.

% Test Data Observed	WordNet			Nell		
	MRR	H@1	H@10	MRR	H@1	H@10
Overall Performance						
50%	0.37	27.50	47.19	0.29	20.77	38.87
100%	0.37	27.67	46.71	0.29	20.82	38.65
On Open-word Queries						
50%	0.16	11.87	20.11	0.09	4.81	16.47
100%	0.18	12.90	22.91	0.13	8.58	19.54

CILK on the overall query dataset. The avg. test query processing time of CILK is 1.25 sec (on a Nvidia Titan RTX GPU). While evaluating a query q_j , if the true answer of q_j does not exist in KB \mathcal{K}_{j+1} and \mathcal{M}_{j+1} rejects q_j , we consider it as a correct prediction. For such q_j , Reciprocal Rank (RR) cannot be computed. Thus, we exclude q_j while computing MRR, but consider it in computing Hits.

Table 2 shows the performance of CILK variants on the query dataset, evaluated in terms of MRR, H@1 and H@10 for both KBs. We present the overall result on the whole query dataset as well as results on subsets of query datasets, denoted as (*Rel-X*, *Ent-Y*), where X and Y can be either *known* (‘K’) or *unknown* (‘UNK’) and ‘*Rel*’ denotes query relation and ‘*Ent*’ denotes query entity. So, here, (*Rel-K*, *Ent-UNK*) denotes the subset of the query dataset that contains query triples involving only known query relations and unknown query entities (with respect to \mathcal{K}_b). For all variants, we fix the maximum number of clue triples and entity fact triples provided by the simulated user for each query (when asked) as 1 and 3 respectively.

From Table 2, we see that, *MaxTh-BTr* (version of CILK in Sec. 3) achieves the overall best results compared to other variants for both KB datasets. Among different threshold versions, *MaxTh-BTr* and *MinTh-BTr* perform better than the rest. The relatively poor result of *RelTh-BTr* shows threshold strategy plays a vital role in performance improvement. Considering different dataset sampling strategies, again we see *MaxTh-BTr* performs better than other versions. As the triples involving both query entity and relation are selected for online training in *MaxTh-BTr*, CILK gets specifically trained on relevant (query-specific) triples before the query is answered. For other variants, either triples involving query relation (for *MaxTh-EntTr*) or triples involving query entity (for *MaxTh-RelTr*) are discarded, causing a drop in performance.

In Table 3, we compare different CILK threshold variants based on how often it predicts (or rejects) the query, when the true answer exists (does not

exist) in its current KB, given by $\Pr(\text{pred} \mid \text{AE})$ [$\Pr(\text{Reject} \mid \neg\text{AE})$]. For both datasets, *EntTh-BTr* has a tendency to predict more and reject less. Whereas, *RelTh-BTr* is more precautious in prediction. *MinTh-BTr* is the least precautious in prediction among all. *MaxTh-BTr* adopts the best of both worlds (*EntTh-BTr* and *RelTh-BTr*), showing moderate strategy in prediction and rejection behavior.

Table 4 shows comparative performances of *MaxTh-BTr* on varying the maximum number of clue triples and entity fact triples provided by the user (when asked). Comparing (1, 1), (1, 2), (1, 3) we see a clear performance improvement in *MaxTh-BTr* with the increase in (acquired) entity fact triples (specially, for WordNet). This shows that if user interacts more and provides more information for a given query, CILK can gradually improve its performance over time [i.e., with more accumulated triples in its KB]. For Nell, performance improves for both (1, 2) and (1, 3) compared to that in (1, 1), (1, 2) variant being the best overall. Comparing (1, 3) and (2, 2) for both KBs, we see that acquiring more entity facts dominates the overall performance improvement compared to acquiring more clues. This is because, as a past query relation is more probable to appear in future query compared to a past query entity, CILK can gradually learn the relation embedding with less clues per query unlike that for an entity. (1, 3)-U denotes the set up, where CILK asks for clues or entity facts only if the query triple has unknown entity and/or relation, i.e. we disable the use of performance buffer \mathcal{P} (see Sec 3.3). Due to lack of sufficient training triples to learn an unknown query relation and entity, the overall performance degrades. This shows the importance and effectiveness of the performance buffer in improving performance of CILK with limited user interactions.

In Table 5, we show the performance of *MaxTh-BTr* on (predicted) test queries over time. Considering overall performance, the improvement is marginal. However, for open-world queries, there is a substantial improvement in performance as CILK relatively acquires more facts for open-world queries than that of closed-world ones.

5 CILK: Use Cases in Dialogue Systems

There are many applications for CILK. Conversational QA systems (Kiyota et al., 2002; Bordes et al., 2014), conversational recommendation systems (Anelli et al., 2018; Zhang et al., 2018), information-seeking conversational agents (Yang

et al., 2018), etc., that deal with real-world facts, are all potential use cases for CILK.

Recently, (Young et al., 2018; Zhou et al., 2018) showed that dialogue models augmented with commonsense facts improve dialogue generation performance. It’s quite apparent that continuous knowledge learning using CILK can help these models grow their KBs over time and thereby, improve their response generation quality.

The proposed version of CILK has been designed based on a set of assumptions (see Sec. 1) to reduce the complexity of the modeling. For example, we do not handle the case of intentional or unintentional false knowledge injection by users to corrupt the system’s KB. Also, we do not deal with fact extraction errors of the peripheral information extraction module or query parsing errors of the semantic parsing modules, which can affect the knowledge learning of CILK. We believe these are separate research problems and are out of the scope of this work. In future, we plan to model an end-to-end approach of knowledge learning where all peripheral components of CILK can be jointly learned with CILK itself. We also plan to solve the cold start problem when there is little training data for a new relation when it is first added to the KB.

Clearly, CILK does not learn all forms of knowledge. For example, it does not learn new concepts and topics, user traits and personality, and speaking styles. They also form a part of our future work.

6 Conclusion

In this paper, we proposed a continuous (or life-long) and interactive knowledge learning engine CILK for dialogue systems. It exploits the situation when the system is unable to answer a WH-question from the user (considering its existing KB) by asking the user for some knowledge and based on it to infer the query answer. We evaluated the engine on two real-world factual KB data sets and observed promising results. This also shows the potentiality of CILK to serve as a factual knowledge learning engine for future conversational agents.

Acknowledgments

This work was partially supported by a grant from National Science Foundation (NSF IIS 1838770) and a research gift from Northrop Grumman.

References

- David Ameixa, Luisa Coheur, Pedro Fialho, and Paulo Quresma. 2014. Luke, i am your father: dealing with out-of-domain requests by using movies subtitles. In *International Conference on Intelligent Virtual Agents*, pages 13–21. Springer.
- Vito Walter Anelli, Pierpaolo Basile, Derek Bridge, Tommaso Di Noia, Pasquale Lops, Cataldo Musto, Fedelucio Narducci, and Markus Zanker. 2018. Knowledge-aware and conversational recommender systems. In *ACM RecSys*.
- Gabor Angeli, Melvin Johnson Premkumar, and Christopher D Manning. 2015. Leveraging linguistic structure for open domain information extraction. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, pages 344–354.
- Rafael E Banchs and Haizhou Li. 2012. Iris: a chat-oriented dialogue system based on the vector space model. In *Proceedings of the ACL 2012 System Demonstrations*, pages 37–42. ACL.
- Antoine Bordes, Sumit Chopra, and Jason Weston. 2014. Question answering with subgraph embeddings. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 615–620.
- Antoine Bordes, Nicolas Usunier, Sumit Chopra, and Jason Weston. 2015. Large-scale simple question answering with memory networks. *arXiv preprint arXiv:1506.02075*.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Advances in neural information processing systems*.
- Antoine Bordes, Jason Weston, Ronan Collobert, and Yoshua Bengio. 2011. Learning structured embeddings of knowledge bases. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*.
- Zhiyuan Chen and Bing Liu. 2018. *Lifelong machine learning*. Morgan & Claypool Publishers.
- Mihail Eric and Christopher D Manning. 2017. Key-value retrieval networks for task-oriented dialogue. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*.
- Matt Gardner, Partha Talukdar, Jayant Krishnamurthy, and Tom Mitchell. 2014. Incorporating vector space similarity in random walk inference over knowledge bases. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 397–406.
- Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih, and Michel Galley. 2018. A knowledge-grounded neural conversation model. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Rudolf Kadlec, Ondrej Bajgar, and Jan Kleindienst. 2017. Knowledge base completion: Baselines strike back. *Proceedings of the 2nd Workshop on Representation Learning for NLP, ACL*.
- Yoji Kiyota, Sadao Kurohashi, and Fuyuko Kido. 2002. Dialog navigator: A question answering system based on large text knowledge base. In *Proceedings of the 19th international conference on Computational linguistics*, pages 1–7. ACL.
- Kazunori Komatani, Tsugumi Otsuka, Satoshi Sato, and Mikio Nakano. 2016. Question selection based on expected utility to acquire information through dialogue. In *International Workshop on Spoken Dialogue Systems (IWSDS)*.
- Ni Lao and William W Cohen. 2010. Relational retrieval using a combination of path-constrained random walks. *Machine learning*, pages 53–67.
- Ni Lao, Tom Mitchell, and William W Cohen. 2011. Random walk inference and learning in a large scale knowledge base. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 529–539. ACL.
- Phong Le, Marc Dymetman, and Jean-Michel Renders. 2016. Lstm-based mixture-of-experts for knowledge-aware dialogues. *arXiv preprint arXiv:1605.01652*.
- Jiwei Li, Alexander H Miller, Sumit Chopra, Marc’Aurelio Ranzato, and Jason Weston. 2017a. Dialogue learning with human-in-the-loop. *International Conference on Learning Representations*.
- Jiwei Li, Alexander H Miller, Sumit Chopra, Marc’Aurelio Ranzato, and Jason Weston. 2017b. Learning through dialogue interactions. *International Conference on Learning Representations*.
- Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, Alan Ritter, and Dan Jurafsky. 2017c. Adversarial learning for neural dialogue generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2157–2169.
- Yinong Long, Jianan Wang, Zhen Xu, Zongsheng Wang, Baoxun Wang, and Zhuoran Wang. 2017. A knowledge enhanced generative conversational service agent. In *Proceedings of the 6th Dialog System Technology Challenges (DSTC6) Workshop*.
- Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. 2015. The ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 285–294.

- Andrea Madotto, Chien-Sheng Wu, and Pascale Fung. 2018. Mem2seq: Effectively incorporating knowledge bases into end-to-end task-oriented dialog systems. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, pages 1468–1478.
- Sahisnu Mazumder and Bing Liu. 2017. Context-aware path ranking for knowledge base completion. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 1195–1201. AAAI Press.
- Sahisnu Mazumder, Nianzu Ma, and Bing Liu. 2018. Towards a continuous knowledge learning engine for chatbots. *arXiv preprint arXiv:1802.06024*.
- T Mitchell, W Cohen, E Hruschka, P Talukdar, J Beteridge, A Carlson, B Dalvi, M Gardner, B Kisiel, J Krishnamurthy, et al. 2015. Never-ending learning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 2302–2310. AAAI Press.
- Arvind Neelakantan, Benjamin Roth, and Andrew McCallum. 2015. Compositional vector space models for knowledge base completion. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, pages 156–166.
- Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabilovich. 2015. A review of relational machine learning for knowledge graphs. *Proceedings of the IEEE*, pages 11–33.
- Kohei Ono, Ryu Takeda, Eric Nichols, Mikio Nakano, and Kazunori Komatani. 2017. Lexical acquisition through implicit confirmations over multiple dialogues. In *Proceedings of the 18th Annual SIGDIAL Meeting on Discourse and Dialogue*.
- Kohei Ono, Ryu Takeda, Eric Nichols, Mikio Nakano, and Kazunori Komatani. 2016. Toward lexical acquisition during dialogues through implicit confirmation for closed-domain chatbots. In *Proceedings of Second Workshop on Chatbots and Conversational Agent Technologies (WOCHAT)*.
- Tsugumi Otsuka, Kazunori Komatani, Satoshi Sato, and Mikio Nakano. 2013. Generating more specific questions for acquiring attributes of unknown concepts from users. In *14th Annual SIGDIAL Meeting on Discourse and Dialogue*.
- Iulian Vlad Serban, Ryan Lowe, Peter Henderson, Laurent Charlin, and Joelle Pineau. 2015. A survey of available corpora for building data-driven dialogue systems. *arXiv preprint arXiv:1512.05742*.
- Baoxu Shi and Tim Wenginger. 2018. Open-world knowledge graph completion. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Sida Wang, Percy Liang, and Christopher D Manning. 2016. Learning language games through interaction. In *54th Annual Meeting of the Association for Computational Linguistics*, pages 2368–2378. ACL.
- Sida I Wang, Samuel Ginn, Percy Liang, and Christopher D Manning. 2017. Naturalizing a programming language via interactive learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 929–938.
- Chen Xing, Wei Wu, Yu Wu, Jie Liu, Yalou Huang, Ming Zhou, and Wei-Ying Ma. 2017. Topic aware neural response generation. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Wenhan Xiong, Mo Yu, Shiyu Chang, Xiaoxiao Guo, and William Yang Wang. 2018. One-shot relational learning for knowledge graphs. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1980–1990.
- Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2014. Embedding entities and relations for learning and inference in knowledge bases. *International Conference on Learning Representations*.
- Liu Yang, Minghui Qiu, Chen Qu, Jiafeng Guo, Yongfeng Zhang, W Bruce Croft, Jun Huang, and Haiqing Chen. 2018. Response ranking with deep matching networks and external knowledge in information-seeking conversation systems. In *ACM SIGIR*.
- Tom Young, Erik Cambria, Iti Chaturvedi, Hao Zhou, Subham Biswas, and Minlie Huang. 2018. Augmenting end-to-end dialogue systems with commonsense knowledge. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Haichao Zhang, Haonan Yu, and Wei Xu. 2017. Listen, interact and talk: Learning to speak via interaction. *arXiv preprint arXiv:1705.09906*.
- Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W Bruce Croft. 2018. Towards conversational search and recommendation: System ask, user respond. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 177–186. ACM.
- Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018. Commonsense knowledge aware conversation generation with graph attention. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 4623–4629. AAAI Press.

Few-Shot Dialogue Generation Without Annotated Data: A Transfer Learning Approach

Igor Shalyminov[†], Sungjin Lee[‡], Arash Eshghi[†], and Oliver Lemon[†]

[†]Heriot-Watt University, UK

[‡]Microsoft Research, US

[†]{is33, a.eshghi, o.lemon}@hw.ac.uk, [‡]sungjinlee.plus@gmail.com

Abstract

Learning with minimal data is one of the key challenges in the development of practical, production-ready goal-oriented dialogue systems. In a real-world enterprise setting where dialogue systems are developed rapidly and are expected to work robustly for an ever-growing variety of domains, products, and scenarios, efficient learning from a limited number of examples becomes indispensable.

In this paper, we introduce a technique to achieve state-of-the-art dialogue generation performance in a few-shot setup, without using any annotated data. We do this by leveraging background knowledge from a larger, more highly represented dialogue source — namely, the MetaLWOz dataset. We evaluate our model on the Stanford Multi-Domain Dialogue Dataset, consisting of human-human goal-oriented dialogues in in-car navigation, appointment scheduling, and weather information domains.

We show that our few-shot approach achieves state-of-the-art results on that dataset by consistently outperforming the previous best model in terms of BLEU and Entity F1 scores, while being more data-efficient by not requiring any data annotation.

1 Introduction

Data-driven dialogue systems are becoming widely adopted in enterprise environments. One of the key properties of a dialogue model in this setting is its *data efficiency*, i.e. whether it can attain high accuracy and good generalization properties when only trained from minimal data.

Recent deep learning-based approaches to training dialogue systems (Ultes et al., 2018; Wen et al., 2017) put emphasis on collecting large amounts of data in order to account for numerous variations in the user inputs and to cover as many

dialogue trajectories as possible. However, in real-world production environments there isn't enough domain-specific data easily available throughout the development process. In addition, it's important to be able to rapidly adjust a system's behavior according to updates in requirements and new product features in the domain. Therefore, data-efficient training is a priority direction in dialogue system research.

In this paper, we build on a technique to train a dialogue model for a new domain in a 'zero-shot' setup (in terms of full dialogues in the target domain) only using *annotated* 'seed' utterances (Zhao and Eskénazi, 2018).

We present an alternative, 'few-shot' approach to data-efficient dialogue system training: we do use complete in-domain dialogues while using approximately the same amount of training data as Zhao and Eskénazi (2018), with respect to utterances. However, in our method, *no annotation is required* — we instead use a latent dialogue act annotation learned in an unsupervised way from a larger (multi-domain) data source, broadly following the model of Zhao et al. (2018). This approach is potentially more attractive for practical purposes because it is easier to collect unannotated dialogues than collecting utterances across various domains under a consistent annotation scheme.

2 Related Work

There is a substantial amount of work on learning dialogue with minimal data — starting with the Dialog State Tracking Challenge 3 (Henderson et al., 2014) where the problem was to adjust a pre-trained state tracker to a different domain using a seed dataset.

In dialogue response generation, there has also been work on bootstrapping a goal-oriented dialogue system from a few examples using a lin-

guistically informed model: (Eshghi et al., 2017) used an incremental semantic parser – DyLan (Eshghi et al., 2011; Eshghi, 2015) – to obtain contextual meaning representations, and based the dialogue state on this (Kalatzis et al., 2016). Incremental response generation was learned using Reinforcement Learning, again using the parser to incrementally process the agent’s output and thus prune ungrammatical paths for the learner. Compared to a neural model — End-to-End Memory Network (Sukhbaatar et al., 2015), this linguistically informed model was superior in a 1-shot setting (Shalymov et al., 2017). At the same time, its main linguistic resource — a domain-general dialogue grammar for English — makes the model inflexible unless wide coverage is achieved.

Transfer learning for Natural Language Processing is strongly motivated by recent advances in vision. When training a convolutional neural network (CNN) on a small dataset for a specific problem domain, it often helps to learn low-level convolutional features from a greater, more diverse dataset. For numerous applications in vision, ImageNet (Deng et al., 2009) became the source dataset for pre-training convolutional models. For NLP, the main means for transfer were Word2Vec word embeddings (Mikolov et al., 2013) which have recently been updated to models capturing contexts as well (Peters et al., 2018; Devlin et al., 2018). While these tools are widely known to improve performance in various tasks, more specialized models could as well be created for specific research areas, e.g. dialogue generation in our case.

The models above are some of the approaches to one of the central issues of efficient knowledge transfer — learning a unified data representation generalizable across datasets, dubbed ‘representation learning’. In our approach, we will use one such technique based on variational autoencoding with discrete latent variables (Zhao et al., 2018). In this paper we present an approach to transfer learning which is more tailored — both model-wise and dataset-wise — to goal-oriented dialogue in underrepresented domains.

3 The approach

3.1 Zero-shot theoretical framework

We first describe the original Zero-Shot Dialogue Generation (ZSDG) theoretical framework of (Zhao and Eskénazi, 2018) which we base our

work on. For ZSDG, there is a set of source dialogue domains and one target domain, with the task of training a dialogue response generation model from all the available source data and a significantly reduced subset of the target data (referred to as *seed* data). The trained system’s performance is evaluated exclusively on the target domain.

More specifically, the data in ZSDG is organized as follows. There are unannotated dialogues in the form of $\{c, \mathbf{x}, d\}_{src/tgt}$ — tuples of dialogue contexts, responses, and domain names respectively for each of the source and target domains. There are also domain descriptions in the form of $\{\mathbf{x}, \mathbf{a}, d\}_{src/tgt}$ — tuples of utterances, slot-value annotations, and domain names respectively for source and target domains.

ZSDG is essentially a hierarchical encoder-decoder model which is trained in a multi-task fashion by receiving two types of data: (1) dialogue batches drawn from all the available source-domain data, and (2) seed data batches, a limited number of which are drawn from domain description data for all of the source and target domains.

ZSDG model optimizes for 2 objectives. With dialogue batches, the model maximizes the probability of generating a response given the context:

$$\mathcal{L}_{dialog} = -\log p_{\mathcal{F}^d}(\mathbf{x} | \mathcal{F}^e(\mathbf{c}, d)) + \lambda \mathcal{D}(\mathcal{R}(\mathbf{x}, d) || \mathcal{F}^e(\mathbf{c}, d)) \quad (1)$$

where \mathcal{F}^e and \mathcal{F}^d are respectively the encoding and decoding components of a hierarchical generative model; \mathcal{R} is the shared recurrent utterance encoder (the *recognition model*); and \mathcal{D} is a distance function (L_2 norm).

In turn, with domain description batches, the model maximizes the probability of generating the utterance given its slot-value annotation, both represented as sequences of tokens:

$$\mathcal{L}_{dd} = -\log p_{\mathcal{F}^d}(\mathbf{x} | \mathcal{R}(\mathbf{a}, d)) + \lambda \mathcal{D}(\mathcal{R}(\mathbf{x}, d) || \mathcal{R}(\mathbf{a}, d)) \quad (2)$$

In this multi-task setup, the latent space of \mathcal{R} is shared between both utterances and domain descriptions across all the domains. Moreover, the distance-based loss terms make sure that (a) utterances with similar annotations are closer together in the latent space (Eq. 2), and (b) utterances are closer to their dialogue contexts (Eq. 1) so that their encodings capture some of the contexts’




Figure 1: Model architecture. At the pre-training stage (1a), we train the discretized LAED dialogue representation on the Transfer dataset. We then train a zero/few-shot dialogue generation model on SMD with this representation incorporated (1b).

meaning. These properties of the model make it possible to achieve better cross-domain generalization.

3.2 Unsupervised representation learning

As was the case with ZSDG, robust representation learning helps achieve better generalization across domains. The most widely-adopted way to train better representations has been to leverage a greater data source. In this work, we consider unsupervised, variational autoencoder-based (VAE) representation learning on a large dataset of unannotated dialogues. The specific approach we refer to is the Latent Action Encoder-Decoder (LAED) model of (Zhao et al., 2018). LAED is a variant of VAE with two modifications: (1) an optimization objective augmented with mutual information between the input and the latent variable for better and more stable learning performance, and (2) discretized latent variable for the interpretability of the resulting latent actions. Just as in ZSDG, LAED is a hierarchical encoder-decoder model with the key component being a discrete-information (DI) utterance-level VAE. Two versions of this model are introduced, with respective optimization objectives:

$$\mathcal{L}_{DI-VAE} = \mathbb{E}_{q_{\mathcal{R}}(z|x)p(x)} [\log p_{\mathcal{G}}(x | z)] - KL(q(z)||p(z)) \quad (3)$$

$$\mathcal{L}_{DI-VST} = \mathbb{E}_{q_{\mathcal{R}}(z|x)p(x)} [\log p_{\mathcal{G}}^n(x_n | z)p_{\mathcal{G}}^p(x_p | z)] - KL(q(z)||p(z)) \quad (4)$$

where \mathcal{R} and \mathcal{G} are recognition and generation components respectively, x is the model’s input, z is the latent variable, and $p(z)$ and $q(z)$ are respectively prior and posterior distributions of z .

DI-VAE works in a standard VAE fashion reconstructing the input x itself, while DI-VST follows the idea of Variational Skip-Thought reconstructing the input’s previous and next contexts: $\{x_n, x_p\}$. As reported by the authors, the two models capture different aspects of utterances, i.e. DI-VAE reconstructs specific words within an utterance, whereas DI-VST captures the overall intent better — see the visualization in Figure 1a.

3.3 Proposed models¹

In our approach, we simplify the ZSDG setup by not using any explicit domain descriptions, therefore we only work with ‘dialogue’ batches. We also make use of Knowledge Base information without loss of generality (see Section 5) — thus we work with data of the form $\{c, x, k, d\}$ where k is the KB information. We refer to this model as Few-Shot Dialogue Generation, or *FSDG*.

For learning a reusable dialogue representation, we use an external multi-domain dialogue dataset, the Transfer dataset (see Section 4).

We perform a preliminary training stage on it where we train 2 LAED models, both DI-VAE and DI-VST. Then, at the main training stage, we use the hierarchical encoders of both models and incorporate them with FSDG’s decoder. Thus, we have the following encoding function (which is

¹Code is available at https://bit.ly/fsdg_sigdial2019

Model \ Domain	Navigation		Weather		Schedule	
	BLEU, %	Entity F1, %	BLEU, %	Entity F1, %	BLEU, %	Entity F1, %
ZSDG	5.9	14.0	8.1	31	7.9	36.9
NLU_ZSDG	6.1 ± 2.2	12.7 ± 3.3	5.0 ± 1.6	16.8 ± 6.7	6.0 ± 1.7	26.5 ± 5.4
NLU_ZSDG+LAED	7.9 ± 1	12.3 ± 2.9	8.7 ± 0.6	21.5 ± 6.2	8.3 ± 1	20.7 ± 4.8
FSDG@1%	6.0 ± 1.8	9.8 ± 4.8	6.9 ± 1.1	22.2 ± 10.7	5.5 ± 0.8	25.6 ± 8.2
FSDG@3%	7.9 ± 0.7	11.8 ± 4.4	9.6 ± 1.8	39.8 ± 7	8.2 ± 1.1	34.8 ± 4.4
FSDG@5%	8.3 ± 1.3	15.3 ± 6.3	11.5 ± 1.6	38.0 ± 10.5	9.7 ± 1.4	37.6 ± 8.0
FSDG@10%	9.8 ± 0.8	19.2 ± 3.2	12.9 ± 2.4	40.4 ± 11.0	12.0 ± 1.0	38.2 ± 4.2
FSDG+VAE@1%	3.6 ± 2.6	9.3 ± 4.1	6.8 ± 1.3	23.2 ± 10.1	4.6 ± 1.6	28.9 ± 7.3
FSDG+VAE@3%	6.9 ± 1.9	15.6 ± 5.8	9.5 ± 2.6	32.2 ± 11.8	6.6 ± 1.7	34.8 ± 7.7
FSDG+VAE@5%	7.8 ± 1.9	12.7 ± 4.2	10.1 ± 2.1	40.3 ± 10.4	8.2 ± 1.7	34.2 ± 8.7
FSDG+VAE@10%	9.0 ± 2.0	18.0 ± 5.8	12.9 ± 2.2	40.1 ± 7.6	11.6 ± 1.5	39.9 ± 6.9
FSDG+LAED@1%	7.1 ± 0.8*	10.1 ± 4.5	10.6 ± 2.1*	31.4 ± 8.1*	7.4 ± 1.2	29.1 ± 6.6
FSDG+LAED@3%	9.2 ± 0.8	14.5 ± 4.8*	13.1 ± 1.7	40.8 ± 6.1	9.2 ± 1.2*	32.7 ± 6.1
FSDG+LAED@5%	10.3 ± 1.2	15.6 ± 4.5	14.5 ± 2.2	40.9 ± 8.6	11.8 ± 1.9	37.6 ± 6.1*
FSDG+LAED@10%	12.3 ± 0.9	17.3 ± 4.5	17.6 ± 1.9	47.5 ± 6.0	15.2 ± 1.6	38.7 ± 8.4

Table 1: Evaluation results. Marked with asterisks are individual results higher than the ZSDG baseline which are achieved with the minimum amount of training data, and in bold is the model consistently outperforming ZSDG in all domains and metrics with minimum data.

then plugged in to the Eq. 1):

$$\begin{aligned} \mathcal{F}^e(\mathbf{c}, \mathbf{k}, d) = & \mathcal{F}_{DI-VAE}^e(\mathbf{c}, \mathbf{k}, d) \\ & \oplus \mathcal{F}_{DI-VST}^e(\mathbf{c}, \mathbf{k}, d) \\ & \oplus \mathcal{F}_{FSDG}^e(\mathbf{c}, \mathbf{k}, d) \end{aligned} \quad (5)$$

where \oplus is the concatenation operator. We refer to this model as **FSDG+LAED**.

We compare this LAED-augmented model to a similar one, with latent representation trained on the same data but using a regular VAE objective and thus providing regular continuous embeddings (we refer to it as **FSDG+VAE**).

$$\begin{aligned} \mathcal{L}_{VAE} = & \mathbb{E}_{q_{\mathcal{R}}(z|\mathbf{x})}[\log p_{\mathcal{G}}(\mathbf{x} | z)] \\ & - KL(q_{\mathcal{R}}(z)||p(z)) \end{aligned} \quad (6)$$

Finally, in order to explore the original ZSDG setup as much as possible, we also consider its version with automatic Natural Language Understanding (NLU) markup instead of human annotations as domain descriptions. Our NLU annotations include Named Entity Recognizer (Finkel et al., 2005), a date/time extraction library (Chang and Manning, 2012), and a Wikidata entity linker (Pappu et al., 2017). We have models with (NLU_ZSDG+LAED) and without LAED representation (NLU_ZSDG). Our entire setup is shown in Figure 1.

4 Datasets

We use the Stanford Multi-Domain (SMD) human-human goal-oriented dialogue dataset

(Eric et al., 2017) in 3 domains: appointment scheduling, city navigation, and weather information. Each dialogue comes with knowledge base snippet from the underlying domain-specific API.

For LAED training, we use MetaLWOz (Lee et al., 2019), a human-human goal-oriented dialogue corpus specifically designed for various meta-learning and pre-training purposes. It contains conversations in 51 domains with several tasks in each of those. The dialogues are collected using the Wizard-of-Oz method where human participants were given a problem domain and a specific task. No domain-specific APIs or knowledge bases were available for the participants, and in the actual dialogues they were free to use fictional names and entities in a consistent way. The dataset totals more than 40,000 dialogues, with the average length of 11.9 turns.

5 Experimental setup and evaluation

Our few-shot setup is as follows. Given the target domain, we first train LAED models (a dialogue-level DI-VST and an utterance-level DI-VAE, both of the size 10×5) on the MetaLWOz dataset — here we exclude from training every domain that might overlap with the target one.

Next, using the LAED encoders, we train a Few-Shot Dialogue Generation model on all the SMD source domains. We use a random sample (1% to 10%) of the target domain utterances together with their contexts as seed data.

We incorporate KB information into our model by simply serializing the records and prepending

them to the dialogue context, ending up with a setup similar to CopyNet in (Eric et al., 2017).

For the NLU_ZSDG setup, we use 1000 random seed utterances from each source domain and 200 utterances from the target domain².

For evaluation, we follow the approach of (Zhao and Eskénazi, 2018) and report BLEU and Entity F1 scores — means/variances over 10 runs.

6 Results and discussion

Our results are shown in Table 1. Our objective here is maximum accuracy with minimum training data required, and it can be seen that few-shot models with LAED representation are the best performing models for this objective. While the improvements can already be seen with simple FSDG, the use of LAED representation helps to significantly reduce the amount of in-domain training data needed: in most cases, the state-of-the-art results are attained with as little as 3% of in-domain data. At 5%, we see that FSDG+LAED consistently improves upon all other models in every domain, either by increasing the mean accuracy or by decreasing the variation. In SMD, with its average dialogue length of 5.25 turns (see Table 4), 5% of training dialogues amounts to approximately 200 in-domain training utterances. In contrast, the ZSDG setup used approximately 150 *annotated* training utterances for each of the 3 domains, totalling about 450 *annotated* utterances. Although in our few-shot approach we use full in-domain dialogues, we end up having a comparable amount of target-domain training data, with the crucial difference that none of those has to be annotated for our approach. Therefore, the method we introduced attains state-of-the-art in both accuracy and data-efficiency.

The results of the ZSDG_NLU setup demonstrate that single utterance annotations, if not domain-specific and produced by human experts, don't provide as much signal as raw dialogues.

The comparison of the setups with different latent representations also gives us some insight: while the VAE-powered FSDG model improves on the baseline in multiple cases, it lacks generalization potential compared to LAED. The reason for that might be inherently more stable training of LAED due to its modified objective function

²The numbers are selected so that the domain description task is kept secondary.

which in turn results in a more informative, generalizable representation.

Finally, we discuss the evaluation metrics. Since we base this paper on the work of (Zhao and Eskénazi, 2018), we have had to fully conform to the metrics they used to enable direct comparison. However, BLEU as the primary evaluation metric, does not necessarily reflect NLG quality in dialogue settings — see examples in Table 2 of the Appendix (see also Novikova et al. (2017)). This is a general issue in dialogue model evaluation since the variability of possible responses equivalent in meaning is very high in dialogue. In future work, instead of using BLEU, we will put more emphasis on the meaning of utterances, for example by using external dialogue act tagging resources, using quality metrics of language generation – e.g. perplexity – as well as more task-oriented metrics like Entity F1. We expect these to make for more meaningful evaluation criteria.

7 Conclusion and future work

In this paper, we have introduced a technique to achieve state-of-the-art dialogue generation performance in a few-shot setup, without using any annotated data. By leveraging larger, more highly represented dialogue sources and learning robust latent dialogue representations from them, we obtained a model with superior generalization to an underrepresented domain. Specifically, we showed that our few-shot approach achieves state-of-the-art results on the Stanford Multi-Domain dataset while being more data-efficient than the previous best model, by not requiring any data annotation.

Although being state-of-the-art, the accuracy scores themselves still suggest that our technique is not ready for immediate adoption for real-world production purposes, and the task of few-shot generalization to a new dialogue domain remains an area of active research. We expect that such initiatives will be fostered by the release of large dialogue corpora such as MetaLWOz.

In our own future work, we will try and find ways to improve the unsupervised representation in order to increase the transfer potential. Adversarial learning can also be beneficial in the setting of limited data. And apart from improving the model itself, it is necessary to consider an alternative criterion to BLEU-score for adequate evaluation of response generation.

References

- Angel X. Chang and Christopher D. Manning. 2012. [Sutime: A library for recognizing and normalizing time expressions](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation, LREC 2012, Istanbul, Turkey, May 23-25, 2012*, pages 3735–3740.
- Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder-decoder for statistical machine translation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1724–1734.
- J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Mihail Eric, Lakshmi Krishnan, Francois Charette, and Christopher D. Manning. 2017. [Key-value retrieval networks for task-oriented dialogue](#). In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue, Saarbrücken, Germany, August 15-17, 2017*, pages 37–49.
- A. Eshghi, M. Purver, and Julian Hough. 2011. Dylan: Parser for dynamic syntax. Technical report, Queen Mary University of London.
- Arash Eshghi. 2015. DS-TTR: An incremental, semantic, contextual parser for dialogue. In *Proceedings of Semdial 2015 (goDial), the 19th workshop on the semantics and pragmatics of dialogue*.
- Arash Eshghi, Igor Shalymov, and Oliver Lemon. 2017. [Bootstrapping incremental dialogue systems from minimal data: the generalisation power of dialogue grammars](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 2220–2230.
- Jenny Rose Finkel, Trond Grenager, and Christopher D. Manning. 2005. [Incorporating non-local information into information extraction systems by gibbs sampling](#). In *ACL 2005, 43rd Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, 25-30 June 2005, University of Michigan, USA*, pages 363–370.
- Matthew Henderson, Blaise Thomson, and Jason D. Williams. 2014. [The third dialog state tracking challenge](#). In *2014 IEEE Spoken Language Technology Workshop, SLT 2014, South Lake Tahoe, NV, USA, December 7-10, 2014*, pages 324–329.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Computation*, 9(8):1735–1780.
- Dimitrios Kalatzis, Arash Eshghi, and Oliver Lemon. 2016. [Bootstrapping incremental dialogue systems: using linguistic knowledge to learn from minimal data](#). In *Proceedings of the NIPS 2016 workshop on Learning Methods for Dialogue*, Barcelona.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *CoRR*, abs/1412.6980.
- Sungjin Lee, Hannes Schulz, Adam Atkinson, Jianfeng Gao, Kaheer Suleman, Layla El Asri, Mahmoud Adada, Minlie Huang, Shikhar Sharma, Wendy Tay, and Xiujun Li. 2019. [Multi-domain task-completion dialog challenge](#).
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. [Distributed representations of words and phrases and their compositionality](#). In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States.*, pages 3111–3119.
- Jekaterina Novikova, Ondrej Dusek, Amanda Cercas-Curry, and Verena Rieser. 2017. [Why we need new evaluation metrics for nlg](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Aasish Pappu, Roi Blanco, Yashar Mehdad, Amanda Stent, and Kapil Thadani. 2017. [Lightweight multilingual entity extraction and linking](#). In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining, WSDM 17, New York, NY, USA. ACM*.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 2227–2237.
- Igor Shalymov, Arash Eshghi, and Oliver Lemon. 2017. [Challenging neural dialogue models with natural data: Memory networks fail on incremental phenomena](#). In *Proceedings of the 21st Workshop on the Semantics and Pragmatics of Dialogue, SemDial 2017*.
- Sainbayar Sukhbaatar, Arthur Szlam, Jason Weston, and Rob Fergus. 2015. [End-to-end memory networks](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 2440–2448.

- Stefan Ultes, Pawel Budzianowski, Iñigo Casanueva, Lina Maria Rojas-Barahona, Bo-Hsiang Tseng, Yen-Chen Wu, Steve J. Young, and Milica Gasic. 2018. [Addressing objects and their relations: The conversational entity dialogue model](#). In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue, Melbourne, Australia, July 12-14, 2018*, pages 273–283.
- Tsung-Hsien Wen, Yishu Miao, Phil Blunsom, and Steve J. Young. 2017. [Latent intention dialogue models](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pages 3732–3741.
- Tiancheng Zhao and Maxine Eskénazi. 2018. [Zero-shot dialog generation with cross-domain latent actions](#). In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue, Melbourne, Australia, July 12-14, 2018*, pages 1–10.
- Tiancheng Zhao, Kyusong Lee, and Maxine Eskénazi. 2018. [Unsupervised discrete sentence representation learning for interpretable neural dialog generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 1098–1107.

Domain	Context	Gold response	Predicted response
schedule	<usr> Remind me to take my pills	Ok setting your medicine appointment for 7pm	Okay, setting a reminder to take your pills at 7 pm.
	<sys> What time do you need to take your pills?		
	<usr> I need to take my pills at 7 pm.		
navigate	<usr> Find the address to a hospital	Have a good day	No problem.
	<sys> Stanford Express Care is at 214 El Camino Real.		
	<usr> Thank you.		
navigate	<usr> What is the weather forecast for the weekend?	For what city would you like to know that?	For what city would you like the weekend forecast for?

Table 2: Selected FSDG+LAED model’s responses

Domain	#Dialogues	Domain	#Dialogues	Domain	#Dialogues
UPDATE_CALENDAR	1991	GUINNESS_CHECK	1886	ALARM_SET	1681
SCAM_LOOKUP	1658	PLAY_TIMES	1601	GAME_RULES	1590
CONTACT_MANAGER	1483	LIBRARY_REQUEST	1339	INSURANCE	1299
HOME_BOT	1210	HOW_TO_BASIC	1086	CITY_INFO	965
TIME_ZONE	951	TOURISM	935	SHOPPING	903
BUS_SCHEDULE_BOT	898	CHECK_STATUS	784	WHAT_IS_IT	776
STORE_DETAILS	737	APPOINTMENT_REMINDER	668	PRESENT_IDEAS	664
GEOGRAPHY	653	SKI_BOT	607	MOVIE_LISTINGS	607
UPDATE_CONTACT	581	ORDER_PIZZA	577	EDIT_PLAYLIST	574
SPORTS_INFO	561	BOOKING_FLIGHT	555	WEATHER_CHECK	551
EVENT_RESERVE	539	RESTAURANT_PICKER	535	LOOK_UP_INFO	533
AUTO_SORT	514	QUOTE_OF_THE_DAY_BOT	513	WEDDING_PLANNER	510
MAKE_RESTAURANT_RESERVATIONS	510	AGREEMENT_BOT	507	NAME_SUGGESTER	499
APARTMENT_FINDER	499	HOTEL_RESERVE	497	PHONE_PLAN_BOT	496
DECIDER_BOT	487	VACATION_IDEAS	486	PHONE_SETTINGS	473
POLICY_BOT	447	PROMPT_GENERATOR	446	MUSIC_SUGGESTER	445
PET_ADVICE	426	BANK_BOT	367	CATALOGUE_BOT	288

Table 3: MetaLWOz domains

A Appendices

A.1 Training details

We train our models with the Adam optimizer (Kingma and Ba, 2014) with learning rate 0.001. Our hierarchical models’ utterance encoder is an LSTM cell (Hochreiter and Schmidhuber, 1997) of size 256, and the dialogue-level encoder is a GRU (Cho et al., 2014) of size 512.

Domain \ Statistic	Navigation	Weather	Schedule
Dialogues	800	797	8285
Utterances	5248	4314	3170
Mean length	6.56	5.41	3.83

Table 4: SMD trainset statistics

A.2 MetaLWOz dialogue example

Domain: Ski

User Task: You want to know if there are good ski hills an hour’s drive from your current location.

Bot Task: Tell the user that there are no ski hills in their immediate location.

Bot Hello how may I help you?

User Is there any good ski hills an hour’s drive from my current location?

Bot I’m sorry to inform you that there are no ski hills in your immediate location

User Can you help me find the nearest?

Bot Absolutely! It looks like you’re about 3 hours away from Bear Mountain. That seems to be the closest.

User Hmm.. sounds good

Bot Alright! I can help you get your lift tickets now! When will you be going?

User Awesome! please get me a ticket for 10pax

Bot You’ve got it. Anything else I can help you with?

User None. Thanks again!

Bot No problem!

SIM: A Slot-Independent Neural Model for Dialogue State Tracking

Chenguang Zhu

Michael Zeng

Xuedong Huang

Microsoft Speech and Dialogue Group, Redmond, WA, USA
{chezhu, nzeng, xdh}@microsoft.com

Abstract

Dialogue state tracking is an important component in task-oriented dialogue systems to identify users’ goals and requests as a dialogue proceeds. However, as most previous models are dependent on dialogue slots, the model complexity soars when the number of slots increases. In this paper, we put forward a slot-independent neural model (SIM) to track dialogue states while keeping the model complexity invariant to the number of dialogue slots. The model utilizes attention mechanisms between user utterance and system actions. SIM achieves state-of-the-art results on WoZ and DSTC2 tasks, with only 20% of the model size of previous models.

1 Introduction

With the rapid development in deep learning, there is a recent boom of task-oriented dialogue systems in terms of both algorithms and datasets. The goal of task-oriented dialogue is to fulfill a user’s requests such as booking hotels via communication in natural language. Due to the complexity and ambiguity of human language, previous systems have included semantic decoding (Mrkšić et al., 2016) to project natural language input into pre-defined dialogue states. These states are typically represented by slots and values: slots indicate the category of information and values specify the content of information. For instance, the user utterance “can you help me find the address of any hotel in the south side of the city” can be decoded as *inform(area, south)* and *request(address)*, meaning that the user has specified the value *south* for slot *area* and requested another slot *address*.

Numerous methods have been put forward to decode a user’s utterance into slot values. Some use hand-crafted features and domain-specific delexicalization methods to achieve strong performance (Henderson et al., 2014; Zilka and Jurci-

cek, 2015). Mrkšić et al. (2016) employs CNN and pretrained embeddings to further improve the state tracking accuracy. Mrkšić and Vulić (2018) extends this work by using two additional statistical update mechanisms. Liu et al. (2018) uses human teaching and feedback to boost the state tracking performance. Zhong et al. (2018) utilizes both global and local attention mechanism in the proposed GLAD model which obtains state-of-the-art results on WoZ and DSTC2 datasets. However, most of these methods require slot-specific neural structures for accurate prediction. For example, Zhong et al. (2018) defines a parametrized local attention matrix for each slot. Slot-specific mechanisms become unwieldy when the dialogue task involves many topics and slots, as is typical in a complex conversational setting like product troubleshooting. Furthermore, due to the sparsity of labels, there may not be enough data to thoroughly train each slot-specific network structure. Rastogi et al. (2017); Ramadan et al. (2018) both propose to remove the model’s dependency on dialogue slots but there’s no modification to the representation part, which could be crucial to textual understanding as we will show later.

To solve this problem, we need a state tracking model independent of dialogue slots. In other words, the network should depend on the semantic similarity between slots and utterance instead of slot-specific modules. To this end, we propose the Slot-Independent Model (SIM). Our model complexity does *not* increase when the number of slots in dialogue tasks go up. Thus, SIM has many fewer parameters than existing dialogue state tracking models. To compensate for the exclusion of slot-specific parameters, we incorporate better feature representation of user utterance and dialogue states using syntactic information and convolutional neural networks (CNN). The refined representation, in addition to cross and self-

attention mechanisms, make our model achieve even better performance than slot-specific models. For instance, on Wizard-of-Oz (WOZ) 2.0 dataset (Wen et al., 2016), the SIM model obtains a joint-accuracy score of 89.5%, 1.4% higher than the previously best model GLAD, with only 22% of the number of parameters. On DSTC2 dataset, SIM achieves comparable performance with previous best models with only 19% of the model size.

2 Problem Formulation

As outlined in Young et al. (2010), the dialogue state tracking task is formulated as follows: at each turn of dialogue, the user’s utterance is semantically decoded into a set of slot-value pairs. There are two types of slots. *Goal* slots indicate the category, e.g. area, food, and the values specify the constraint given by users for the category, e.g. South, Mediterranean. *Request* slots refer to requests, and the value is the category that the user demands, e.g. phone, area. Each user’s turn is thus decoded into *turn goals* and *turn requests*. Furthermore, to summarize the user’s goals so far, the union of all previous turn goals up to the current turn is defined as *joint goals*.

Similarly, the dialogue system’s reply from the previous round is labeled with a set of slot-value pairs denoted as *system actions*. The dialogue state tracking task requires models to predict turn goal and turn request given user’s utterance and system actions from previous turns.

Formally, the *ontology* of dialogue, O , consists of all possible slots S and the set of values for each slot, $V(s), \forall s \in S$. Specifically, *req* is the name for *request* slot and its values include all the requestable category information. The dialogue state tracking task is that, given the user’s utterance in the i -th turn, U , and system actions from the $(i - 1)$ -th turn, $A = \{(s_1, v_1), \dots, (s_q, v_q)\}$, where $s_j \in S, v_j \in V(s_j)$, the model should predict:

1. Turn goals: $\{(s_1, v_1), \dots, (s_b, v_b)\}$, where $s_j \in S, v_j \in V(s_j)$,
2. Turn requests: $\{(req, v_1), \dots, (req, v_t)\}$, where $v_j \in V(req)$.

The joint goals at turn i are then computed by taking the union of all the predicted turn goals from turn 1 to turn i .

Usually this prediction task is cast as a binary classification problem: for each slot-value

pair (s, v) , determine whether it should be included in the predicted turn goals/requests. Namely, the model is to learn a mapping function $f(U, A, (s, v)) \rightarrow \{0, 1\}$.

3 Slot-Independent Model

To predict whether a slot-value pair should be included in the turn goals/requests, previous models (Mrkšić et al., 2016; Zhong et al., 2018) usually define network components for each slot $s \in S$. This can be cumbersome when the ontology is large, and it suffers from the insufficient data problem: the labelled data for a single slot may not suffice to effectively train the parameters for the slot-specific neural networks structure.

Therefore, we propose that in the classification process, the model needs to rely on the semantic similarity between the user’s utterance and slot-value pair, with system action information. In other words, the model should have only a single global neural structure independent of slots. We heretofore refer to this model as Slot-Independent Model (SIM) for dialogue state tracking.

3.1 Input Representation

Suppose the user’s utterance in the i -th turn contains m words, $U = (w_1, w_2, \dots, w_m)$. For each word w_i , we use GloVe word embedding e_i , character-CNN embedding c_i , Part-Of-Speech (POS) embedding POS_i , Named-Entity-Recognition (NER) embedding NER_i and exact match feature EM_i . The POS and NER tags are extracted by spaCy and then mapped into a fixed-length vector. The exact matching feature has two bits, indicating whether a word and its lemma can be found in the slot-value pair representation, respectively. This is the first step to establish a semantic relationship between user utterance and slots. To summarize, we represent the user utterance as $X^U = \{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_m\} \in \mathbb{R}^{m \times d_u}$, $\mathbf{u}_i = [e_i; c_i; POS_i; NER_i; EM_i]$.

For each slot-value pair (s, v) either in system action or in the ontology, we get its text representation by concatenating the contents of slot and value¹. We use GloVe to embed each word in the text. Therefore, each slot-value pair in system actions is represented as $X^A \in \mathbb{R}^{a \times d}$ and each slot-value pair in ontology is represented as

¹To align with previous work, we prepend the word “inform” to goal slot.

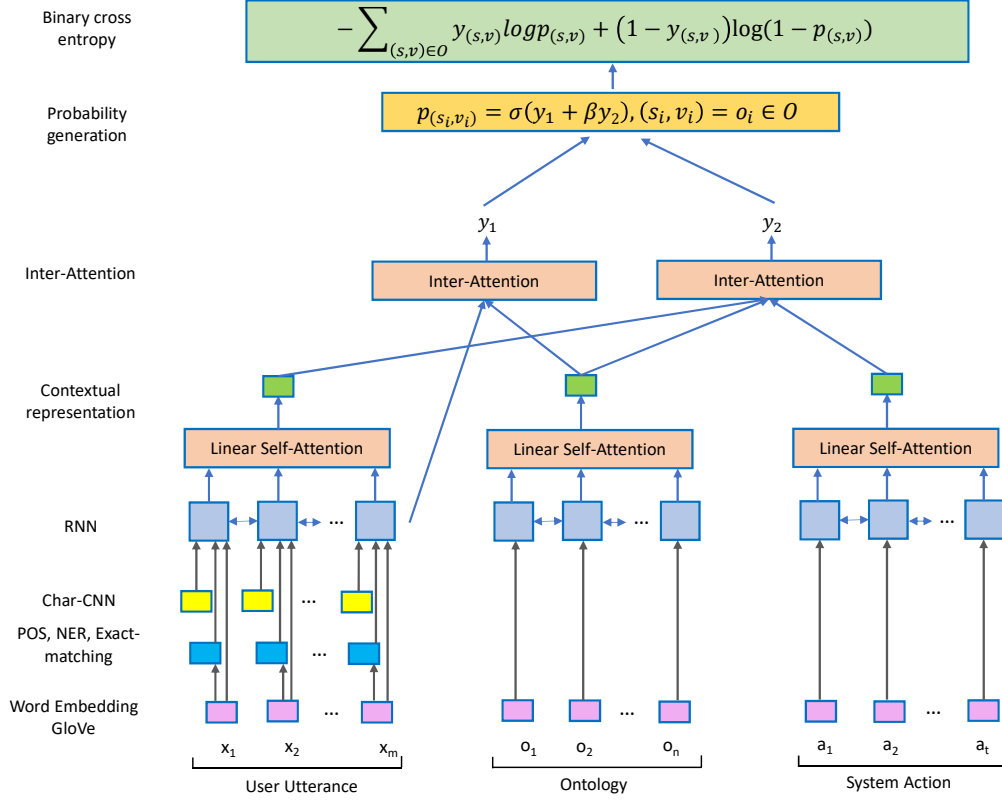


Figure 1: SIM model structure.

$X^O \in \mathbb{R}^{o \times d}$, where a and o is the number of words in the corresponding text.

3.2 Contextual Representation

To incorporate contextual information, we employ a bi-directional RNN layer on the input representation. For instance, for user utterance,

$$R^U = \text{BiLSTM}(X^U) \in \mathbb{R}^{m \times d_{rnn}} \quad (1)$$

We apply variational dropout (Kingma et al., 2015) for RNN inputs, i.e. the dropout mask is shared over different timesteps.

After RNN, we use linear self-attention to get a single summarization vector for user utterance, using weight vector $w \in \mathbb{R}^{d_{rnn}}$ and bias scalar b :

$$\alpha = R^U w + b \in \mathbb{R}^m \quad (2)$$

$$p = \text{softmax}(\alpha) \in \mathbb{R}^m \quad (3)$$

$$s^U = (R^U)^T p \in \mathbb{R}^{d_{rnn}} \quad (4)$$

For each slot-value pair in the system actions and ontology, we conduct RNN and linear self-attention summarization in a similar way. As the slot-value pair input is not a sentence, we only keep the summarization vector $s^A \in \mathbb{R}^{d_{rnn}}$ and $s^O \in \mathbb{R}^{d_{rnn}}$ for each slot-value pair in system actions and ontology respectively.

3.3 Inter-Attention

To determine whether the current user utterance refers to a slot-value pair (s, v) in the ontology, the model employs inter-attention between user utterance, system action and ontology. Similar to the framework in Zhong et al. (2018), we employ two sources of interactions.

The first is the semantic similarity between the user utterance, represented by embedding R^U and each slot-value pair from ontology (s, v) , represented by embedding s^O . We linearly combine vectors in R^U via the normalized inner product with s^O , which is then employed to compute the similarity score y_1 :

$$\alpha = R^U s^O \in \mathbb{R}^m \quad (5)$$

$$p_1 = \text{softmax}(\alpha) \in \mathbb{R}^m \quad (6)$$

$$q_1 = (R^U)^T p_1 \in \mathbb{R}^{d_{rnn}} \quad (7)$$

$$y_1 = w_1^T q_1 + b_1 \in \mathbb{R} \quad (8)$$

The second source involves the system actions. The reason is that if the system requested certain information in the previous round, it is very likely that the user will give answer in this round, and the answer may refer to the question, e.g. “yes” or

“no” to the question. Thus, we first attend to system actions from user utterance and then combine with the ontology to get similarity score. Suppose there are L slot-values pairs in the system actions from previous round², represented by s_1^A, \dots, s_L^A :

$$p_2 = \text{softmax}(\{s_j^{AT} s^U\}_{j=1}^L) \in \mathbb{R}^L \quad (9)$$

$$q_2 = \sum_{j=1}^L p_j s_j^A \in \mathbb{R}^{d_{rnn}} \quad (10)$$

$$y_2 = q_2^T s^O \in \mathbb{R} \quad (11)$$

The final similarity score between the user utterance and a slot-value pair (s, v) from the ontology is a linear combination of y_1 and y_2 and normalized using sigmoid function.

$$p_{(s,v)} = \sigma(y_1 + \beta y_2) \in \mathbb{R}, \quad (12)$$

where β is a learned coefficient. The loss function is the sum of binary cross entropy over all slot-value pairs in the ontology:

$$L(\theta) = - \sum_{(s,v) \in O} y_{(s,v)} \log p_{(s,v)} + \quad (13)$$

$$(1 - y_{(s,v)}) \log(1 - p_{(s,v)}), \quad (14)$$

where $y_{(s,v)} \in \{0, 1\}$ is the ground truth. We illustrate the model structure of SIM in Figure 1.

4 Experiment

4.1 Dataset

We evaluated our model on Wizard of Oz (WoZ) (Wen et al., 2016) and the second Dialogue System Technology Challenges (Williams et al., 2013). Both tasks are for restaurant reservation and have slot-value pairs of both goal and request types. WoZ has 4 kinds of slots (*area, food, price range, request*) and 94 values in total. DSTC2 has an additional slot *name* and 220 values in total. WoZ has 800 dialogues in the training and development set and 400 dialogues in the test set, while DSTC2 dataset consists of 2118 dialogues in the training and development set, and 1117 dialogues in the test set.

4.2 Metrics

We use accuracy on the joint goal and turn request as the evaluation metrics. Both are sets of

²This includes a special sentinel action which refers to ignoring the system action.

slot-value pairs, so the predicted set must exactly match the answer to be judged as correct. For joint goals, if a later turn generates a slot-value pair where the slot has been specified in previous rounds, we replace the value with the latest content.

4.3 Training Details

We fix GloVe (Pennington et al., 2014) as the word embedding matrix. The models are trained using ADAM optimizer (Kingma and Ba, 2014) with an initial learning rate of 1e-3. The dimension of POS and NER embeddings are 12 and 8, respectively. In character-CNN, each character is embedded into a vector of length 50. The CNN window size is 3 and hidden size is 50. We apply a dropout rate of 0.1 for the input to each module. The hidden size of RNN is 125.

During training, we pick the best model with highest joint goal score on development set and report the result on the test set.

For DSTC2, we adhere to the standard procedure to use the N-best list from the noisy ASR results for testing. The ASR results are very noisy. We experimented with several strategies and ended up using only the top result from the N-best list. The training and validation on DSTC2 are based on noise-free user utterance. The WoZ task does not have ASR results available, so we directly use noise-free user utterance.

4.4 Baseline models and result

We compare our model SIM with a number of baseline systems: delexicalization model (Wen et al., 2016; Henderson et al., 2014), the neural belief tracker model (NBT) (Mrkšić et al., 2016), global-locally self-attentive model GLAD (Zhong et al., 2018), large-scale belief tracking model LSBT (Ramadan et al., 2018) and scalable multi-domain dialogue state tracking model SMDST (Rastogi et al., 2017).

Table 1 shows that, on WoZ dataset, SIM achieves a new state-of-the-art joint goal accuracy of 89.5%, a significant improvement of 1.4% over GLAD, and turn request accuracy of 97.3%, 0.2% above GLAD. On DSTC2 dataset, where noisy ASR results are used as user utterance during test, SIM obtains comparable results with GLAD. Furthermore, the better representation in SIM makes it significantly outperform previous slot-independent models LSBT and SMDST.

Model	WoZ		DSTC2	
	Joint goal	Turn request	Joint goal	Turn request
SMDST	/	/	70.3%	/
Delex. Model + Semantic Dictionary	83.7%	87.6%	72.9%	95.7%
Neural Belief Tracker (NBT)	84.2%	91.6%	73.4%	96.5%
LSBT	85.5%	/	/	/
GLAD	88.1%	97.1%	74.5%	97.5%
SIM	89.5%	97.3%	74.7%	96.2%

Table 1: Joint goal and turn request accuracies on WoZ and DSTC2 restaurant reservation datasets.

Furthermore, as SIM has no slot-specific neural network structures, its model size is much smaller than previous models. Table 2 shows that, on WoZ and DSTC2 datasets, SIM model has the same number of parameters, which is only 23% and 19% of that in GLAD model.

Ablation Study. We conduct an ablation study of SIM on WoZ dataset. As shown in Table 3, the additional utterance word features, including character, POS, NER and exact matching embeddings, can boost the performance by 2.4% in joint goal accuracy. These features include POS, NER and exact match features. This indicates that for the dialogue state tracking task, syntactic information and text matching are very useful. Character-CNN captures sub-word level information and is effective in understanding spelling errors, hence it helps with 1.2% in joint goal accuracy. Variational dropout is also beneficial, contributing 0.9% to the joint goal accuracy, which shows the importance of uniform masking during dropout.

Model	WoZ	DSTC2
SIM	1.47M	1.47M
GLAD (Zhong et al., 2018)	6.41M	7.69M

Table 2: Model size comparison between SIM and GLAD (Zhong et al., 2018) on WoZ and DSTC2.

Model	Joint Goal	Turn Request
SIM	89.5	97.3
-Var. dropout	88.6	97.1
-Char. CNN	88.3	97.0
-Utt. features	87.1	97.1

Table 3: Ablation study of SIM on WoZ. We pick the model with highest joint goal score on development set and report its performance on test set.

5 Conclusion

In this paper, we propose a slot-independent neural model, SIM, to tackle the dialogue state tracking problem. Via incorporating better feature representations, SIM can effectively reduce the model complexity while still achieving superior or comparable results on various datasets, compared with previous models.

For future work, we plan to design general slot-free dialogue state tracking models which can be adapted to different domains during inference time, given domain-specific ontology information. This will make the model more agile in real applications.

Acknowledgement

We thank the anonymous reviewers for the insightful comments. We thank William Hinthorn for proof-reading our paper.

References

- Matthew Henderson, Blaise Thomson, and Steve Young. 2014. Word-based dialog state tracking with recurrent neural networks. In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 292–299.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Diederik P Kingma, Tim Salimans, and Max Welling. 2015. Variational dropout and the local reparameterization trick. In *Advances in Neural Information Processing Systems*, pages 2575–2583.
- Bing Liu, Gokhan Tur, Dilek Hakkani-Tur, Pararth Shah, and Larry Heck. 2018. Dialogue learning with human teaching and feedback in end-to-end trainable task-oriented dialogue systems. *arXiv preprint arXiv:1804.06512*.

- Nikola Mrkšić, Diarmuid O Séaghdha, Tsung-Hsien Wen, Blaise Thomson, and Steve Young. 2016. Neural belief tracker: Data-driven dialogue state tracking. *arXiv preprint arXiv:1606.03777*.
- Nikola Mrkšić and Ivan Vulić. 2018. Fully statistical neural belief tracking. *arXiv preprint arXiv:1805.11350*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Osman Ramadan, Paweł Budzianowski, and Milica Gašić. 2018. Large-scale multi-domain belief tracking with knowledge sharing. *arXiv preprint arXiv:1807.06517*.
- Abhinav Rastogi, Dilek Hakkani-Tür, and Larry Heck. 2017. Scalable multi-domain dialogue state tracking. In *Automatic Speech Recognition and Understanding Workshop (ASRU), 2017 IEEE*, pages 561–568. IEEE.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrksic, Milica Gasic, Lina M Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2016. A network-based end-to-end trainable task-oriented dialogue system. *arXiv preprint arXiv:1604.04562*.
- Jason Williams, Antoine Raux, Deepak Ramachandran, and Alan Black. 2013. The dialog state tracking challenge. In *Proceedings of the SIGDIAL 2013 Conference*, pages 404–413.
- Steve Young, Milica Gašić, Simon Keizer, François Mairesse, Jost Schatzmann, Blaise Thomson, and Kai Yu. 2010. The hidden information state model: A practical framework for pomdp-based spoken dialogue management. *Computer Speech & Language*, 24(2):150–174.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2018. Global-locally self-attentive dialogue state tracker. *arXiv preprint arXiv:1805.09655*.
- Lukas Zilka and Filip Jurcicek. 2015. Incremental lstm-based dialog state tracker. In *Automatic Speech Recognition and Understanding (ASRU), 2015 IEEE Workshop on*, pages 757–762. IEEE.

Simple, Fast, Accurate Intent Classification and Slot Labeling for Goal-Oriented Dialogue Systems

Arshit Gupta*

Amazon AI
Seattle

arshig@amazon.com

John Hewitt*†

Stanford University
Palo Alto

johnhew@stanford.edu

Katrin Kirchhoff

Amazon AI
Seattle

katrinki@amazon.com

Abstract

With the advent of conversational assistants like Amazon Alexa, Google Now, etc., dialogue systems are gaining a lot of traction, especially in industrial settings. These systems typically include a Spoken Language understanding component which consists of two tasks: Intent Classification (IC) and Slot Labeling (SL). Generally, these two tasks are modeled together jointly to achieve best performance. However, this joint modeling adds to model obfuscation. In this work, we first design framework for a modularization of joint IC-SL task to enhance architecture transparency. Then, we explore a number of self-attention, convolutional, and recurrent models, contributing a large-scale analysis of modeling paradigms for IC+SL across two datasets. Finally, using this framework, we propose a class of ‘label-recurrent’ models that are non-recurrent apart from a 10-dimensional representation of the label history, and show that our proposed systems are highly accurate (achieving over 30% error reduction in SL over the state-of-the-art on the Snips dataset), as well as fast, at 2x the inference and 2/3 to 1/2 the training time of comparable recurrent models, thus giving an edge in critical real-world systems.

1 Introduction

At the core of task-oriented dialogue systems are spoken language understanding (SLU) models, tasked with determining the intent of users’ utterances and labeling semantically relevant words at each turn of the conversation. Performance on these tasks, known as intent classification (IC) and slot labeling (SL), upper-bounds the utility of such dialogue systems. A large body of recent research has improved these models through the use of recurrent neural networks, encoder-decoder architectures, and attention mechanisms. However, for

*Equal Contribution

†Work performed while at Amazon AI




Figure 1: A general framework of joint IC+SL, decoupling modeling tasks to permit the analysis of each component independently.

production dialogue systems in particular, system speed is at a premium, both during training and in real-time inference.

In this work, we propose fully non-recurrent and label-recurrent model paradigms including self-attention and convolution for comparison to state-of-the-art recurrent models in terms of accuracy and speed. To achieve this, we design a framework for joint IC-SL models that is modularized into different components and makes the task agnostic to type of neural network used. This, in turn, makes the model architecture simpler, easy to understand and renders the task network agnostic, allowing for easier plug and play using existing components, such as pre-trained contextual word embeddings (Devlin et al., 2019), etc. This is essential for easier model debugging and quicker experimentation, especially in industrial settings.

Using this framework, we identify three distinct model families of interest: fully recurrent,

label-recurrent, and non-recurrent. Recent state-of-the-art models fall into the first category, as encoder-decoder architectures have recurrent encoders to perform word context encoding, and predict slot label sequences using recurrent decoders that use both word and label information as they decode (Hakkani-Tür et al., 2016; Liu and Lane, 2016; Li et al., 2018). In second category, we have ‘non-recurrent’ networks: fully feed-forward, attention-based, or convolutional models, for example. Lastly, we have a class of label-recurrent models, inspired by structured sequential models like conditional random fields on top of non-recurrent word contextualization components. In this class of models, slot label decoding proceeds such that label sequences are encoded by a recurrent component, but word sequences are not.

Our contributions are:

- A class of label-recurrent convolutional models that achieve state-of-the-art performance on Snips and competitive performance on ATIS while maintaining faster training and inference speeds than fully-recurrent models
- A new modular framework for joint IC-SL models that permits the analysis of individual modeling components that decomposes these joint models into separate components for *word context encoding*, *summarization of the sentence* into a single vector for intent classification, and *modeling of dependencies in the output space* of slot label sequences.
- In-depth analysis of different word contextualizations for Spoken Language Understanding task (for instance, providing evidence for the intuition that explicitly focusing on local context is a useful architectural inductive prior for slot labeling)

2 Prior Work

There is a large body of research in applying recurrent modeling advances to intent classification and slot labeling (frequently called spoken language understanding). Traditionally, for intent classification, word n-grams were used with SVM classifier (Haffner et al., 2003) and Adaboost (Schapire and Singer, 2000). For the SL task, CRFs (Gorin et al., 1997) have been used in the past.

Recently, a larger focus has been on joint modeling of IC and SL tasks. Long short-term memory recurrent neural networks (Hochreiter and

Schmidhuber, 1997) and Gated Recurrent Unit models (Cho et al.) were proposed for slot labeling by Yao et al. (2014) and Zhang and Wang (2016) respectively, while Guo et al. (2014) used recursive neural networks. Subsequent improvements to recurrent neural modeling techniques, like bidirectionality and attention (Bahdanau et al., 2014) were incorporated into IC+SL in recent years as well (Hakkani-Tür et al., 2016; Liu and Lane, 2016). Li et al. (2018) introduced a self-attention based joint model where they used self-attention and LSTM layers along with the gating mechanism for this task.

Non-recurrent modeling for language has been re-visited recently, even as recurrent techniques continue to be dominant. Dilated CNNs (Yu and Koltun, 2015) with CRF label modeling were applied to named entity recognition by Strubell et al. (2017), and earlier were applied to SL by Xu and Sarikaya (2013). Convolutional and attention-based sentence encoders have been applied in complex tasks, including machine translation, natural language inference, and parsing. (Gehring et al., 2017; Vaswani et al., 2017; Shen et al., 2017; Kitaev and Klein, 2018) We draw from both of these bodies of work to propose a simple yet highly effective family of IC+SL models.

3 A general framework of joint IC+SL

Intent classification and slot labeling take as input an utterance $\mathbf{x}_{1:T} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$, composed of words \mathbf{x}_i and of length T . Models construct a distribution over intents and slot label sequences given the utterance. One intent is assigned per utterance and one slot label is assigned per word:

$$P(l_{1:T}, c | \mathbf{x}_{1:T}) \quad (1)$$

where $c \in \mathcal{I}$, a fixed set of intents, and $l_i \in \mathcal{L}$, a fixed set of slot labels. Models are trained to minimize the cross-entropy loss between the assigned distribution and the training data. To the end of constructing this distribution, our framework explicitly separates the following components, which are explicitly or implicitly present in all joint IC+SL systems (Figure 1):

3.1 Word contextualization

We first assume words are encoded through an embedding layer, providing context-independent word vectors. Overloading notation, we denote the embedded sequence $\mathbf{x}_{1:T}$, with $\mathbf{x}_i \in \mathbb{R}^{d_x}$.

In this component, word representations are enriched with sentential context. Each word \mathbf{x}_i is assigned a contextualized representation \mathbf{h}_i . To ease layering these components, we keep the dimensionality the same as the word embeddings; $\mathbf{h}_i \in \mathbb{R}^{d_x}$. Our study consists mainly of varying this component across models, which are described in detail in Section 4. In all models, we assume independence of intent classification and slot labeling given the learned representations:

$$P(l_{1:T}, c | \mathbf{h}_{1:T}) = P(l_{1:T} | \mathbf{h}_{1:T}) P(c | \mathbf{h}_{1:T}) \quad (2)$$

3.2 Sentence representation

In this component, the output of the word contextualization component is summarized in a single vector,

$$\mathbf{s} = \text{SentenceRepr}(\mathbf{h}_{1:T}) \quad (3)$$

where $\mathbf{s} \in \mathbb{R}^{d_x}$. For all our experiments, we keep this component constant, using a simple attention-like pooling which is the weighted sum of word contextualization for each position in the sentence. These weights are computed using softmax over these individual word contextualizations

While simple, this model permits word contextualization components freedom in how they encode sentential information; for example, self-attention models may spread full-sentence information across all words, whereas 1-directional LSTMs may focus full-sentence information in the last word’s vector.

3.3 Intent prediction

In this component, the sentence representation is used as features to predict the intent of the utterance. For all experiments, we keep this component fixed as well, using a simple two-layer feed-forward block on top of \mathbf{s} .

3.4 Slot label prediction

In this component, the output of the word contextualization component is used to construct a distribution over slot label sequences for the utterance. We decompose the joint probability of the label sequence given the contextualized word representations into a left-to-right labeling:

$$P(l_{1:T} | \mathbf{h}_{1:T}) = \prod_{i=1}^T P(l_i | \mathbf{h}_{1:T}, l_{1:i-1}) \quad (4)$$

In our experiments, we explore two models for slot prediction, one fully-parallelizable because of strong independence assumptions, the other permitting a constrained dependence between labeling decisions that we call ‘label-recurrent’.

Independent slot prediction The first is a non-recurrent model, which assumes independence between all labeling decisions once given $\mathbf{h}_{1:T}$, as well as independence from all word representations except that of the word being labeled:

$$P(l_i | \mathbf{h}_{1:T}, l_{1:i-1}) = P(l_i | \mathbf{h}_i) \quad (5)$$

This model is fully parallelizable on GPU architectures, and the probability of each labeling decision is modeled according to

$$P(l_i | \mathbf{h}_{1:T}) = \text{softmax}(W^{(3)} \mathbf{p}_{i,SL} + \mathbf{b}^{(3)}) \quad (6)$$

$$\mathbf{p}_{i,SL} = \tanh(W^{(4)} \mathbf{h}_i + \mathbf{b}^{(4)}) \quad (7)$$

hence, SL prediction features are learned using each contextualized word independently.

Label-recurrent slot prediction The second class of slot prediction models we consider lead to our classification, ‘label-recurrent.’¹ These models permit dependence of labeling decisions on the sequence of decisions made so far, but keep the independence assumption on the word representations:

$$P(l_i | \mathbf{h}_{1:T}, l_{1:i-1}) = P(l_i | l_{1:i-1}, \mathbf{h}_i) \quad (8)$$

Notably, this family of models excludes traditional encoder-decoder models, since the decoder component uses labeling decisions $l_{1:i-1}$ and earlier word representations $\mathbf{h}_{1:i-1}$ to influence the predictor features $\mathbf{p}_{i,SL}$. However, it includes models such as CNN-CRF.

The space of label sequences in slot labeling is much smaller than the space of word sequences. This adds minimal computational burden and permits the model to benefit from GPU parallelism during $\mathbf{h}_{1:T}$ computation.

For our experiments, we propose a single label-recurrent model, which encodes labeling histories $l_{1:i}$ using only a 10-dimensional LSTM. First, slot labels are embedded, such that for each $l \in \mathcal{L}$, we have $\mathbf{l} \in \mathbb{R}^{d_l}$. An initial tag history state, h_0^{tag} , is randomly initialized. Each tag decision is fed

¹We use this term for clarity in language, not to claim that no such models have been explored in the past.

along with the previous tag history state to the LSTM, which returns the next tag history state:

$$\mathbf{h}_i^{\text{tag}} = \text{LSTM}(\mathbf{l}_{i-1}, \mathbf{h}_{i-1}^{\text{tag}}). \quad (9)$$

We omit a precise description of the LSTM model here for space, referring the reader to (Hochreiter and Schmidhuber, 1997).

The tag history is used at each prediction step as additional inputs to construct the predictor features $\mathbf{p}_{i,\text{SL}}$, replacing Eqn. 7 with:

$$\mathbf{p}_{i,\text{SL}} = \tanh(W^{(5)}[\mathbf{h}_i; \mathbf{h}_i^{\text{tag}}] + \mathbf{b}^{(5)}) \quad (10)$$

where $[a; b]$ denotes concatenation. This model and other label-recurrent models are not only parallelizable more than fully-recurrent models, but also provide an architectural inductive bias, separating modeling of tag sequences from modeling of word sequences. In our experiments, we perform greedy decoding to maintain a high decoding speed.

4 Word contextualization models

In this section, we describe word contextualization models with the goal of identifying non-recurrent architectures that achieve high accuracy and faster speed than recurrent models.

4.1 Feed-forward model

In this model, we set $\mathbf{h}_{1:T} = \mathbf{x}_{1:T} + \mathbf{a}_{1:T}$, where $\mathbf{a}_{1:T}$ is a learned absolute position representation, with one vector learned per absolute position, as used in (Gehring et al., 2017). While extremely simple, this model provides a useful baseline as a totally context-free model. It also permits us to analyze the contribution of a label-recurrent component in such a context-deprived scenario.

4.2 Self-attention models

Recent work in non-recurrent modeling has surfaced a number of variants of attention-based word context modeling.

The simplest constructs each \mathbf{h}_i by incorporating a weighted average of the rest of the sequence, $\mathbf{x}_{1:T} \setminus \mathbf{x}_i$. We use a bilinear attention mechanism with a residual connection while masking out the

identity in the attention weights.

$$\mathbf{h}_i = \text{relu}(\sqrt{.5}(\mathbf{c}_i + \mathbf{x}_i)) \quad (11)$$

$$\mathbf{c}_i = \sum_{j=1, j \neq i}^T \alpha_j \mathbf{x}_j \quad (12)$$

$$\alpha_j = \frac{\exp(\mathbf{x}_i^T W^{(5)} \mathbf{x}_j)}{\sum_{j'=1}^T \exp(\mathbf{x}_i^T W^{(5)} \mathbf{x}_{j'})} \quad (13)$$

In this and all subsequent models, we optionally stack multiple layers, feeding the word representations from each layer into the next; in this case we denote the models ATTN-1L, ATTN-2L, etc.

We also analyze multi-head attention models, drawing from (Vaswani et al., 2017). For a model with k heads, we construct one matrix of the form $A \in \mathbb{R}^{d_x/k}$ for each head, and transform each \mathbf{x}_i , $\mathbf{x}_i^{k'} = A^{k'} \mathbf{x}_i$ for $k' \in \{1, \dots, k\}$. These are passed into the attention equations above, generating context vectors $\mathbf{c}_i^1, \dots, \mathbf{c}_i^k \in \mathbb{R}^{d_x/k}$, which are then concatenated to form a vector in \mathbb{R}^{d_x} . These context layers are usually sent through a linear transformation to combine features between the heads, the output of which is \mathbf{c}_i , but we found that omitting this combination transformation leads to significantly improved results, so we do so in all experiments. We denote these models K-HEAD ATTN.

4.2.1 Relative position representations

We found in early experiments that the absolute position embeddings in self-attention models are insufficient for representing order. Hence, in all attention models except when explicitly noted, we use relative position representations as follows. We follow Shaw et al. (2018), who improved the absolute position representations of the Transformer model (Vaswani et al., 2017) by learning vector representations of relative positions and incorporating them into the self-attention mechanism as follows:

$$\mathbf{c}_i = \sum_{i'=1, i' \neq i}^T \alpha_{i'} (\mathbf{x}_{i'} + \mathbf{v}_{f(i,i')}) \quad (14)$$

$$\alpha_j = \frac{\exp(\mathbf{x}_i^T W^{(5)} \mathbf{x}_j + b_{f(i,j)})}{\sum_{j'=1}^T \exp(\mathbf{x}_i^T W^{(5)} \mathbf{x}_{j'} + b_{f(i,j)})} \quad (15)$$

where $\mathbf{v}_{f(i,j)}$ is a learned vector representing how the relative positions i and j should be incorporated, and $b_{f(i,j)}$ is a learned bias that determines how the relative position should affect the

weight given to position j when contextualizing position i . The function f determines which relative positions to group together with a single relative position vector. Given the generally small datasets in IC+SL, we use the following relative position function, which buckets relative positions together in exponentially larger groups as distance increases, following the results of [Khandelwal et al. \(2018\)](#), that LSTMs represent position fuzzily at long relative distances.

$$f(i, j) = \begin{cases} \pm 1, |j - i| = 1 \\ \pm 2, |j - i| \in \{2, 3\} \\ \pm 3, |j - i| \in \{4..7\} \\ \dots \end{cases} \quad (16)$$

This is similar to the preprint of [Bilan and Roth \(2018\)](#), who use linearly increasing bucket sizes; we found exponentially increasing sizes to work well compared to the constant bucket sizes of [Shaw et al. \(2018\)](#).

4.3 Convolutional models

Convolution incorporates local word context into word representations, where kernel width parameter specifies the total size (in words) of local context considered. Each convolutional layer produces a vector representation of each word,

$$\mathbf{h}_{1:T} = \text{relu}(\sqrt{.5} * [\text{CNN}(\mathbf{x}_{1:T}) + \mathbf{x}_{1:T}]) \quad (17)$$

and includes a residual connection, and variance normalization, following ([Gehring et al., 2017](#)). To maintain the dimensionality of \mathbf{h}_i as \mathbb{R}^{d_x} , we use a filter count of d_x . We vary the number of CNN layers as well as the kernel width, and for all models use a variant known as dilated CNNs. These CNNs incorporate distant context into word representations by skipping an increasing number of nearby words in each subsequent convolutional pass. We use an exponentially increasing dilation size; in the first layer, words of distance 1 are incorporated; at layer two, words of distance 2, then 4, etc. This permits large contexts to be incorporated into word representations while keeping kernel sizes and the number of layers low.

4.4 Recurrent models

We also construct a recurrent word contextualization model, more or less identical to encoders of recent state-of-the-art models. We use a bidirectional LSTM to encode word contexts, $\mathbf{h}_{1:T} =$

$\text{BiLSTM}(\mathbf{x}_{1:T})$. As with all other models, we report the performance of this model with feed-forward slot label prediction as well as with label-recurrent slot label prediction. Though similar to earlier work, both models are new; though the latter is recurrent both in word contextualization and slot label prediction, it is distinct from past models in that the two recurrent components are completely decoupled until the prediction step.

5 Datasets

We evaluate our framework and models on the ATIS data set ([Hemphill et al., 1990](#)) of spoken airline reservation requests and the Snips NLU Benchmark set ([Coucke et al., 2018](#)). The ATIS training set contains 4978 utterances from the ATIS-2 and ATIS-3 corpora; the test set consists of 893 utterances from the ATIS-3 NOV93 and DEC94 data sets. The number of slot labels is 127, and the number of intent classes is 18. Only the words themselves are used as input; no additional tags are used.

The Snips 2017 dataset is a collection of 16K crowdsourced queries, with about 2400 utterances per each of 7 intents. These intents range from ‘Play Music’ to ‘Get Weather’. Training data contains 13784 utterances and the test data consists of 700 utterances. The utterance tokens are mixed case unlike the ATIS dataset, where all the tokens are lowercased. Total number of slot labels are 72. We use IOB tagging, and split 10% of the train set off to form a development set. Utterances in Snips are, on average, short, with 9.15 words per utterance compared to ATIS’ 11.2. However, slot label sequences themselves are longer in Snips, averaging 1.8 tokens per span to ATIS’ 1.2, making span-level slot labeling more difficult. For our development experiments, we use the casing and tokenization provided by Snips. Co, but to compare to prior work, in one test experiment we use the lowercased, tokenized version of ([Goo et al., 2018](#))².

6 Experiments

We evaluate multiple models from each of our model paradigms to help determine what modeling structures are necessary for SLU, and where the best accuracy-speed tradeoffs are. First, we report extensive evaluation across the Snips and ATIS development sets, tracking inference speed and time to convergence along with the usual IC

²<https://github.com/MiuLab/SlotGated-SLU>

Model	label recurrent	IC acc		SL F1		Inference ms/utterance	Epochs to converge	s/epoch	#
		Snips	ATIS	Snips	ATIS				
FEED-FORWARD	No	98.56	97.14	53.59	69.68	0.61	48	1.82	17k
FEED-FORWARD	Yes	98.54	97.46	75.35	88.72	1.82	83	2.52	19k
CNN, 5KERNEL, 1L	No	98.56	98.40	85.88	94.11	0.82	23	1.90	42k
CNN, 5KERNEL, 3L	No	99.04	98.42	92.21	96.68	1.37	55	2.16	91k
CNN, 3KERNEL, 4L	No	98.81	98.32	91.65	96.75	1.28	57	2.29	76k
CNN, 5KERNEL, 1L	Yes	98.85	98.36	93.12	96.39	2.13	51	2.77	43k
CNN, 5KERNEL, 3L	Yes	99.10	98.36	94.22	96.95	2.68	59	3.34	93k
CNN, 3KERNEL, 4L	Yes	98.96	98.32	93.71	96.95	2.60	53	3.43	78k
ATTN, 1HEAD, 1L, NO-POS	No	98.50	97.51	53.61	69.31	1.95	25	1.94	22k
ATTN, 1HEAD, 1L	No	98.53	97.74	75.55	93.22	4.75	117	4.34	23k
ATTN, 1HEAD, 3L	No	98.74	98.10	81.51	94.07	7.68	160	4.32	33k
ATTN, 2HEAD, 3L	No	98.31	98.10	83.02	94.61	7.86	79	4.87	47k
ATTN, 1HEAD, 1L, NO POS	Yes	98.63	97.68	74.94	88.60	3.24	60	2.66	24k
ATTN, 1HEAD, 1L	Yes	98.61	98.00	86.72	94.53	6.12	89	5.53	24k
ATTN, 1HEAD, 3L	Yes	98.51	98.26	88.04	94.99	9.03	109	6.06	34k
ATTN, 2HEAD, 3L	Yes	98.48	98.26	89.31	95.86	9.17	93	6.54	49k
LSTM, 1L	No	98.82	98.34	91.83	97.28	2.65	45	2.91	47k
LSTM, 2L	No	98.77	98.20	93.10	97.36	4.72	58	5.09	77k
LSTM, 1L	Yes	98.68	98.36	93.83	97.37	3.98	54	4.62	49k
LSTM, 2L	Yes	98.71	98.30	93.88	97.28	6.03	69	6.82	79k

Table 1: Development results on the Snips 2017 and ATIS datasets, comparing models from feed-forward, convolutional, self-attention, and recurrent paradigms, as well as comparing non-recurrent, label-recurrent, and fully recurrent architectures, on IC, SL, inference speed, and training time. Inference speed, convergence time, and parameter count are drawn from Snips experiments, but the trends hold on ATIS. The best IC and SL for each dataset is bolded within each model paradigm to help compare between paradigms.

accuracy and SL F1. Second, we pick a small number of our best-performing models to evaluate on ATIS and Snips test sets, to compare against prior work.

For each experiment below, we train until convergence, where convergence is defined by an early stopping criterion with a patience of 30 epochs and an average of development set IC accuracy and token-level SL F1 used as the performance metric.

6.1 Modeling study experiments

In our first category of experiments, we evaluate variants of each word contextualization paradigm introduced.

We evaluate one feed-forward word contextualization module (labeled as FEED-FORWARD) to provide a baseline performance. As with all subsequent models, we evaluate this word contextualization module with and without our proposed label-recurrent decoder. This baseline should help us determine the extent to which each dataset requires the modeling of context.

We evaluate 3 convolutional word contextualization modules. The first has 1 layer with a kernel size of 5, and is intended to provide intuition as to whether a relatively large local

context can sufficiently model SL behavior. We label this model CNN, 5KERNEL, 1L, and name all other CNN models similarly. The next model has 3 layers with kernel size 5, and is dilated. This model incorporates long-distance context hierarchically, and is shorter and wider-per-layer than the otherwise-similar 3rd CNN model, with 4 layers and kernel size 3.

We evaluate 4 attention-based word contextualization modules. The first is simple, with 1 attention head and 1 layer. Unlike all others we analyze, it does not use relative position embeddings. Thus, this model is word order-invariant except for a simple absolute position embedding. If it improves over FEEDFORWARD, then, it provides strong evidence that semantic information from the context words, irrespective of order, is useful in making tagging decisions. We label this model with the flag NO-POS. To evaluate the utility of relative position embeddings, we also compare a model with 1 head and one layer, labeled ATTN, 1HEAD, 1L. We then test two increasingly complex models, first with 3 layers and 1 head, the second with 3 layers and 2 heads per layer.

We evaluate 2 LSTM-based word contextualization modules; one uses a single LSTM layer, whereas the other stacks a second on top of the

first. As with all other models, we test these two models both with independent slot prediction and label-recurrent slot prediction.

6.2 Comparison to prior work

For our second category of experiments, we take a few high-performing models from our analysis and evaluate them on the Snips and ATIS test sets for comparison to prior work. For these models, we report not only the average IC accuracy and SL F1 across random initializations, but also the standard deviation and best model, as most work has not reported average values. We keep all hyperparameters fixed across all experiments, potentially hindering performance but providing a stronger analysis of robustness.

Note on pre-trained contextual word embedding: Although our framework allows easy integration of contextual pre-trained embeddings like BERT (Devlin et al., 2019) and EMLo (Gardner et al., 2017) by replacing the word contextualization component, we exclude them in our experimentation in order to reduce model obfuscation and have fair comparison against baselines.

7 Results and discussion

In this section, we draw from results reported in Table 1, on the development sets of Snips and ATIS. It is easy to see that very little in the way of modeling is necessary for IC task, so we focus our analysis on SL task. We emphasize that ATIS has shorter spans than Snips, averaging 1.2 and 1.8 tokens, respectively, leading to differing modeling requirements.

7.1 Minimal modeling for SLU

By analyzing three simple models - FEED-FORWARD, ATTN-1HEAD-1L-NO-POS, and CNN-5KERNEL-1L - we conclude that explicitly incorporating local features is a useful inductive bias for high SL accuracy. The purely feed-forward model achieves 53.59 SL F1 on Snips, whereas one layer of convolution improves that number to 85.88. The story is similar for ATIS SL. However, a single layer of attention without position information fails to improve over the feed-forward model whatsoever which we believe is due to the order-invariant nature of self-attention. This also emphasizes the fact that focusing on local context is useful inductive prior for SL task.

For each of these simple models, switching

from independent slot label prediction to label-recurrent prediction provides large gains on both datasets. We find an approximate 1.3ms/utterance slowdown from using label recurrence across all models. Thus, in terms of accuracy-for-speed, very simple models can achieve much of the results of more expensive models as long as they are label-recurrent and incorporate local context.

7.2 High-performing convolutional models

The larger convolutional models provide very high accuracy while maintaining fast inference and training speeds. In particular, our best CNN model, CNN-5KERNEL-3L, achieves 94.22 SL F1 on Snips, compared to the two-layer LSTM with label-recurrence, which achieves 93.88. The model achieves this modest improvement with over 2x the inference speed, training in under 1/2 the time, and demonstrating even stronger results on the test sets, discussed below.

On ATIS, where utterances are longer but slot label spans are shorter, LSTMs outperform CNNs on the development sets.

7.3 Issues with self-attention

Our strongest self-attention model underperforms CNNs and LSTMs on both Snips and ATIS, with a maximum SL of 89.31 and 95.86 on the datasets, respectively. Though self-attention models have seen success in complex tasks with lots of training data, we suggest in this study that they lack the inductive biases to perform well on these small datasets.

Relative position embeddings go a long way in improving self-attention models; adding them to a 1-layer attentional encoder improves ATIS and Snips SL by approximately 24 and 22 points, respectively. We find that adding attention heads does not add considerably to the computational complexity of attention models, while increasing accuracy; thus in a speed-accuracy tradeoff, it is likely better to add heads rather than layers as each layer adds $O(n^2 * d_x)$ additional computations.

7.4 Word and label recurrence in LSTMs

Our LSTM word contextualization modules show that with recurrent word context modeling, label-recurrence is less important. For instance, 2-layer LSTM achieves only .78 increase in SL with label recurrence over independent prediction.

Model	Recurrence	Snips			
		IC Acc		SLR F1	
		Mean	Max	Mean	Max
16 LSTM* (Hakkani-Tür et al., 2016)	full	96.9	-	87.3	-
'16 seq2seq+attn* (Liu and Lane, 2016)	full	96.7	-	87.8	-
LSTM+attn+gates (Goo et al., 2018)	full	97.0	-	88.8	-
OUR CNN, 5KERNEL, 3L	none	97.65±0.28	97.57	89.57±0.54	90.66
OUR CNN, 5KERNEL, 3L	label	97.57±0.41	98.29	92.30±0.40	93.11
OUR LSTM, 2L	word	97.28±0.36	97.57	90.66±0.55	91.53
OUR LSTM, 2L	full (decoupled)	97.22±0.32	97.14	91.53±0.50	92.62

Table 2: Test set results on the Snips dataset. (*) indicates numbers reported by (Goo et al., 2018)

Model	Recurrence	ATIS			
		IC Acc		SLR F1	
		Mean	Max	Mean	Max
LSTM+attn+gates (Goo et al., 2018)	full	94.10	-	95.20	-
'18 Two LSTMs (Wang et al., 2018)	full	-	98.99	-	96.89
'18 self-attn+LSTM (Li et al., 2018)	full	-	98.77	-	96.52
OUR CNN, 5KERNEL, 3L	none	97.04±0.62	97.98	94.84±0.22	94.95
OUR CNN, 5KERNEL, 3L	label	97.37±0.57	98.10	95.27±0.19	95.54
OUR LSTM, 2L	word	96.84±0.49	97.65	95.13±0.29	95.41
OUR LSTM, 2L	full (decoupled)	97.00±0.44	97.98	95.15±0.25	95.21

Table 3: Test set results on the ATIS dataset, compared to recent recurrent models.

7.5 Best models compared to prior work

We report test set results on Snips and ATIS in Tables 2 and 3. Our best models from our validation study, CNN-5KERNEL-3L and LSTM-2L, outperform the state-of-the-art on the Snips dataset, with label-recurrence proving crucial, especially for Snips. In particular, CNN-5KERNEL-3L with label recurrence achieves an average SL F1 of 92.30, improving over the previous state-of-the-art of 88.8, by reducing error rate by 30%, and .57-point improvement on IC.

On ATIS, our label-recurrent models outperform slot-gated LSTM model of Goo et al. (2018) on both IC and SL tasks.³ Wang et al. (2018) attribute their result to using IC and SL-specific LSTMs and use 300-dimensional word embedding and 200-dimensional LSTMs, but with an ATIS vocabulary of 867 words (suggesting a relatively simple sequence space), we are unable to determine the source of the improvement from a modeling standpoint. Similar observation was made for (Li et al., 2018) where 264-dimension embeddings is used.

We hypothesize that our models perform better on Snips because much of Snips slot labeling depends on consistency within long spans,

³We note that, since this work was performed, considerable efforts have been put into the Snips dataset, including the use of ELMo (Siddhant et al., 2019), BERT (Chen et al., 2019b), and capsule networks (Zhang et al., 2019), among other methods (Chen et al., 2019a).




Figure 2: Visualization of the weight given to each token representation by the attention-based pooling for sentence representation. Lighter colors indicate greater attention.

whereas ATIS slot labels have longer-distance dependencies, for example between `to_city` and `from_city` tags.

7.6 Attention Visualization

We note that anecdotally, few words in each utterance are useful in indicating the intent. In the example given in Figure 2, presence of possible departure and arrival cities may be distracting, but the attention mechanism correctly learns to focus on words that indicate `atis.aircraft` intent.

8 Conclusion

We presented a general family of joint IC+SL neural architectures that decomposes the task into modules for analysis. Using this framework, we conducted an extensive study of word contextualization methods (including utility of recurrence in the representation and output space) and determined that label-recurrent models, with non-recurrent word representation and a recurrent model of slot label dependencies, are a good fit for

high performance in both accuracy and speed.

With the results of this study, we proposed a convolution-based joint IC+SL model for SLU that achieves new state-of-the-art results on Snips dataset while maintaining a simple design, shorter training, and faster inference than comparable recurrent methods.

9 Implementation details

All models are implemented in MXNet (Chen et al., 2015). For all models, we randomly initialize word embeddings and use $d_x = 70$. We optimize using Adadelta algorithm (Zeiler, 2012), with initial learning rate, .01. We clip and pad all training and development sentences to length 30, with clipping affecting a small number of utterances. Dropout (Srivastava et al., 2014) probability of .3 is used in all models. We train using a batch size of 128 split across 4 GPUs on a p3.8xlarge EC2 instance, and perform inference using CPUs on same machine.

Acknowledgements

We would like to thank the reviewers for their comments, as well as the Lex team at Amazon AI for helpful conversations and feedback.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *ICLR*.
- Ivan Bilan and Benjamin Roth. 2018. Position-aware self-attention with relative positional encodings for slot filling. *arXiv preprint arXiv:1807.03052*.
- Mengyang Chen, Jin Zeng, and Jie Lou. 2019a. A self-attention joint model for spoken language understanding in situational dialog applications. *arXiv preprint arXiv:1905.11393*.
- Qian Chen, Zhu Zhuo, and Wen Wang. 2019b. BERT for joint intent classification and slot filling. *arXiv preprint arXiv:1902.10909*.
- Tianqi Chen, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, Chiyuan Zhang, and Zheng Zhang. 2015. Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems. In *LearningSys at Neural Information Processing Systems 2015*.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *EMNLP*.
- Alice Coucke, Alaa Saade, Adrien Ball, Théodore Bluche, Alexandre Caulier, David Leroy, Clément Doumouro, Thibault Gisselbrecht, Francesco Calta-girone, Thibaut Lavril, et al. 2018. Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces. *arXiv preprint arXiv:1805.10190*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke S. Zettlemoyer. 2017. *AllenNLP: A deep semantic natural language processing platform*.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. 2017. Convolutional sequence to sequence learning. In *ICML*, pages 1243–1252.
- Chih-Wen Goo, Guang Gao, Yun-Kai Hsu, Chih-Li Huo, Tsung-Chieh Chen, Keng-Wei Hsu, and Yun-Nung Chen. 2018. Slot-gated modeling for joint slot filling and intent prediction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, volume 2, pages 753–757.
- Allen L Gorin, Giuseppe Riccardi, and Jeremy H Wright. 1997. How may i help you? *Speech communication*, 23(1-2):113–127.
- D. Guo, G. Tur, W.T. Yih, and G. Zweig. 2014. Joint semantic utterance classification and slot filling with recursive neural networks. In *Proceedings of Spoken Language Technology Workshop (SLT)*, page 554–559.
- Patrick Haffner, Gokhan Tur, and Jerry H Wright. 2003. Optimizing svms for complex call classification. In *ICASSP*, volume 1, pages I–I. IEEE.
- Dilek Hakkani-Tür, Gökhan Tür, Asli Celikyilmaz, Yun-Nung Chen, Jianfeng Gao, Li Deng, and Ye-Yi Wang. 2016. Multi-domain joint semantic frame parsing using bi-directional RNN-LSTM. In *Inter-speech*, pages 715–719.
- C.T. Hemphill, J.J. Godfrey, and G.R. Doddington. 1990. The ATIS spoken language systems pilot corpus. In *Proceedings of the DARPA Speech and Natural Language Workshop*, page 96–101.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Urvashi Khandelwal, He He, Peng Qi, and Dan Jurafsky. 2018. Sharp Nearby, Fuzzy Far Away: How Neural Language Models Use Context. *Association for Computational Linguistics (ACL)*.

- Nikita Kitaev and Dan Klein. 2018. [Constituency parsing with a self-attentive encoder](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2676–2686. Association for Computational Linguistics.
- Changliang Li, Liang Li, and Ji Qi. 2018. A self-attentive model with gate mechanism for spoken language understanding. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3824–3833.
- B. Liu and I. Lane. 2016. Attention-based recurrent neural network models for joint intent detection and slot filling. In *Proceedings of Interspeech*.
- Robert E Schapire and Yoram Singer. 2000. Boostexter: A boosting-based system for text categorization. *Machine learning*, 39(2-3):135–168.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. [Self-attention with relative position representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 464–468. Association for Computational Linguistics.
- T. Shen, T. Zhou, G. Long, J. Jiang, S. Pan, and C. Zhang. 2017. DiSAN: Directional self-attention network for RNN/CNN-free language understanding.
- Aditya Siddhant, Anuj Goyal, and Angeliki Metallinou. 2019. Unsupervised transfer learning for spoken language understanding in intelligent agents.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958.
- E. Strubell, P. Verga, D. Belanger, and A. McCallum. 2017. Fast and accurate entity recognition with iterated dilated convolutions. In *Proceedings of EMNLP*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Yu Wang, Yilin Shen, and Hongxia Jin. 2018. A bi-model based RNN semantic frame parsing model for intent detection and slot filling. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, volume 2, pages 309–314.
- P. Xu and R. Sarikaya. 2013. Convolutional neural network based triangular CRF for joint intent detection and slot labeling. In *Proceedings of IEEE ASRU Workshop*.
- Kaisheng Yao, Baolin Peng, Yu Zhang, Dong Yu, Geoffrey Zweig, and Yangyang Shi. 2014. Spoken language understanding using long short-term memory neural networks. In *Spoken Language Technology Workshop (SLT), 2014 IEEE*, pages 189–194. IEEE.
- Fisher Yu and Vladlen Koltun. 2015. Multi-scale context aggregation by dilated convolutions. In *ICLR*.
- Matthew D Zeiler. 2012. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.
- Chenwei Zhang, Yaliang Li, Nan Du, Wei Fan, and Philip S Yu. 2019. Joint slot filling and intent detection via capsule neural networks.
- Xiaodong Zhang and Houfeng Wang. 2016. A joint model of intent determination and slot filling for spoken language understanding. In *IJCAI*, pages 2993–2999.

Time Masking: Leveraging Temporal Information in Spoken Dialogue Systems

Rylan Conway and Lambert Mathias

Amazon Alexa AI

{conrylan, mathiasl}@amazon.com

Abstract

In a spoken dialogue system, dialogue state tracker (DST) components track the state of the conversation by updating a distribution of values associated with each of the slots being tracked for the current user turn, using the interactions until then. Much of the previous work has relied on modeling the natural order of the conversation, using distance based offsets as an approximation of time. In this work, we hypothesize that leveraging the wall-clock temporal difference between turns is crucial for finer-grained control of dialogue scenarios. We develop a novel approach that applies a *time mask*, based on the wall-clock time difference, to the associated slot embeddings and empirically demonstrate that our proposed approach outperforms existing approaches that leverage distance offsets, on both an internal benchmark dataset as well as DSTC2.

1 Introduction

Modern spoken dialogue systems – such as Intelligent Personal Digital Assistants (IPDAs) like Google Assistant, Siri, and Alexa – provide users a natural language interface to help complete tasks such as reserving restaurants, checking the weather, playing music etc. Spoken language understanding (SLU) is a central component in such dialogue systems, and is responsible for parsing the natural language text to semantic frames. In task-oriented spoken dialogue systems, a key challenge is tracking entities the user introduced in previous dialogue turns. For example, if a user request for *what’s the weather in arlington* is followed by *how about tomorrow*, the dialogue system has to keep track of the entity *arlington* being referenced. Typically, this is formulated as a dialogue state tracking (DST) task (Henderson et al., 2014b; Mrkšić et al., 2016).

Previous approaches to dialogue state tracking have mostly focused on dialogue representa-

tions (Mrkšić et al., 2016), dealing with noisy input (Henderson et al., 2012; Mesnil et al., 2015), or tracking slots from multiple domains (Henderson et al., 2014b; Rastogi et al., 2017; Naik et al., 2018). In this paper, we focus on temporal information associated with each dialogue turn. Although the dialogue representations – typically encoded using LSTMs – are able to implicitly capture the temporal order in the sequence of dialogue turns, we hypothesize that explicitly and accurately encoding temporal information is essential for resolving ambiguity in dialogue state tracking. Recently, (Naik et al., 2018) presented work that models the slot distance offset from the current turn using a one-hot representation input to the DST module. Alternatively, (Su et al., 2018) leverage the distance offset in an attention mechanism. We posit that the notion of time based on distance offset relative to the current turn is too coarse-grained and often insufficient for resolving ambiguities associated with more complex multi-turn dialogues. For example, in a dialogue “*how far is issaquah?*” followed by “*what is the weather like?*” we could have two possible interpretations – a follow-up utterance issued within 10 seconds would indicate that the user is referring to the city slot of “Issaquah” from the previous turn, whereas, if the follow-up utterance is more than 30 seconds apart there is a good chance that the user was just inquiring about the weather in their current location. In this case, a dialogue system that only encodes the distance offset will be unable to correctly disambiguate the aforementioned situation. Based on this intuition, we develop a novel approach for incorporating temporal information in dialogue state tracking by using a time mask over the slots.

To summarize, we introduce the notion of a *time mask* to incorporate temporal information into the embedding for slots. In contrast to previous ap-

proaches using distance offsets, we propose leveraging the wall-clock time difference between the current turn and the previous turns in the dialogue to explicitly model temporal information. Furthermore, we demonstrate how domain and intent information can be mixed in with the temporal information in this framework to improve DST accuracy. We demonstrate empirically that our proposed approach improves over the baseline that only encodes distance offsets as temporal information.

2 Approach

2.1 Slot Carryover Task Description

In this paper, we build on the approach in (Naik et al., 2018). For completeness, we define the carryover task formulation here, but refer readers to the original work for architecture details. A dialogue turn at time t is defined as the tuple $\{a_t, \mathcal{S}_t, \mathbf{w}_t\}$, where $\mathbf{w}_t \in \mathcal{W}$ is a sequence of words $\{w_{it}\}_{i=1}^{N_t}$; $a_t \in \mathcal{A}$ is the dialogue act; and \mathcal{S}_t is a set of slots, where each slot s is a key-value pair $s = \{k, v\}$, with $k \in \mathcal{K}$ being the slot name (or slot key), and $v \in \mathcal{V}$ being the slot value. A user turn is represented by $\mathbf{u}_t = \{a_t^u, \mathcal{S}_t^u, \mathbf{w}_t^u\}$ and a system turn is represented by $\mathbf{v}_t = \{a_t^v, \mathcal{S}_t^v, \mathbf{w}_t^v\}$. Given a sequence of D user turns $\{\mathbf{u}_{t-D}, \dots, \mathbf{u}_{t-2}, \mathbf{u}_{t-1}\}$; and their associated system turns $\{\mathbf{v}_{t-D}, \dots, \mathbf{v}_{t-2}, \mathbf{v}_{t-1}\}$ ¹; and the current user turn \mathbf{u}_t , we construct a candidate set of slots from the context as

$$C(\mathcal{S}) = \bigcup_{\substack{j=t-D \\ i \in u, v}}^t \mathcal{S}_j^i. \quad (1)$$

For a candidate slot $s \in C(\mathcal{S})$, for the dialogue turn at time t , the probability to carryover the slot is defined as

$$P(+|s, d(s), \mathbf{u}_t, \mathbf{u}_{t-D}^{t-1}, \mathbf{v}_{t-D}^{t-1}), \quad (2)$$

where $d(s) \in [0, D]$ is an integer value describing the offset of the candidate slot from the current turn \mathbf{u}_t . The final carryover decision is determined by comparing the carryover probability to a tunable decision threshold

An encoder-decoder model is used to evaluate each slot candidate, as shown in Figure 1. The current turn, past user turns, and past system turns are all encoded using an LSTM layer with atten-

¹For simplicity we assume a turn taking model - a user turn and system turn alternate.

tion. Each slot (key and value) and intent are also encoded by averaging the word embeddings contained in each. Finally, the *slot distance* is encoded by counting the number of turns back that the slot appeared in (this would equal zero for slots from the current turn) and one-hot encoding that value. This is shown by the "Recenty One-Hot" input in the diagram. The final encoded slot candidate is passed to the decoder which produces a final carryover probability that determines whether or not the slot should be carried over to the current turn².

2.2 Simple Time Mask (STM)

Inspired by (Li et al., 2018), we introduce the concept of *masked* embeddings so that irrelevant dimensions are suppressed in the embedding of the slots. We start by constructing a *time embedding* based on the *temporal distance*, $d_{\Delta t}$, of each candidate slot³. This is shown in Fig. 1 as the bottom input, in the red box. The time embedding is given by

$$\mathbf{d}_t = \phi(W_t d_{\Delta t} + \mathbf{b}_t), \quad (3)$$

where \mathbf{d}_t is a nonlinear transformation implemented as a single layer feedforward neural network with weight matrix $W_t \in \mathbb{R}^{N_t \times 1}$ and N_t is dimensionality of the time embedding vector. The time mask, \mathbf{m}_t , is computed by passing the time embedding, \mathbf{d}_t , through another feedforward neural network

$$\mathbf{m}_t = \sigma(W_{d_t} \mathbf{d}_t + \mathbf{b}_{d_t}), \quad (4)$$

where $W_{d_t} \in \mathbb{R}^{N_s \times N_t}$ and N_s is the dimensionality of the candidate slot embedding, \mathbf{h}_s .

Finally, we apply the time mask to the encoded slot embedding:

$$\mathbf{h}'_s = \mathbf{h}_s \odot \mathbf{m}_{d_t}, \quad (5)$$

The updated candidate slot embedding, \mathbf{h}'_s , is now passed to the decoder in the exact same way as in the baseline slot carryover model as described in (Naik et al., 2018).

Temporal dialogue behavior can vary by domain. Figure 2 shows how much the distribution of $d_{\Delta t}$ can differ between three different domains in an internal IPDA dataset (described in more detail in 3.1). Therefore, we consider two exten-

²The inputs shown in red are not part of the original formulation in (Naik et al., 2018).

³Defined as number of seconds in the past that the turn which contains the slot occurred relative to the current utterance.

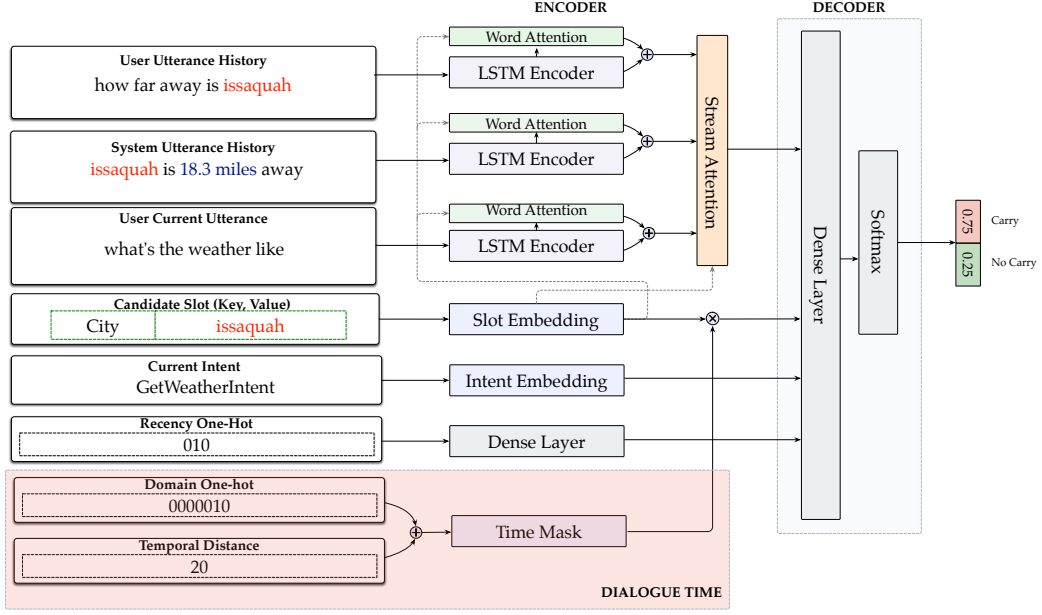


Figure 1: Slot carryover architecture from (Naik et al., 2018) augmented with a temporal component using domain-specific time masking as described in Section 2.2.2.

sions of the time masking approach that take into account the multi-domain nature of IPDAs.

2.2.1 Intent Specific Time Mask (ITM)

We leverage the dialogue act or intent associated with the current turn in the time mask model. In this formulation the time embedding is now given by

$$\mathbf{d}_t = \phi(W_t \mathbf{d}_{\Delta t, a} + \mathbf{b}_t), \quad (6)$$

where $\mathbf{d}_{\Delta t, a} = d_{\Delta t} \oplus \mathbf{h}_a$ is just the temporal distance concatenated with the existing intent embedding \mathbf{h}_a and now $W_t \in \mathbb{R}^{N_t \times (N_a + 1)}$, where N_a is the number of dimensions used in the intent embedding.

2.2.2 Domain Specific Time Mask (DTM)

We also try more coarse-grained, domain-level information as input to the time embedding. Here we use a one-hot encoded representation of the domains, which gives us:

$$\mathbf{d}_t = \phi(W_t \mathbf{d}_{\Delta t, D} + \mathbf{b}_t), \quad (7)$$

where $\mathbf{d}_{\Delta t, D} = d_{\Delta t} \oplus \mathbf{1}_D$ is the concatenation of the temporal distance with the one-hot-encoded domain, $\mathbf{1}_D$, and $W_t \in \mathbb{R}^{N_t \times (N_D + 1)}$, where N_D is the number of dimensions used in the one-hot-encoded Domain representation. This architecture is shown in in Figure 1.




Figure 2: Distribution of temporal distance, $d_{\Delta t}$, for all carryover slots from three of the largest domains in the IPDA dataset.

2.2.3 Time-Decay Attention (TDA)

For comparison, we re-implemented the time-decay attention (TDA) model introduced in (Su et al., 2018). However, the original work does not actually use time as a feature input but rather the *ordinal distance* (equivalent to slot distance in our formulation) of each dialogue turn from the current utterance. To compare with our methods we use the actual temporal difference between dialogue turns in our implementation of the TDA

model. The parameters are learned in the end-to-end training process.

3 Experiments

3.1 Dataset

We present the results on 2 datasets. The **IPDA dataset**, in Table 1, is an internal benchmark dataset collected from an IPDA for the en-US locale based on real usage. It consists of interactions over 7 domains - Music, Weather, LocalSearch, SmartHome, Video, MovieShowTimes, and Question Answering. The data is transformed into individual candidate slots that are presented to the model, which determines whether or not they are relevant for the given turn. For benchmarking against a public corpora, we also measure performance on DSTC2 (Henderson et al., 2014a) dataset. We post-process the dataset similar to the internal dataset and only consider the top ASR and SLU hypothesis in addition to the system turn, dialogue acts and the associated slots.

	Train	Dev	Test
Total	264148	32437	33747
Positive Carryover	92084	11389	11769
Avg. $d_{\Delta t}$	15.33s	15.58s	15.31s

Table 1: IPDA dataset statistics. Here ‘positive carryover’ slots is the number of candidate slots that are relevant for the current turn.

Figure 3 shows the distribution of time between turns for both datasets. If a slot candidate came from a context turn that was spoken 20 seconds before the current turn then $d_{\Delta t} = 20$. Based on human judged ground-truth, the slots that should be carried over to the current turn are shown in orange and the slots that should not are shown in blue. One clear difference between the two distributions in the IPDA dataset is the long tail of the non-carryover distribution, indicating carryovers are more likely from a recent turn. The domain specific distributions further indicate that leveraging dialogue time could be useful.

3.2 Results on IPDA Dataset

From Table 2, we can see that the TDA models offer a slight improvement over the baseline model. Both models incorporate slot distance offset but the attention mechanism provides an additional boost. The time mask models show additional gains demonstrating that leveraging dialogue time from each turn is important. Moreover, the time

information provides complementary information over the distance offset based measure, as shown by the improvements of the time masking models over the baseline model. The DTM model performs the best overall in terms of F1, which suggests that adding domain information into the time mask provides additional disambiguation power. Interestingly, we see that the ITM model does not improve much over the STM model, possibly because the intent embeddings do not necessarily distinguish between temporal behavior, and are already being leveraged by the slot carryover model.

3.2.1 Investigating longer temporal distance

Here, we investigate the ability of the models to maintain higher accuracy over longer time windows in the dialogue context. The overall F1 scores for each model are binned by $d_{\Delta t}$ for each candidate slot. The results are shown in Table 2. The domain specific time mask model performs the best in each $d_{\Delta t}$ bin. The effect of adding dialogue time information significantly improves performance in the largest $d_{\Delta t}$ range. This is likely due to the model learning that older slots are less relevant to the current turn, which is impossible for the baseline model to do. Additionally, we can see that the TDA model performs comparably to the STM model in the range $0 < d_{\Delta t} \leq 30$ but in the highest bin ($30 < d_{\Delta t} \leq 60$) we see that it falls well short of all of the time-masked models.

3.3 Results on DSTC2 Dataset

Since there is only one domain in DSTC2, we chose to only implement the STM model. From the last column in Table 2, we can see that the STM model produces the best result. The TDA model, contrary to previously reported results on DSTC4, does not perform as well. Our hypothesis is that the temporal distribution across turns is not monotonically decaying, which is an assumption made in their approach.

4 Related Work

Previous work on leveraging temporal information for dialogue state tracking has focused mostly on using distance offsets. (Chen et al., 2017) presented a time-aware attention network to leverage both contextual and ordinal distance information (i.e. the number of turns back from the current turn) and saw significant improvement. Subsequently, (Su et al., 2018) improved upon this by

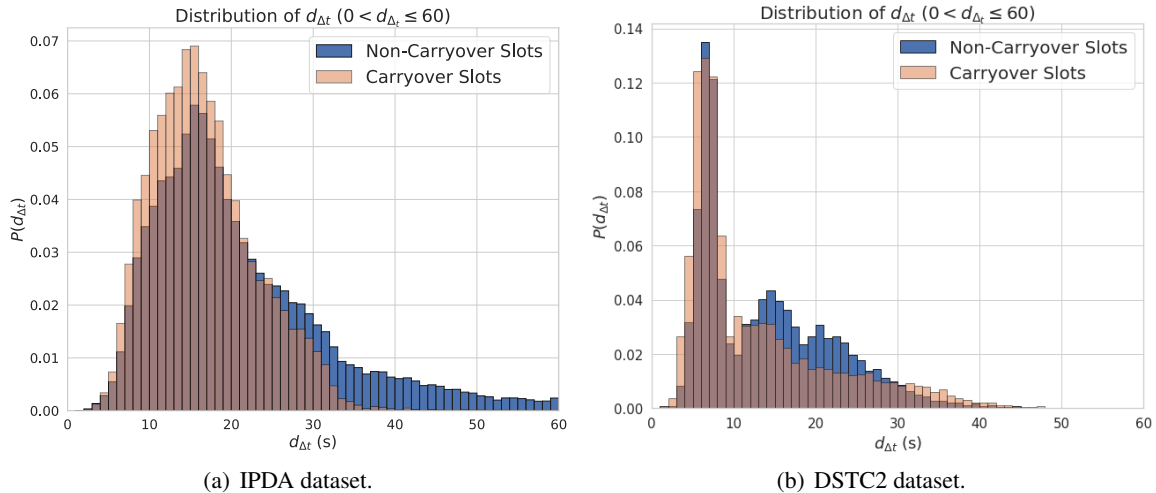


Figure 3: Distribution of temporal distance, $d_{\Delta t}$, for all candidate slots. The “Carryover Slots”, shown in orange, represent slot candidates found in the context that *should* be carried over to the current turn.

Model	Overall	$d_{\Delta t} = (0, 15]$	$d_{\Delta t} = (15, 30]$	$d_{\Delta t} = (30, 60]$	Overall DSTC2 F1
Baseline (Naik et al., 2018)	87.8	89.3	86.8	74.3	95.0
STM	88.4	89.7	87.5	77.5	96.1
ITM	88.6	89.8	87.8	76.9	-
DTM	89.2	90.5	88.3	80.0	-
TDA (Su et al., 2018)	88.4	90.0	87.5	72.8	94.6

Table 2: Overall F1 scores on the IPDA and DSTC2 dataset as well as F1 scores binned by $d_{\Delta t}$ for the IPDA dataset, which is measured in seconds. Note: the DSTC2 dataset only contains a single domain

designing a more flexible data-driven time attention mechanism that applied continuously decaying weights to past utterances before being fed into a contextual encoder. The attention weight was determined based on the distance offset relative to the current turn. However, distance offset is unable to capture complex dialogue scenarios, and our work improves upon this by modeling the actual wall-clock time difference between the current turn and the contextual turns. This is particularly important in a multi-domain setting where a few second pause between consecutive user turns can be interpreted very differently depending on the dialogue scenario, and our experiments support our hypothesis.

Embedding masks have been explored in machine translation. (Choi et al., 2016) showed that contextualized word embeddings could be constructed from static word embeddings by applying a learned context mask. This context mask allows the word to have different representations depending on the source sentence context around the word that is being translated, and the authors demonstrated improvements in machine translation tasks with this approach. The approach of masking word representations was also explored

in (Ruseti et al., 2016) for categorizing words into their wordnet classes. We extend this masking concept to dialogue state tracking task, where we encode the temporal information in the dialogue as the masking operation over slots.

5 Conclusion

In this work we presented a novel approach for incorporating dialogue time information in multi-domain large-scale SLU systems. We showed that our proposed time masking strategy provided gains over baseline systems that simply encode dialogue distance. We presented several methods for incorporating additional information such as domain and intents into the time mask, and showed that this approach improved over competing approaches that indirectly incorporate time, particularly for multi-domain dialogues. In the future, we want to investigate more contextualized representations of the domain and intent in order to capture more subtle variations in the dialogue for multi-domain scenarios.

References

- Po-Chun Chen, Ta-Chung Chi, Shang-Yu Su, and Yun-Nung Chen. 2017. Dynamic time-aware attention to speaker roles and contexts for spoken language understanding. In *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 554–560. IEEE.
- Heeyoul Choi, Kyunghyun Cho, and Yoshua Bengio. 2016. [Context-dependent word representation for neural machine translation](#). *arXiv e-prints*, abs/1607.00578.
- Matthew Henderson, Milica Gašić, Blaise Thomson, Pirros Tsiakoulis, Kai Yu, and Steve Young. 2012. Discriminative spoken language understanding using word confusion networks. In *2012 IEEE Spoken Language Technology Workshop (SLT)*, pages 176–181. IEEE.
- Matthew Henderson, Blaise Thomson, and Jason D Williams. 2014a. The second dialog state tracking challenge. In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 263–272.
- Matthew Henderson, Blaise Thomson, and Steve Young. 2014b. Word-based dialog state tracking with recurrent neural networks. In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 292–299.
- Yang Li, Nan Du, and Samy Bengio. 2018. [Time-dependent representation for neural event sequence prediction](#). *Proceedings of the Sixth International Conference on Learning Representations*. To appear.
- Grégoire Mesnil, Yann Dauphin, Kaisheng Yao, Yoshua Bengio, Li Deng, Dilek Hakkani-Tur, Xiaodong He, Larry Heck, Gokhan Tur, Dong Yu, et al. 2015. Using recurrent neural networks for slot filling in spoken language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(3):530–539.
- Nikola Mrkšić, Diarmuid O Séaghdha, Tsung-Hsien Wen, Blaise Thomson, and Steve Young. 2016. Neural belief tracker: Data-driven dialogue state tracking. *arXiv preprint arXiv:1606.03777*.
- Chetan Naik, Arpit Gupta, Hancheng Ge, Lambert Mathias, and Ruhi Sarikaya. 2018. Contextual slot carryover for disparate schemas. In *Interspeech*.
- Abhinav Rastogi, Dilek Hakkani-Tür, and Larry Heck. 2017. Scalable multi-domain dialogue state tracking. In *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 561–568. IEEE.
- Stefan Ruseti, Traian Rebedea, and Stefan Trausan-Matu. 2016. Using embedding masks for word categorization. In *Proceedings of the 1st Workshop on Representation Learning for NLP*, pages 201–205.
- Shang-Yu Su, Pei-Chieh Yuan, and Yun-Nung Chen. 2018. How time matters: Learning time-decay attention for contextual spoken language understanding in dialogues. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Long Papers)*, volume 1, pages 2133–2142.

To Combine or Not To Combine? A *Rainbow* Deep Reinforcement Learning Agent for Dialog Policy

Dirk Văth Ngoc Thang Vu

University of Stuttgart, Institute for Natural Language Processing (IMS)

Pfaffenwaldring 5B, 70569 Stuttgart, Germany

{dirk.vaeth, thang.vu}@ims.uni-stuttgart.de

Abstract

We explore state-of-the-art deep reinforcement learning methods such as prioritized experience replay, double deep Q-Networks, dueling network architectures, distributional learning methods for dialog policy. Our main findings show that each individual method improves the rewards and the task success rate but combining these methods in a *Rainbow* agent, which performs best across tasks and environments, is a non-trivial task. We, therefore, provide insights about the influence of each method on the combination and how to combine them to form the *Rainbow* agent.

1 Introduction

Dialog system can be designed for generic purposes, e.g. smalltalk (Weizenbaum, 1966) or a specific task such as finding restaurants or booking flights (Bobrow et al., 1977; Wen et al., 2017). This paper focuses on task-oriented dialog systems, which interact with a user to aid achieving their goals. The systems have several modules which solve different subtasks (Williams et al., 2016) starting with natural language understanding (NLU) module (De Mori et al., 2008). Its output is then passed to a belief tracking module (Mrkšić et al., 2017) that holds the state of the dialog, i.e. all relevant information provided by the user. This belief state is then passed to the dialog policy module (Williams and Young, 2007) which has to decide how the system should reply. Depending on the ontology of the task, e.g. the restaurant search, the size of the input space for the policy can quickly become very large. Furthermore, the belief state might be wrong due to noisy inputs, e.g. the user could be misunderstood because of NLU errors or in general, language ambiguity. Therefore, building such policies by hand is rather time consuming. Reinforcement learning (RL) can alleviate this task by allowing

to learn such policies automatically (Williams and Young, 2007) with a user simulator such as proposed in Schatzmann et al. (2007) within a task (Dhingra et al., 2017; Peng et al., 2018), between task and non-task (Yu et al., 2017) and also in multimodal dialog systems (Manuvinakurike et al., 2017; Zhang et al., 2018).

Deep RL has been proven to be successful with Deep Q-Learning (DQN) (Mnih et al., 2013) introducing the idea of using neural networks as a Q-function approximator. It has been widely used in the context of dialog policy learning (Fatemi et al., 2016; Dhingra et al., 2017; Casanueva et al., 2017). However according to a recent comparison (Casanueva et al., 2017) in the context of dialog policy learning, it performed worse than other RL methods such as Gaussian Process in many testing conditions. Recently, several advances in deep RL such as distributional RL (Bellemare et al., 2017), dueling network architectures (Wang et al., 2016) and their combination (Hessel et al., 2018) - a *Rainbow* agent - have been shown to be promising for further improvements of deep RL agents in benchmark environments, e.g. Atari 2600. However, it is still unclear whether these methods could advance dialog policies.

This paper attempts to provide insights motivated from dialog policy modeling perspectives how to use state-of-the-art deep RL methods such as prioritized experience replay (Schaul et al., 2015), double DQN (Van Hasselt et al., 2016), dueling network architecture, distributional learning method and how to combine them to train the *Rainbow* agent for dialog policy learning¹. Moreover, we explore the influence of each method w.r.t the resulting rewards and the number of successful dialogs, highlighting methods with the biggest and the smallest impact.

¹Agent code: <https://github.com/DigitalPhonetics/adviser>

Task	Env. 1			Env. 2			Env. 3			Env. 4			Env. 5			Env. 6		
	T1.1	T1.2	T1.3	T2.1	T2.2	T2.3	T3.1	T3.2	T3.3	T4.1	T4.2	T4.3	T5.1	T5.2	T5.3	T6.1	T6.2	T6.3
Domain	CR	SFR	LAP	CR	SFR	LAP	CR	SFR	LAP	CR	SFR	LAP	CR	SFR	LAP	CR	SFR	LAP
SER		0%			0%			15%			15%			15%				30%
Masks		On			Off			On			Off			On				On
User		Standard			Standard			Standard			Standard			Unfriendly				Standard

Table 1: Benchmarking environments with several domains, semantic error rates (SERs), action masks and user models (Casanueva et al., 2017).

2 Proposed Method

For value-based reinforcement learning methods like Q-learning, potentially large state spaces as in the dialog setting require the use of function approximators. The DQN-Algorithm (Mnih et al., 2013) is an example of such a method where the action-value function is approximated by a neural network which takes a state vector as input and outputs a value for each possible action. Loss is calculated with the squared temporal difference (TD) error. Efficient off-policy batch-training is enabled by a replay buffer which records the agent’s turn-level experiences and allows the drawing of uncorrelated training samples.

Prioritized experience replay Drawing samples from this buffer uniformly is straightforward but problematic: important state transitions might never be drawn from the buffer or at least too few times to have an impact on the network weights. Motivated by the insight that a high absolute TD-error of an experience means that the current action-value is not an accurate estimate yet, prioritized experience replay (Schaul et al., 2015) samples experiences having higher TD-errors with greater probability than those with lower TD-error. This method is relevant because it is expected to increase learning efficiency. In the context of dialog policy, there are some system actions which are crucial to the outcome of the dialog and should have a higher probability for being used as training data if they are not well approximated. For example, if systems end the dialog before the user’s goal is completed by telling the user *goodbye*, this will immediately terminate a dialog with a negative reward and without any chance of recovery.

Double DQN Another improvement mitigates the overestimation bias inherent to Q-learning by introducing a second action-value network (Van Hasselt et al., 2016) which copies the parameters from the online action-value network periodically and is held fix otherwise. This additional network is then used to evaluate the action-value of the action selected greedily w.r.t. the online

Q-function, thereby decoupling action choice and evaluation which could increase stability of the learning process.

Dueling network architecture In comparison to the action-value function, the state-value function is a simpler estimate - it is the expectation over a state’s action-values and therefore only a single value. But in states where the action choice does not matter, or to avoid visiting states with a low state value in general, an estimate of the value function should be sufficient. Dueling network architecture (Wang et al., 2016) therefore splits the calculation of the action-value function into separate layers of a neural network, one group computing the value function and another an advantage function chosen so that their combination results in the action-value function again. This approach also has the benefit that the state value estimation is updated every time when a state is observed by the network, regardless of the chosen action. As a result, it should encourage generalization across actions. In dialog settings, there are many states where generalization across actions could prove beneficial, e.g. exact action choice is not important, just the choice between action classes. For example at the beginning of a dialog, when users greet the system without providing any information, the only appropriate action for the system is to ask for more information. The exact type of information should not matter and all other actions except for the dialog ending action should be about equally unsuitable.

Distributional learning method One of the latest additions to reinforcement learning is the quantile regression distributional reinforcement learning algorithm (Dabney et al., 2018). Instead of learning only the expected value for each state-action pair, as in regular Q-learning, the distribution of rewards is approximated instead, thereby modeling the randomness of the reward over multiple turns induced by action selection and random state transitions. A noisy environment like dialog could benefit from better knowledge about the distribution of rewards.

The Rainbow agent Following the methodology from (Hessel et al., 2018), we extend the DQN algorithm (Mnih et al., 2013) with prioritized experience replay, double DQN, and dueling network architecture. Furthermore in contrast to (Hessel et al., 2018), we apply the following changes to successfully train the *Rainbow* agent: 1) we drop the multi-step method (Sutton, 1988) because it seems to diminish the obtained rewards. As the step size gets larger, the rewards are decreased more. A possible explanation could be that the noise generated by the user simulator leads to accumulation of noise in rewards over multiple steps, which could lead to higher variance in value estimates. 2) we discard the noisy linear layers (Fortunato et al., 2018), relying on ϵ -greedy exploration instead. The first reason could be the additional parameters, which usually would require more training samples. Since the agent was already required to learn environmental noises from the user simulator, a complementary explanation could be that the inclusion of a second noise distribution might have been too difficult to learn, especially when considering the relatively small amount of training episodes. 3) we swap the categorical DQN approach (Bellemare et al., 2017) with the quantile regression Q-learning algorithm (Dabney et al., 2018), now consistent with the theoretic results from (Bellemare et al., 2017), no longer restricting the values of the value distribution to a uniform resolution and also no longer requiring knowledge about their bounds.

3 Resources

We used PyDial toolkit (Ultes et al., 2017) as a test-bed for experiments and evaluation. It includes a configurable user simulator and provides multiple dialog ontologies like Cambridge Restaurants (CR), Laptops (LAP) and San Francisco Restaurants (SFR). The ontologies used for the benchmarks in this paper together with their properties are listed in table 2.

Domain	#slots	#requests	#values
CR	3	9	268
SFR	6	11	636
LAP	11	21	257

Table 2: Benchmark domains with #slots the user can provide or #request from the system as well as #values of each requestable slot (Casanueva et al., 2017).

Casanueva et al. (2017) propose six different environmental models, varying in user friendliness, simulated input channel noise and the presence or absence of action masks, which, when enabled, simplify learning by masking some of the possible actions. An overview of all these environmental configurations and their assignment to tasks is given in Table 1. Evaluation results in Casanueva et al. (2017) with several dialog policy types, e.g. a handcrafted policy and the best reported policies serve as baselines in our experiments.

4 Experimental Results

Training and evaluation with the PyDial user simulator follows the PyDial benchmarking tasks (Casanueva et al., 2017), where each task (see Table 1) is trained on 10000 dialogs split into ten training iterations of 1000 dialogs each. We evaluate policies after each training iteration on 1000 test dialogs. All of the following results were obtained by averaging over the outcome of ten different random seeds using the parameters described in appendix A.

4.1 The Rainbow Agent

The first row of Table 3 and 4 show the results of the highest scoring policy from the PyDial benchmark (Casanueva et al., 2017) to serve as baselines. Evaluations of the handcrafted policies follow in the last line. The results show that *Rainbow* agent outperforms reward of the best PyDial agents in all 18 conditions and success rate in 16 out of 18 setting. Compared to the basic DQN agent, *Rainbow* agent is better in all 18 conditions w.r.t both reward and success rate. When averaged across all 18 tasks, *Rainbow* agent (mean reward 10.1) scores more than 29% higher rewards compared to the best PyDial agent (DQN, mean reward 7.8) and more than 9.7% compared to our DQN agent. An average success rate of 90.4% is superior to the best PyDial agent (GP-Sarsa, 80.2%). Mean deviation across all tasks and random seeds is 0.4 in reward and 1.6% in successful dialogs.

4.2 Model Ablation Analysis

Figure 1 shows the averaged success rates for each of our *Rainbow* agents leaving out one particular method after training with 10000 dialogs. Each of the plotted values has been evaluated on 1000 dialogs per random seed and averaged over all tasks. Regarding learning speed w.r.t. success rate, the

Agent	Task																	
	T1.1	T1.2	T1.3	T2.1	T2.2	T2.3	T3.1	T3.2	T3.3	T4.1	T4.2	T4.3	T5.1	T5.2	T5.3	T6.1	T6.2	T6.3
best PyDial	13.5 ¹	12.3 ²	11.0 ²	12.7 ³	10.1 ¹	9.1 ³	12.2 ³	8.6 ²	6.5 ²	11.1 ³	8.2 ³	5.8 ¹	10.5 ³	6.5 ²	3.8 ²	9.9 ³	3.6 ³	3.2 ²
DQN	13.0	10.8	9.5	13.1	11.0	9.5	12.7	9.7	7.5	11.9	7.9	5.1	11.4	7.3	4.3	10.7	5.7	4.7
<i>Rainbow</i>	14.0	12.4	11.2	13.6	11.8	10.1	12.8	9.8	8.1	12.2	10.0	8.9	11.8	7.8	4.9	10.9	6.5	4.8
handcrafted	14.0	12.4	11.7	14.0	12.4	11.7	11.0	9.0	8.7	11.0	9.0	8.7	9.7	6.4	5.5	9.3	6.0	5.3

Table 3: Rewards per task and agent (¹GP-SARSA, ²eNAC, ³DQN).

Agent	Task																	
	T1.1	T1.2	T1.3	T2.1	T2.2	T2.3	T3.1	T3.2	T3.3	T4.1	T4.2	T4.3	T5.1	T5.2	T5.3	T6.1	T6.2	T6.3
best PyDial	99.4 ¹	97.3¹	92.1 ²	97.9 ¹	95.4¹	87.5 ¹	95.8 ¹	84.1 ²	73.3 ²	92.6 ¹	81.1 ³	74.0 ¹	92.6 ³	82.3 ²	72.8 ²	89.6 ¹	64.8 ³	61.2 ²
DQN	95.1	89.4	83.7	96.9	91.4	87.6	97.1	89.6	79.6	94.3	79.8	68.0	95.6	84.9	74.4	91.7	75.8	71.1
<i>Rainbow</i>	99.7	97.3	93.4	98.8	94.4	90.3	97.2	90.5	83.6	96.5	88.8	87.3	97.0	87.9	78.0	92.4	80.4	73.0
handcrafted	100.0	98.2	97.0	100.0	98.2	97.0	96.7	90.9	89.6	96.7	90.9	89.6	95.9	87.7	85.1	89.6	79.0	76.1

Table 4: Success rates per task and agent (¹GP-SARSA, ²eNAC, ³DQN).

results show that *Rainbow* agent without distributional learning method learns the fastest, surpassing the final success rate of the DQN and non-dueling agents after only 2000 dialogs. The reward plot displays similar characteristics.




Figure 1: Avg. success rate for policies after training with 10000 dialogs (1000 dialogs per iteration).

Results in Table 5 show that there is almost no difference between the distributional and the non-distributional approach. Their final rewards are the same and their success rates differ by 0.1% when averaged across all tasks. A possible reason could be that the diversity of the training dialogs was too little and rewards too sparse to show a benefit by using the distributional reinforcement learning method. This coincides with the findings in Hessel et al. (2018) which found their combined agent without distributional learning performing similar to the combined agent with distributional learning for the first 40 million frames on the Atari benchmark.

The strongest benefits to final performance come with the dueling architecture. For some scenarios like the previously described dialog start without any user-provided information, we examined the action-state values by clustering them and

observed fewer clusters and smaller within-cluster variance for the dueling agents, indicating better generalization and simpler action-value functions. Prioritized experience replay helped with learning efficiency but had no significant effect on final performance, as expected. Only a small improvement can be attributed to double DQN, but overall performance seems to be slightly more stable.

Overall, Table 5 shows that the final best *Rainbow* agent performs considerably better than the best reported PyDial agent and the DQN agent across all the tasks and testing environments and is on par with handcrafted policy performance.

Agent	CR		SFR		LAP	
	Suc.	Rew.	Suc.	Rew.	Suc.	Rew.
best PyDial	94.5% ¹	10.7 ¹	79.7% ¹	6.8 ²	66.8% ²	5.0 ²
DQN	95.1%	12.2	85.1%	8.7	77.4%	6.8
<i>Rainbow</i>	96.9%	12.6	89.9%	9.7	84.3%	8.0
- distributional	96.6%	12.4	89.6%	9.6	85.4%	8.2
- double	96.0%	12.3	88.5%	9.4	85.2%	8.3
- dueling	95.5%	12.1	84.7%	8.4	77.0%	6.4
- prioritization	97.4%	12.7	89.0%	9.7	82.1%	8.0
handcrafted	90.8%	9.2	90.8%	9.2	89.1%	8.6

Table 5: Success rates and rewards per domain (¹GP-SARSA, ²DQN).

5 Conclusions

We explored state-of-the-art deep RL methods for dialog policy on different domains with various noise levels and user behaviours. Our findings are that not all extensions to DQN prove beneficial in dialog policy settings, especially when learning speed is concerned: distributional reinforcement learning method requires more training time to reach the success rates and final rewards of the non-distributional agent. The *Rainbow* agent that makes use of prioritized experience replay, double DQN and dueling network architecture is stable across domains and evaluation settings and learns fastest (when excluding distributional method).

References

- Stefan Ultes et al. 2017. PyDial: A Multi-domain Statistical Dialogue System Toolkit. In *ACL, System Demonstrations*.
- Marc G Bellemare, Will Dabney, and Rémi Munos. 2017. A distributional perspective on reinforcement learning. In *International Conference on Machine Learning*.
- Daniel G. Bobrow, Ronald M. Kaplan, Martin Kay, Donald A. Norman, Henry Thompson, and Terry Winograd. 1977. GUS: a Frame-Driven Dialog System. *Artificial Intelligence*, 8.
- Iñigo Casanueva, Paweł Budzianowski, Pei-Hao Su, Nikola Mrkšić, Tsung-Hsien Wen, Stefan Ultes, Lina Rojas-Barahona, Steve Young, and Milica Gašić. 2017. A benchmarking environment for reinforcement learning based task oriented dialogue management. In *Deep Reinforcement Learning Symposium*.
- Will Dabney, Mark Rowland, Marc G Bellemare, and Rémi Munos. 2018. Distributional reinforcement learning with quantile regression. In *AAAI*.
- Renato De Mori, Frédéric Bechet, Dilek Hakkani-Tur, Michael McTear, Giuseppe Riccardi, and Gokhan Tur. 2008. Spoken language understanding. *IEEE Signal Processing Magazine*, 25(3).
- Bhuvan Dhingra, Lihong Li, Xiujun Li, Jianfeng Gao, Yun-Nung Chen, Faisal Ahmed, and Li Deng. 2017. Towards end-to-end reinforcement learning of dialogue agents for information access. In *ACL (Vol. 1: Long Papers)*.
- Mehdi Fatemi, Layla El Asri, Hannes Schulz, Jing He, and Kaheer Suleman. 2016. Policy networks with two-stage training for dialogue systems. In *SigDial*.
- Meire Fortunato, Mohammad Gheshlaghi Azar, Bilal Piot, Jacob Menick, Ian Osband, Alex Graves, Vlad Mnih, Remi Munos, Demis Hassabis, Olivier Pietquin, et al. 2018. Noisy networks for exploration. In *International Conference on Learning Representations*.
- Matteo Hessel, Joseph Modayil, Hado Van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Azar, and David Silver. 2018. Rainbow: Combining improvements in deep reinforcement learning. In *AAAI*.
- Ramesh Manuvinakurike, David DeVault, and Kalliroi Georgila. 2017. Using reinforcement learning to model incrementality in a fast-paced dialogue game. In *SigDial*.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. In *NIPS Deep Learning Workshop*.
- Nikola Mrkšić, Diarmuid O Séaghdha, Tsung-Hsien Wen, Blaise Thomson, and Steve Young. 2017. Neural belief tracker: Data-driven dialogue state tracking. In *ACL (Vol. 1: Long Papers)*.
- Baolin Peng, Xiujun Li, Jianfeng Gao, Jingjing Liu, Yun-Nung Chen, and Kam-Fai Wong. 2018. Adversarial advantage actor-critic model for task-completion dialogue policy learning. In *IEEE International Conference on Acoustics, Speech and Signal Processing*.
- Jost Schatzmann, Blaise Thomson, Karl Weilhammer, Hui Ye, and Steve Young. 2007. Agenda-based user simulation for bootstrapping a pomdp dialogue system. In *NAACL (Vol. 2: Short Papers)*.
- Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. 2015. Prioritized experience replay. In *International Conference on Learning Representations*.
- Richard S Sutton. 1988. Learning to predict by the methods of temporal differences. *Machine learning*.
- Hado Van Hasselt, Arthur Guez, and David Silver. 2016. Deep reinforcement learning with double q-learning. In *AAAI*.
- Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Van Hasselt, Marc Lanctot, and Nando De Freitas. 2016. Dueling network architectures for deep reinforcement learning. In *International Conference on Machine Learning*.
- Joseph Weizenbaum. 1966. ELIZA: A Computer Program for the Study of Natural Language Communication Between Man and Machine. *Communications of the ACM*, 9(1).
- Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gašić, Lina M Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017. A network-based end-to-end trainable task-oriented dialogue system. In *EACL*.
- Jason Williams, Antoine Raux, and Matthew Henderson. 2016. The dialog state tracking challenge series: A review. *Dialogue & Discourse*, 7(3):4–33.
- Jason D Williams and Steve Young. 2007. Partially observable markov decision processes for spoken dialog systems. *Computer Speech & Language*, 21(2):393–422.
- Zhou Yu, Alan W Black, and Alexander I Rudnicky. 2017. Learning conversational systems that interleave task and non-task content. In *International Joint Conference on Artificial Intelligence*.
- Jiaping Zhang, Tiancheng Zhao, and Zhou Yu. 2018. Multimodal hierarchical reinforcement learning policy for task-oriented visual dialog. In *SigDial*.

A Hyperparameters

All neural network layers are fully connected linear layers with ReLUs as activation functions. In case of the dueling network architecture, the shared layer consists of 256 neurons, followed by two value layers, each with 300 neurons, and two advantage layers with 400 neurons per layer. Distributional agents use an atom count of 50. Where the dueling architecture is replaced by a standard architecture in the evaluation process, three layers of sizes 256, 700 and 700 are used to guarantee a fair comparison to the dueling case by providing the same model capacity. For prioritized replay, the prioritization exponent α is set to 0.525 and importance sampling exponent β to 0.4 (Schaul et al., 2015). To train the networks we use the Adam optimizer with a learning rate of 10^{-4} . Exploration is performed ϵ -greedy with linear ϵ decay, starting at 0.3. Whenever an agent makes use of double Q-learning, it updates its target network after 6 dialogs. All agents use an experience replay buffer capacity of 16384 transitions, a discount factor $\gamma = 0.99$, minibatch size 256 and the Huber Loss κ is set to 1. Following the PyDial benchmarking process, we leave all hyperparameters constant across all environments and dialog domains (Casanueva et al., 2017), thus also evaluating the generalization capabilities of the agents.

B Example Dialogs

Tables 6 and 7 show example dialogs between the *Rainbow* agent as well as the handcrafted policy and a user recorded in the LAP domain as qualitative evaluation. The *Rainbow* agent was trained in a noisy environment with 15% semantic error rate. The user’s task was to find a cheap, light notebook with medium-sized hard drive and arbitrary battery rating. User responses were kept identical for all tested policies whenever possible. Both the combined and the handcrafted agents succeeded in fulfilling the user’s goal with the same amount of dialog turns.

Turn	Actor	Utterance
1	System	Hello, welcome to the laptop information system, what kind of laptop are you after?
1	User	Hi, I’m looking for a cheap and light laptop
2	System	Would you like the hard drive to be small, medium, or large?
2	User	Medium should suffice
3	System	Sorry would you like a product for business computing or not?
3	User	Yes, please
4	System	The satellite-pro-nb10-a-10p is in the budget price range, is in the light weight range and is in the medium drive range.
4	User	Thank you, bye
5	System	Thank you, goodbye.

Table 6: Recorded dialog between a human and the *Rainbow* agent for task T4.3.

Turn	Actor	Utterance
1	System	Hello, welcome to the laptop information system, what kind of laptop are you after?
1	User	Hi, I’m looking for a cheap and light laptop
2	System	Would you like your battery rating to be standard, good, or exceptional?
2	User	I don’t care
3	System	Would you like the hard drive to be small, medium, or large?
3	User	Medium should suffice
4	System	The satellite-pro-nb10-a-10p is in the light weight range, has a standard battery rating, is in the budget price range and is in the medium drive range.
4	User	Thank you, bye
5	System	Thank you, goodbye.

Table 7: Recorded dialog between a human and PyDial’s handcrafted policy for task T4.3.

Contextualized Representations for Low-resource Utterance Tagging

Bhargavi Paranjape
Carnegie Mellon University
bvp@cs.cmu.edu

Graham Neubig
Carnegie Mellon University
gneubig@cs.cmu.edu

Abstract

Utterance-level analysis of the speaker’s intentions and emotions is a core task in conversational understanding. Depending on the end objective of the conversational understanding task, different categorical dialog-act or affect labels are expertly designed to cover specific aspects of the speakers’ intentions or emotions respectively. Accurately annotating with these labels requires a high level of human expertise, and thus applying this process to a large conversation corpus or new domains is prohibitively expensive. The resulting paucity of data limits the use of sophisticated neural models. In this paper, we tackle these limitations by performing unsupervised training of utterance representations from a large corpus of spontaneous dialogue data. Models initialized with these representations achieve competitive performance on utterance-level dialogue-act recognition and emotion classification, especially in low-resource settings encountered when analyzing conversations in new domains.

1 Introduction

Spontaneous human conversations have been collected in different domains to support research in data-driven dialogue systems (Serban et al., 2015), affective computing (Zadeh et al., 2018; Busso et al., 2008; Park et al., 2014), clinical psychology (Althoff et al., 2016) and tutoring systems (Sinha et al., 2015). These conversations are analyzed by segmenting transcriptions into each speaker’s utterances (Traum and Heeman, 1996), which are often labeled with different types of information. The exact type of label to be used depends on the downstream task or research questions to be answered, and thus the tagging paradigms are varied and numerous. For example, the speaker’s intention can be specified using a dialogue acts (DAs) or speech acts (Searle and Searle, 1969), which capture the pragmatic or semantic function of the utterance.

Utterance	DA
A: Hi	Greeting
B: Hi, How are you?	Greeting
A: Are you done with your homework?	Question
B: Yeah	Yes Answer
B: How about you?	Question
A: I’m having trouble with Q4	Statement
B: Yeah	Backchannel
A: so it’s going to take some time	Statement

Table 1: Snippets of conversation with dialogue act tags. “Yeah” is tagged differently in different contexts.

Utterances may also be tagged with traits such as sentiment, emotion and valence labels (Busso et al., 2008; Zadeh et al., 2018), speaker persuasiveness (Park et al., 2014), speaker dominance (Busso et al., 2008) and other characteristics at the utterance and conversational level.

While these labels vary greatly, one constant is that they are often ambiguous and context-dependent (Table 1), making it challenging for humans to annotate efficiently and accurately. Thus, curating large corpora is labor-intensive, and we are always faced with a paucity of data in new domains and labeling paradigms of interest.

Moreover, the label assigned to an utterance depends on the current state of the dialogue (Stone, 2005) and prediction of an utterance’s label benefits from referring to other utterances in context and their labels (Jaiswal et al., 2019). Deep learning models like RNNs and CNNs have proven effective tools to encode neighbouring utterances (Chen et al., 2018; Liu et al., 2017; Blunsom and Kalchbrenner, 2013; Bothe et al., 2018; Kumar et al., 2017). However such models rely on large annotated corpora that are prohibitively expensive to procure, especially for niche domains.

One recently popular method to overcome the dearth of supervised data in NLP is unsupervised pretraining over large unlabeled corpora. For ex-

ample, Melamud et al. (2016); Peters et al. (2018); Devlin et al. (2018) use language modeling as an unsupervised task to learn word embeddings in context, and demonstrate remarkable improvements on a number of downstream NLP tasks. However, these methods learn representations for individual words, whereas for dialog analysis tasks, we need representations for *utterances* in the context of the entire dialog.

In this paper, we adapt the technique of learning contextualized representations using unsupervised pretraining to learn representations for utterances in the context of the dialogue. We first introduce a general model architecture consisting of a token, utterance, and conversation encoder. We then present a method to efficiently train this model by predicting the *bag-of-words* vectors of previous and next utterances over a large heterogeneous corpus of spoken dialogue transcripts. We quantify the effectiveness of learnt contextual utterance representations on two downstream utterance-labeling tasks: DA tagging and emotion recognition. We obtain competitive performance on two popular DA tagging tasks (SwitchBoard and ICSI Meeting Recorder) and an emotion labeling task (IEMO-CAP). Particularly, we observe significant improvements over training complex utterance tagging models from scratch for simulated low-resource settings for these tasks as well as for considerably smaller DA datasets such as LEGO and Map Task.

2 Methodology

We consider a large collection of conversations, where each conversation \mathcal{C} is an ordered list of N utterances $\mathcal{C} = \{u_1, u_2, \dots, u_N\}$ and each utterance is a list of tokens, $u_i = \{w_1, w_2, \dots, w_{|u_i|}\}$. Conversations may also have labels for every utterance: $Y = \{y_1, y_2, \dots, y_N\}$ where each $y_i \in \mathcal{T}$, a finite set of labels expertly defined for a domain.

2.1 Unsupervised Pretraining

Contextualized Utterance Representations

We adopt a hierarchical encoder model consisting of a token encoder, an utterance encoder and a conversation encoder, followed by an output layer. The token encoder layer ENC_{tok} encodes every token w_j in utterance u_i into a fixed-size embedding $e_{w_j}^{tok}$, while the utterance encoder ENC_{utt} encodes token embeddings of an utterance u_i into a fixed-sized utterance representation $e_{u_i}^{utt}$. For our specific instantiation, we combine both en-

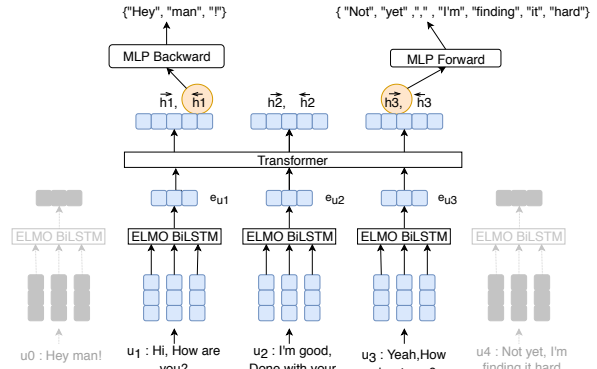


Figure 1: Hierarchical conversation encoder model

coders: we use the pretrained ELMo (Peters et al., 2018) model to encode the sequence of tokens in an utterance u_i and take the final state of the forward and backward LSTMs (concatenated) as our utterance representation $e_{u_i}^{utt}$, $i = 1, 2, \dots, N$. We specifically choose ELMo because it is a strong general-purpose encoder and its character-based representations may be more robust to noise and OOV words in spontaneous conversations. This is followed by a conversation encoder ENC_{conv} , which further converts this sequence of context-independent representations of utterances to a context-dependent sequence of utterance representations. For ENC_{conv} , we use an architecture identical to the decoder variant of the Transformer (Vaswani et al., 2017) with $N = 2$ layers. We specifically choose the self-attentional Transformer for this purpose, as it is efficient to train, can easily capture long-distance dependencies over the entire conversation, and empirically outperformed other architectures such as LSTMs in preliminary experiments. The outputs, h_{u_i} , $i = 1, 2, \dots, N$, of this hierarchical encoder of Figure 1 can be used as contextualized representations for utterances.

Predicting Utterance Bag-of-words In order to learn contextualized representations, the hierarchical encoder is trained to predict the bag-of-words of the previous and next utterances in the conversation using these representations. This training is done in the forward and backward direction respectively by allowing the self-attention layer of the transformer to only attend to earlier positions and later positions in the utterance sequence respectively (Figure 1). Hence, we learn *contextual utterance embeddings* in both directions: $\overleftarrow{h_{u_i}}, \overrightarrow{h_{u_i}}$; $i = 1, 2, \dots, N$. We use an MLP followed by sigmoid function as the output layer over

Corpus	# Utterances	# Tokens
SwitchBoard	460K	3M
Meeting Recorder	105K	11K
CALLHOME	27K	1M
AMI Meeting Corpus	150K	1M
BNC	1M	10M

Table 2: List of dialogue corpora for pretraining contextualized utterance representations

h_{u_i} to predict the set of words in the neighboring utterance. u_{i-1} is reconstructed from $\overleftarrow{h_{u_i}}$ and u_{i+1} from $\overrightarrow{h_{u_i}}$. We use binary cross entropy (BCE) loss, where the target is a vocabulary-sized binary vector with words present in the utterance marked 1 and others 0. Notably this formulation reduces training time by relaxing word-order in the reconstruction loss, unlike other methods that predict words in order for surrounding utterances (Kiros et al., 2015). For utterances u_{i-1} and u_{i+1} with vocabulary vectors U_{i-1} and $U_{i+1} \in \{0, 1\}^{|V|}$ respectively, the *bag-of-word* loss for utterance u_i is given by:

$$\mathcal{L}_{BOW}(u_i) = BCE(\text{MLP}(\overleftarrow{h_{u_i}}), U_{i-1}) + BCE(\text{MLP}(\overrightarrow{h_{u_i}}), U_{i+1}). \quad (1)$$

where,

$$BCE(\mathbf{x}, \mathbf{y}) = \sum_n^{|V|} [y_n \log(x_n) + (1 - y_n) \log(1 - x_n)]$$

For conversation $\mathcal{C} = \{u_1, u_2, \dots, u_N\}$,

$$\mathcal{L}_{BOW}(\mathcal{C}) = \frac{1}{N} \sum_{i=0}^N \mathcal{L}_{BOW}(u_i). \quad (2)$$

2.2 Utterance Tagging

Once we have learned contextualized utterance representations, we can use them to predict the sequence of labels $Y = \{y_1, y_2, \dots, y_N\}$, such as dialogue acts, for utterances in the conversation. In this work we use a linear-chain conditional random field (Lafferty et al., 2001) as used in previous state-of-the-art models for DA tagging (Kumar et al., 2017; Chen et al., 2018) to predict one of the $|\mathcal{T}|$ tags for each u_i , where the utterance is represented as the concatenation of the forward and backward contextualized vectors: $\overleftarrow{h_{u_i}}, \overrightarrow{h_{u_i}}$.

3 Experiments

Pretraining Datasets and Hyperparameters

We train contextualized utterance representations

on transcriptions of spontaneous human-human conversation corpora (Serban et al., 2015). We choose the corpora presented in Table 2 for this work. A majority of the conversations are dialogues, and utterances across all corpora are 10 words long on average. However, the chosen corpora have conversations of widely varying lengths (no. of utterances/conversation). For computation/memory efficiency, and also because more distant utterances likely have diminishing influence on discourse modeling, we divide each conversation into conversational snippets of length 64¹ by moving a 64-length window over the conversation with stride 1 and train the bag-of-word loss on each snippet thus obtained. For the conversational encoder, we use 2 layers of the transformer with 8 attention heads of 64 dimensions each. All feed-forward networks use 2 layers with hidden size of 512. For training and fine-tuning, we use the Adam (Kingma and Ba, 2014) with learning rate 0.0001.

Tasks We evaluate performance of our model on these utterance-level tagging tasks:

SwDA, the Switchboard Dialogue Act Corpus, annotates 1,155 telephonic conversations (224K utterances) with one of the 42 DAs in the DAMSL (Jurafsky, 1997) taxonomy.

MRDA, the ICSI Meeting Recorder Dialogue Act corpus annotates 75 multi-party meetings (105K utterances) with DAs according to 5 domain-specific tags (Dhillon et al., 2004).

IEMOCAP, an emotion recognition dataset of 12 hours of dyadic improvisations or scripted scenarios, with eight categorical emotion labels (Park et al., 2014) (10K utterances).

LEGO, a subset (14K utterances) of the Lets Go bus-information dialogue system corpus (Raux et al., 2006) annotated with the ISO 24617-2 standard for conversation functions of task by (Ribeiro et al., 2016).

Map Task, (Carletta et al., 1997; Anderson et al., 1991) is 18 hrs of dialogue where speakers collaborate to complete a map (5K utterances).

To simulate low-resource settings for the larger datasets like SWDA and MRDA, we experiment with different sizes of the training datasets and evaluate on the standard test set for these. For LEGO and MapTask, we use 10-fold cross validation.

Experimental Settings We use four different experimental settings to measure the efficacy of our

¹tuned model hyper-parameter




Figure 2: Performance by training data sizes. SOTA: comparable state-of-the-art model trained on tagging task for entire dataset.

pretrained utterance representations : *No Context* - With no conversational encoder (i.e. independently encoding every utterance using ELMo); *Random Initialization* - with the conversational encoder randomly initialized and trained on only the downstream tagging task; *Freeze Network* - the conversational encoder initialized using the model pretrained on our bag-of-words objective and kept fixed for downstream task; *Pre-trained Initialization* - the initialized conversation encoder fine-tuned on the downstream task. These settings are used to isolate the gains from using (1) contextualized representations, (2) pretraining them and then (3) fine-tuning them on the downstream task.

4 Result and Discussion

We observe that using pretrained utterance representations shows improved performance over random initialization and is competitive with existing state-of-the-art works by Kumar et al. (2017) for SwDA and MRDA, and Poria et al. (2017) for IEMOCAP that use similar hierarchical architec-

DA Category	% Increase in accuracy	Example
Agree/Accept	43	That's exactly it.
Summarize/Reformulate	180	Oh, you mean you switched schools..
Statement-Opinion	55	I think it's great.
Yes-Answer	33	Yes
Hold before answer or agreement	300	I'm drawing a blank

Table 3: SwDA DA categories that improve using pretrained utterance embeddings with % improvements in accuracy over other experimental settings.

DA Corpus	Pretrained	Random	SOTA
LEGO	93.70	92.98	88.75
Map Task	79.34	77.91	72.50

Table 4: Results on LEGO and Map Task

tures but are only trained on the task (Random initialization setting). From Figure 2, we observe that the pretraining-based initialization is especially helpful when the amount of training data is significantly reduced for SWDA, MRDA and IEMOCAP, over other experimental settings. The improved performance of the random initialization setting over fixing the pretrained conversational encoder parameters underscores the need to fine-tune for downstream tasks. Our pretrained model also outperforms random initialization and existing best results (Ribeiro et al., 2015; Sridhar et al., 2009) for truly low-resource datasets like LEGO and Map Task, as shown in Table 4. We also analyze the gain in accuracy by dialogue act category for the pretrained model over other experimental settings. We find that the pretrained model shows improvements in the categories listed in Table 3 over random initialization. These acts typically require models to keep track of longer contexts than other DAs like questions and back-channels. Dialogue examples in Table 5 further illustrate this.

5 Conclusion

We show that using large dialogue corpora to train contextualized utterance embeddings using a bag-of-words reconstruction loss is beneficial for utterance-level tagging in the low-resource setting, indicating that these embeddings learn useful and generalizable properties of conversational discourse. Future work involves incorporating speaker identity, utterance duration and speech/prosody features.

Utterance	Gold	Pre-trained	Random	No Context
B: where are you going to move to?	Wh-Question	Wh-Question	Wh-Question	Yes-No-Q
A: Uh, Maryland.	Statement	Statement	Statement	Hedge
B: Oh, are you?	Backchannel	Backchannel	Backchannel	Yes-No-Q
	question	question	question	
A: Uh-huh.	Yes answers	Yes answers	Yes answers	Backchannel
B: Do you have friends there?	Yes-No-Q	Yes-No-Q	Yes-No-Q	Yes-No-Q
B: or,	Abandoned	Abandoned	Abandoned	Uninterpret.
A: My fiancée is down there ⟨laughter⟩.	Statement	Statement	Statement	Statement
B: Oh, I see.	Resp. Ack	Resp. Ack	Resp. Ack	Resp. Ack
B: So, does he work for a company down there?	Yes-No-Q	Yes-No-Q	Yes-No-Q	Yes-No-Q
A: Yeah,	Yes answers	Yes answers	Yes answers	Yes answers
A: he works for the government.	Statement	Statement	Statement	Statement
B: Oh, I see.	Resp. Ack	Resp. Ack	Resp. Ack	Resp. Ack
B: Oh, the big company.	Summarize/ reformulate	Summarize/ reformulate	Statement	Statement
A: Yeah	Agree/Accept	Yes Answer	Yes Answer	Yes Answer
A: and I said no, I'm just twenty-three,	Statement	Statement	Statement	Statement
B: Uh-huh.	Backchannel	Backchannel	Backchannel	Abandoned
A: you know, because I don't think of myself as needing to have children	Statement	Statement	Statement	Statement
A: but the first thing he says is, well, don't you miss that part of your life.	Statement	Statement	Statement	Statement
A: And I just,	Abandoned	Abandoned	Abandoned	Uninterpret.
A: my, my mind just went,	Statement	Statement	Statement	Statement
B: You didn't know what you're going to be missing.	Collaborative Completion	Collaborative Completion	Statement	Statement
A: I went, what.	Statement	Statement	Statement	Statement
B: ⟨Laughter⟩.	Non-verbal	Non-verbal	Non-verbal	Non-verbal

Table 5: Dialogue Examples from SwitchBoard with dialogue acts as labelled under different experimental settings. The pre-trained network performs better on categories like Summarizing and Collaborative Completion

Acknowledgments

We thank the reviewers for their insightful comments. This work was supported by the National Institute of Health (NIH) grant no. R01MH096951-08.

References

Tim Althoff, Kevin Clark, and Jure Leskovec. 2016. Large-scale analysis of counseling conversations: An application of natural language processing to mental health. *Transactions of the Association for Computational Linguistics*, 4:463–476.

Anne H Anderson, Miles Bader, Ellen Gurman Bard, Elizabeth Boyle, Gwyneth Doherty, Simon Garrod, Stephen Isard, Jacqueline Kowtko, Jan McAllister,

Jim Miller, et al. 1991. The hcrc map task corpus. *Language and speech*, 34(4):351–366.

Phil Blunsom and Nal Kalchbrenner. 2013. Recurrent convolutional neural networks for discourse compositionality. In *Proceedings of the 2013 Workshop on Continuous Vector Space Models and their Compositionality*. Proceedings of the 2013 Workshop on Continuous Vector Space Models and their .

Chandrakant Bothe, Cornelius Weber, Sven Magg, and Stefan Wermter. 2018. A context-based approach for dialogue act recognition using simple recurrent neural networks. *arXiv preprint arXiv:1805.06280*.

Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeanette N Chang, Sungbok Lee, and Shrikanth S Narayanan. 2008. Iemocap: Interactive emotional dyadic motion capture database. *Language resources and evaluation*, 42(4):335.

- Jean Carletta, Amy Isard, Stephen Isard, Jacqueline C Kowtko, Gwyneth Doherty-Sneddon, and Anne H Anderson. 1997. The reliability of a dialogue structure coding scheme. *Computational Linguistics*.
- Zheqian Chen, Rongqin Yang, Zhou Zhao, Deng Cai, and Xiaofei He. 2018. Dialogue act recognition via crf-attentive structured network. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 225–234. ACM.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Rajdip Dhillon, Sonali Bhagat, Hannah Carvey, and Elizabeth Shriberg. 2004. [Meeting recorder project: Dialog act labeling guide](#).
- Mimansa Jaiswal, Zakaria Aldeneh, Cristian-Paul Bara, Yuanhang Luo, Mihai Burzo, Rada Mihalcea, and Emily Mower Provost. 2019. Muse-ing on the impact of utterance ordering on crowdsourced emotion annotations. *arXiv preprint arXiv:1903.11672*.
- Dan Jurafsky. 1997. Switchboard swbd-damsl shallow-discourse-function annotation coders manual. *Institute of Cognitive Science Technical Report*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. In *Advances in neural information processing systems*, pages 3294–3302.
- Harshit Kumar, Arvind Agarwal, Riddhiman Dasgupta, Sachindra Joshi, and Arun Kumar. 2017. Dialogue act sequence labeling using hierarchical encoder with crf. *arXiv preprint arXiv:1709.04250*.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. [Conditional random fields: Probabilistic models for segmenting and labeling sequence data](#). In *Proceedings of the Eighteenth International Conference on Machine Learning, ICML '01*, pages 282–289, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Yang Liu, Kun Han, Zhao Tan, and Yun Lei. 2017. Using context information for dialog act classification in dnn framework. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2170–2178.
- Oren Melamud, Jacob Goldberger, and Ido Dagan. 2016. [context2vec: Learning generic context embedding with bidirectional lstm](#). In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 51–61, Berlin, Germany. Association for Computational Linguistics.
- Sunghyun Park, Han Suk Shim, Moitreyia Chatterjee, Kenji Sagae, and Louis-Philippe Morency. 2014. [Computational Analysis of Persuasiveness in Social Multimedia: A Novel Dataset and Multimodal Prediction Approach](#). In *Proceedings of the 16th International Conference on Multimodal Interaction*, pages 50–57. ACM Press.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, volume 1, pages 2227–2237.
- Soujanya Poria, Erik Cambria, Devamanyu Hazarika, Navonil Majumder, Amir Zadeh, and Louis-Philippe Morency. 2017. Context-dependent sentiment analysis in user-generated videos. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 873–883.
- Antoine Raux, Dan Bohus, Brian Langner, Alan W Black, and Maxine Eskenazi. 2006. Doing research on a deployed spoken dialogue system: One year of let’s go! experience. In *Ninth International Conference on Spoken Language Processing*.
- Eugénio Ribeiro, Ricardo Ribeiro, and David Martins de Matos. 2015. The influence of context on dialog act recognition. *arXiv preprint arXiv:1506.00839*.
- Eugénio Ribeiro, Ricardo Ribeiro, and David Martins de Matos. 2016. Mapping the dialog act annotations of the lego corpus into the communicative functions of iso 24617-2. *arXiv preprint arXiv:1612.01404*.
- John R Searle and John Rogers Searle. 1969. *Speech acts: An essay in the philosophy of language*, volume 626. Cambridge university press.
- Iulian Vlad Serban, Ryan Lowe, Peter Henderson, Laurent Charlin, and Joelle Pineau. 2015. A survey of available corpora for building data-driven dialogue systems. *arXiv preprint arXiv:1512.05742*.
- Tanmay Sinha, Ran Zhao, and Justine Cassell. 2015. Exploring socio-cognitive effects of conversational strategy congruence in peer tutoring. In *Proceedings of the 1st Workshop on Modeling INTERPERSONAL Synchrony And Influence*, pages 5–12. ACM.
- Vivek Kumar Rangarajan Sridhar, Srinivas Bangalore, and Shrikanth Narayanan. 2009. Combining lexical, syntactic and prosodic cues for improved online dialog act tagging. *Computer Speech & Language*, 23(4):407–422.
- Matthew Stone. 2005. Communicative intentions and conversational processes in humanhuman and human-computer dialogue. *Approaches to studying world-situated language use*, pages 39–70.

David R Traum and Peter A Heeman. 1996. Utterance units in spoken dialogue. In *Workshop on Dialogue Processing in Spoken Language Systems*, pages 125–140. Springer.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.

AmirAli Bagher Zadeh, Paul Pu Liang, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 2018. Multimodal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 2236–2246.

Capturing Dialogue State Variable Dependencies with an Energy-based Neural Dialogue State Tracker

Anh Duong Trinh [†], Robert J. Ross [†], John D. Kelleher [‡]

[†] School of Computer Science

[‡] Information, Communications & Entertainment Institute

Technological University Dublin, Ireland

ADAPT Centre, Science Foundation Ireland

anhduong.trinh@mydit.ie, {robert.ross, john.d.kelleher}@dit.ie

Abstract

Dialogue state tracking requires the population and maintenance of a multi-slot frame representation of the dialogue state. Frequently, dialogue state tracking systems assume independence between slot values within a frame. In this paper we argue that treating the prediction of each slot value as an independent prediction task may ignore important associations between the slot values, and, consequently, we argue that treating dialogue state tracking as a structured prediction problem can help to improve dialogue state tracking performance. To support this argument, the research presented in this paper is structured into three stages: (i) analyzing variable dependencies in dialogue data; (ii) applying an energy-based methodology to model dialogue state tracking as a structured prediction task; and (iii) evaluating the impact of inter-slot relationships on model performance. Overall, we demonstrate that modelling the associations between target slots with an energy-based formalism improves dialogue state tracking performance in a number of ways.

1 Introduction

Dialogue management for spoken dialogue systems is a challenging research domain due in part to difficulties arising from limited resources, the imperfection of technologies on which dialogue management is dependent, and of course the complexities of natural human conversation (Glass, 1999; Ward and DeVault, 2015). Within a conventional dialogue manager, an explicit dialogue state tracker is a key component that attempts to track both interlocutors' contributions to the exchange. The dialogue state tracker in particular suffers due to errors introduced by other components such as an automatic speech recognizer, and, where used, natural language understanding components (Ross and Bateman, 2009). The diffi-

culties also lie within the uncertainties of spoken interactions, and the complexity of conversation context (Paek and Horvitz, 2000; DeVault, 2008).

To reduce the complexity of designing and parameterising a dialogue state tracker, it is typically necessary to limit application of a dialogue state tracker to a specific domain, and to cast the dialogue state as sets of slot-value pairs that are arranged into frames. This structure in its base case is best exemplified by the well-known dialogue state tracking datasets such as Let's Go (Raux et al., 2005), though the structure can also be made more complex as is the case in the tracking of multiple frames of dialogue states throughout the conversation history (El Asri et al., 2017). By casting the dialogue representation as a set of slots to be tracked, the dialogue state tracking process itself is most frequently tackled as a multi-task classification problem.

In recent years, various deep learning approaches that track dialogue states as a combination of individual classification tasks have been proposed (Ren et al., 2018; Perez and Liu, 2017; Vodolan et al., 2017; Mrksic et al., 2017; Ras-togi et al., 2017). However, while these systems achieve state-of-the-art results, there remains notable room for improvement (Liu et al., 2018). Our work begins with the hypothesis that by treating dialogue state tracking as a simple multi-label classification task, we are not taking into account the relationships between dialogue state slot variables. This hypothesis is based in part on experience from other applications of machine learning that have demonstrated that taking target variable dependencies into account is useful, but is also based on the intuition that a human interlocutor would of course take multiple target variables into account while interpreting language (Landragin, 2013).

Given the above argument, in this paper we

present an end-to-end investigation into the impact of domain variable dependencies on the dialogue state tracking process. For practical purposes, we focus our work on the Dialogue State Tracking Challenge (DSTC) series that were introduced to help the research community focus on the specific task and subsequently improve the quality of spoken dialogue systems (Williams et al., 2016). Specifically, our investigation is conducted with respect to the second and the third dialogue state tracking challenges (Henderson et al., 2014a,b), and is presented in three stages:

- **Data analysis** - We perform statistical tests on the dialogue data to determine whether there are indeed dependencies between slot variables and to what extent are these dependencies present.
- **Model development** - Tracking dialogue states while considering the relationships between target variables casts the problem into a structured prediction task. We develop a deep learning based tracker that incorporates an energy-based modelling approach that is notably efficient for structured predictions.
- **Result analysis** - Our model performance is evaluated and analyzed using a number of metrics to provide insights into the impact of variable dependencies on the dialogue state tracking process. We benchmark our energy-based approach against results for a number of state-of-the-art systems (Vodolan et al., 2017; Mrksic et al., 2015; Henderson et al., 2014c,d).

To our knowledge there has been no detailed analysis previously on the role of variable dependencies in dialogue states. The contributions of this paper are, thus, that systematic analysis, and our energy-based structured prediction model for dialogue state tracking.

2 Categorical Data Analysis

The investigation presented in this paper is predicated on the existence of associations between target variables in a dialogue state. Therefore, in this section we provide a concrete analysis of variable dependencies between domain slots in DSTC data.

2.1 DSTC 2 & 3 Datasets

The Dialogue State Tracking Challenge 2 & 3 datasets contain phone calls in the restaurant and

tourism information domains (Henderson et al., 2014a,b). Within the datasets, the main task is referred to as *Joint Goals* and requires systems to estimate the value of each slot in the set of informable slots at every turn of the call. The value constraint is retrieved from the set of possible values predefined in a specified domain ontology.

The DSTC2 dataset is split into three subsets: 1612 dialogues for training; 506 for validation; and 1117 for a test set. The DSTC2 ontology pre-defines four informable slots.

The DSTC3 dataset contains 2275 dialogues that are not split into subsets; the dataset defines nine informable slots in the ontology. Four of the nine slot types also appear in the DSTC2 data, but the value sets are different.

A preliminary analysis shows that these slots are not equally distributed in both datasets (see Table 1). The informable slots are divided into two groups with one group including highly frequent slots ($f > 50\%$) and the other one containing very low frequencies ($f < 10\%$). Therefore, we follow the precedent of other researchers and design our models to track only highly frequent slots. Following this reduction, the DSTC2 *Joint Goals* consist of three slot-value pairs (*food*, *price range*, *area*), and DSTC3 *Joint Goals* consist of four slot-value pairs (*food*, *price range*, *area*, and *type*).

Slot	DSTC2		DSTC3	
	call	turn	call	turn
food	87.9	79.3	63.5	55.4
price range	73.5	62.6	68.3	60.8
area	81.8	72.3	59.5	50.6
type	-	-	98.5	91.0
name	0.8	0.5	1.5	0.6
near	-	-	8.5	6.8
has tv	-	-	7.3	5.8
has internet	-	-	7.6	5.9
children allowed	-	-	4.9	3.6

Table 1: The analysis of informable slot proportions (%) in DSTC 2 & 3 summarised over the number of calls and turns in the whole dataset.

2.2 Variable Dependencies

To test the independence of the slot variables, we apply Pearson’s chi-square tests on labels of the informable slots in a pairwise fashion to generate bivariate statistics. The dependencies between slots are confirmed if and only if the significance

DSTC2		food - price	food - area	price - area
	χ^2	9430.5	12739.0	3937.9
Chi-square	\mathcal{V}	176	180	24
	p	< 2.2e-16	< 2.2e-16	< 2.2e-16
Coefficients	ϕ	0.6081	0.7068	0.3930
	C	0.5196	0.5772	0.3657
	V	0.2720	0.2671	0.1757

Table 2: Statistical tests on DSTC2 dataset.

DSTC3		food - price	food - area	food - type	price - area	price - type	area - type
	χ^2	5792.6	7985.6	6762.5	5070.7	2873.0	3626.5
Chi-square	\mathcal{V}	145	464	116	80	20	64
	p	< 2.2e-16	< 2.2e-16	< 2.2e-16	< 2.2e-16	< 2.2e-16	< 2.2e-16
Coefficients	ϕ	0.5547	0.6513	0.5994	0.5190	0.3907	0.4389
	C	0.4851	0.5458	0.5141	0.4607	0.3639	0.4019
	V	0.2265	0.1580	0.2680	0.2119	0.1747	0.1963

Table 3: Statistical tests on DSTC3 dataset.

value $p < 0.05$. The chi-square test results are reported with the χ^2 statistic, degree of freedom \mathcal{V} , and statistical significance p . The statistic is calculated with the formula:

$$\chi_{\mathcal{V}}^2 = \sum_i \sum_j \frac{(O_{ij} - E_{ij})^2}{E_{ij}} \quad (1)$$

where O_{ij} and E_{ij} are observed and expected frequencies of categories i and j being activated for the observed variables at the same time in the whole dataset.

Furthermore, it is necessary to measure the strength of these dependencies as the chi-square test can only detect the presence of the dependencies without saying if they are strong or not. Fortunately, there exist several chi-square test-based measurements of association strength between variables. We report three such measures: ϕ coefficient, contingency coefficient C , and Cramer’s V coefficient. All the coefficients are calculated through adjustment of the chi-square statistic to account for the dataset size; for instance:

$$V = \sqrt{\frac{\chi^2}{N \min(r-1, c-1)}} \quad (2)$$

where χ^2 is the chi-square statistic, N is the number of samples in the dataset, and r and c are the number of rows and columns in the contingency table. These measures are scaled between 0 and 1 indicating that 1 is the perfect relationship and

0 indicates the lack of any relationship between variables.

We report the statistics analysis of DSTC2 data in Table 2 and DSTC3 data in Table 3. In the results, all variables showed significance values $p < 0.05$, that indicate that there are indeed variable dependencies in the dialogue domains of the DSTC series. The association strength measured by the chi-square based coefficients show different level of variable dependencies ranging from a very strong dependency ($\phi \geq 0.7$, $V \geq 0.25$) to a moderate one ($0.3 \leq \phi < 0.39$, $0.11 \leq V < 0.15$). For example in DSTC2 data, the dependencies between the slot *food* and the other two, *price range* and *area*, are strong.

While the existence of dependencies across our labels may not be surprising, the consistency of their strong occurrence indicates that tracking systems could achieve more accurate results if judgements on trackable slots were made with reference to the information contained within hypotheses for neighbouring slots.

3 Energy-Based Structured State Tracking

The data analysis performed on the DSTC series data suggests that incorporating label dependencies in the dialogue state prediction process would be beneficial. Formally this indicates that we should cast the dialogue state tracking process as a structured prediction problem (Smith, 2009). This

in itself should not be a surprise to the research community, as several researchers have built dialogue state trackers around models that can in principle be thought of as structured classifiers (Zhong et al., 2018; Hori et al., 2016; Jang et al., 2016; Ren et al., 2013).

One of the challenges for previous approaches to structured prediction for dialogue state classification is that they relied on methods that had difficulty integrating a structural component that took inter-slot dependencies into account with a robust underlying classifier that facilitated powerful feature representations from individual contributions to the dialogue. Recently the application of energy-based methods that are implemented through neural architectures have provided one promising avenue for structured prediction. The idea underpinning this approach is that we learn to rate the association between configurations of target variables and our inputs via a so-called energy function (LeCun et al., 2006) rather than attempt to learn to predict the structured output directly.

Below we first introduce the key principles behind energy-based structured prediction, then detail the energy-based dialogue state tracker that we have constructed.

3.1 Energy-Based Structured Prediction

Let us denote the input and structured output variables as X and Y respectively. For us, X can be thought of as the representation of a turn, while Y is a complete dialogue state representation – not the representation of an individual slot. Given X and Y , a function $E(X, Y)$ must be trained to assign some scalar value called energy to any configuration of variables X and Y . This function is called the energy function, and is traditionally designed to assign low energy to correct variable configurations, and higher energy to incorrect configurations. In other words we have low energy when a hypothesis for Y comes close to the ground truth given an input X . At run-time some interpretation process moves through the space of target configurations to find the most appropriate output configuration for a given input.

While the energy function can be thought of as some arbitrary scalar that is to be low for acceptable configurations, the form of the function and training of the function are important. Specifically the energy function takes the following form:

$$E(x, y, \theta) = E_{global}(y, \theta) + E_{local}(x, y, \theta) \quad (3)$$

where θ are trainable parameters of the energy network, $E_{global}(y, \theta)$ is the global energy term for labels y , and $E_{local}(x, y, \theta)$ is the local energy adjustment of both input and output variables. Thus the global energy function specifically considers the acceptability of configurations of the structured target, while the local energy estimates the appropriateness of the input with respect to individualised elements of the prediction.

During training the parameters θ for the energy function are estimated. This is most efficiently done by coupling the energy function to an oracle loss that estimates the loss between a hypothesised output Y and the ground truth label Y^* for a given input X .

Finding the parameters of a good energy function between X and Y directly is generally however not feasible, and historically was one of the key limitations for energy-based structured prediction. Instead it is generally more appropriate to first generate some feature function $F(X)$ that transforms the input to an appropriate representation form that better supports the inference process. Thus more commonly we denote the energy function as $E(F(X), Y)$. Both the feature representation and the energy function itself can be trained through a deep neural network model either dependently or independently.

3.2 Dialogue State Tracker

Based on the principles of energy-based structured prediction, we have designed an energy-aware dialogue state tracker. The framework for training and applying the energy-based method is based specifically on the Deep Value Network architecture proposed by Gygli et al. (2017). The architecture of our tracking model is illustrated in Figure 1.

The energy-based dialogue state tracker can be thought of as consisting of four key elements with associated training and inference processes; we detail these below.

3.2.1 Feature Function Network

The Feature Function Network $F(X)$ is a deep learning network to process raw DSTC dialogue data into a representation that is suitable for feeding into the energy network. As DSTC dialogues contain different input channels we implement different techniques to accommodate the variety of input variables.

In detail, each input of a dialogue turn consists




Figure 1: Deep Value Network-based Dialogue State Tracking Model.

of machine acts in a semantic format and user utterance transcribed by an automatic speech recognizer. We parse the machine dialogue acts with the parsing technique by Henderson et al. (2014d) before reducing the dimensionality of the machine act vectors with two dense neural layers. Meanwhile, all the words in user utterances are embedded with an online trained embedding layer, then passed into a bidirectional LSTM layer (Hochreiter and Schmidhuber, 1997). The output vectors of this bidirectional LSTM layer represent user utterances as real-valued tensors. Following that, the machine act and utterance vectors are concatenated, and fed into a unidirectional LSTM layer that processes dialogue by turn and returns fixed-size dialogue vector representations.

We pre-train this feature network as a multi-task classification model following the method proposed by Trinh et al. (2018). The dialogue representations retrieved from this network are treated as input features for the energy function.

3.2.2 Energy Function Network

The energy function network $E(F(X), Y)$ is implemented as a feed-forward network (Belanger and McCallum, 2016) where the general function form as illustrated in the previous section is hard coded and the parameters are acquired during training. Based on the energy function proposed by Belanger and McCallum (2016), the general forms of the global and local energy functions are:

$$E_{global}(Y) = W_2^\top f(W_1^\top Y) \quad (4)$$

$$E_{local}(X, Y) = \sum_{i=1}^L y_i W_i^\top F(X) \quad (5)$$

where $\theta = \{W, W_1, W_2\}$ are the energy network's trainable parameters, $f(\cdot)$ is a non-linearity function, $F(\cdot)$ is the feature function described in the previous section, and L is the number of classes in the target.

This energy network produces a scalar energy value that is the sum of global and local energy terms for each input and output configuration.

3.2.3 Oracle Function

To train the energy function we need a signal that estimates the compatibility of an input variable X with an output configuration Y . We achieve this by making use of an oracle function $E^*(Y, Y^*)$ that measures the quality of any output variable configuration Y with respect to the ground truth label Y^* . We implement the oracle value function in our model with the F_1 metric:

$$E_{F_1}^*(y, y^*) = \frac{2(y \cap y^*)}{(y \cap y^*) + (y \cup y^*)} \quad (6)$$

where $y \cap y^* = \sum_i \min(y_i, y_i^*)$; and $y \cup y^* = \sum_i \max(y_i, y_i^*)$, that are extended for continuous output variables.

3.2.4 Objective Function

To train and estimate the energy function, we make use of an objective function $L(E, E^*)$. This function calculates the error between predicted energy $E(X, Y)$ and ground truth energy value that is tied

to the oracle value $E^*(Y, Y^*)$. Since the F_1 score falls into the range $[0, 1]$ we design the objective function as a cross entropy loss function:

$$L = -E^* \log E - (1 - E^*) \log(1 - E) \quad (7)$$

3.3 Training Process

The training process for the energy-based dialogue state tracking model is summarized in Algorithm 1. The learning objective is to train the energy function to predict correct quality of output by shaping the energy values to oracle F_1 values. All the trainable parameters of the network are updated via standard backpropagation techniques.

Algorithm 1: Learning process algorithm

Function *TRAIN_EPOCH* (*dataset* \mathcal{D} , *initial weights* θ , *learning rate* λ)

while *not end of* \mathcal{D} **do**

Training sample

$(x, y^*) \in \mathcal{D}$

Output generation

$y \leftarrow \text{GENERATE}(x, \theta)$

Ground truth energy

$E^* \leftarrow E^*(y, y^*)$

Predicted energy

$E \leftarrow E(x, y, \theta)$

Objective function

$L \leftarrow L(E, E^*)$

Backpropagation

$\theta \leftarrow \theta - \lambda \nabla_{\theta} L$

end

end

In detail, for each iteration in a training epoch we generate a batch of dialogues from the dataset. A structured output of each turn in the dialogue is then generated through an inference process (see Section 3.4). The system predicts energy terms for these variable configurations, and calculated oracle values as the ground truth energies. We compute the loss value of the batch, and backpropagate the model based on this loss.

3.4 Inference Process

In the training process a *GENERATE*(\cdot) function was used to come up with a candidate value for Y given a network and input X . This generation process is based in part on the inference process that is used at both training time and run-time to determine a candidate Y for a given X .

The inference process predicts structured output starting from a random initial prediction. The inference process is based on the principle that the gradient of energy with respect to Y can be calculated directly and used to direct a process for selecting Y .

In short, this prediction is generated through an inference loop with the gradient ascent technique for a number of steps:

$$y^{(t+1)} = \mathcal{P}_Y \left(y^{(t)} + \eta \nabla_y E(x, y^{(t)}, \theta) \right) \quad (8)$$

where \mathcal{P}_Y is the projection operation to shape the predicted output to the output variable space $Y = \{y_i\}^L \in \{[0, 1]\}^L$, and η is the learning rate for gradient ascent.

4 Experimental Design

To evaluate the usefulness of the energy-based approach we implemented and trained a tracker based on the model outlined in the previous section against both the DSTC2 and DSTC3 datasets. Training is a two phase process. First, we trained the feature network independently of the energy-based components by casting the feature network as a standard multi-task learning system where each target variable is assumed to be independent of the others. We present the results of this multi-task based model independently, but critically we also then make use of the trained network prior to the output layer as the feature network that is available for training the full energy network. Thus, the second stage of training targets the parameterisation of the energy network once the feature network has already been learned.

As mentioned, the DSTC2 dataset is divided into three subsets for training, validation, and test purposes, while DSTC3 data are provided in a whole set only for the test purpose. Thus we apply DSTC2 directly, but split the DSTC3 dataset into five folds and use cross-validation in the training process. All experiments are run for at least five times to ensure the stability of our results; we report the average performance.

5 Result & Error Analysis

We report our multi-task and energy-based models performance on both the DSTC2 and DSTC3 datasets, and benchmark their performance against state-of-the art systems in Table 4. A state-of-the-art system is selected if it produces highest to date

Model	Entry	DSTC2	DSTC3
Hybrid system (Vodolan et al., 2017)		0.796	-
Web-style ranking system (Williams, 2014)	✓	0.784	-
Multi-domain system (Mrksic et al., 2015)		0.774	0.671
Word-based system (Henderson et al., 2014d)	✓	0.768	-
Unsupervised RNN-based system (Henderson et al., 2014c)	✓	-	0.646
<i>Our work</i>			
Multi-task feature system		0.709	0.531
Energy-based system		0.760	0.622
DSTC baseline	✓	0.719	0.575

Table 4: Performances of state-of-the-art and our dialogue state tracking systems on DSTC 2 & 3 data. The results for *Joint Goals* are reported with Accuracy metric featured in the challenge. The column *Entry* marks the systems submitted to blind evaluation during the competition period.

accuracy on the *Joint Goals* task either during the DSTC competition time or after the competition. We also included the model by Henderson et al. (2014d,c) as it includes data processing techniques that we adopted in our work.

Overall, we find that applying an energy-based algorithm on top of the LSTM enabled slot tracker improves the dialogue state tracking results in term of accuracy by a big margin, 5% for DSTC2 and 9% for DSTC3.

Comparing our work with state-of-the-art systems like the hybrid tracker (Vodolan et al., 2017) and a multi-domain system (Mrksic et al., 2015), the energy-based model has not yet reached their level of performance. However, the hybrid tracker also consists of a feature network and an algorithm to refine the prediction. Vodolan et al. (2017) designed this algorithm with a set of manual rules, while we design the refinement with a deep neural structure and let it learn from the data. With respect to the multi-domain system, we believe it outperforms our energy-based model because of the wider range of data processed by the multi-domain system. Mrksic et al. (2015) trained and combined their models on six datasets of different domains, while we train our energy-based system on a single domain at a time only.

It is also important to note that the web-style ranking system of Williams (2014) was the best entry during the DSTC2 competition, and is not neural-based. It is followed by the word-base tracker (Henderson et al., 2014d) that was developed with a special recurrent neural network architecture. Besides, the word-based system is also notable for its feature parsing technique that is reused in a number of later systems (Henderson

et al., 2014c; Vodolan et al., 2015, 2017; Trinh et al., 2017, 2018) and our work.

5.1 Slot-based Result Analysis

We argue that the feature Accuracy metric in the DSTC series do not provide a full picture of how well a model performs for each slot. Therefore it is necessary to evaluate our work for the individual slots as well as for the joint dialogue states. We conduct a separate evaluation on the result track file and report it in Table 5. Overall our models achieve high accuracy across all informable slots and the joint goals. Here the joint goals accuracy is higher than evaluated with the DSTC evaluation scripts due to the absence of low frequent slots that we omit in our experiments.

We observe that the energy-based model improves the tracking results of all slots both as individual and a joint set. The improvement margin of joint goals is similar to the results measured by the DSTC feature Accuracy metric. The tracking result of individual slots varies from a very small change of 0.3% to a big jump of 7%. These change differences are related to the relative difficulties of the slot. For example slot *food* has the biggest set of possible values, which in turn makes it the most difficult slot to track; it is for this slot that we see the greatest improvement.

5.2 Proportional Reduction in Error

Proportional reduction in error is a statistical test to measure association between two variables on how one can influence the other in the prediction process. For example given variables A and B , this method attempts to evaluate the prediction of A in two ways: predicting A independently; and predicting A with the knowledge of B . Reduction

Dataset	Model	Slot				Joint goals
		food	price	area	type	
DSTC2	Feature system	0.825	0.929	0.919	-	0.717
	Energy-based	0.872	0.938	0.923	-	0.768
DSTC3	Feature system	0.730	0.844	0.781	0.937	0.587
	Energy-based	0.802	0.860	0.817	0.940	0.666

Table 5: Performances of our energy-based dialogue state tracking system. The results are reported per slot and for Joint slots of those present in the task.

in error can be formulated mathematically.

$$\lambda = \frac{E_A - E_{A|B}}{E_A} \quad (9)$$

where E_A is the number of errors in predicting A , and $E_{A|B}$ is the number of errors in predicting A while taking into account B . All errors are assumed to be absolute numbers.

From this formula we can see that λ has the value in the range $[0, 1]$ because $E_{A|B} \leq E_A$ in all cases. If $\lambda = 0$, A and B are completely independent, thus knowing B does not help predicting A better. On the other hand, when $\lambda = 1$, the relationship between A and B is absolute, i.e., that the knowledge of B gives us the perfect prediction of A .

To apply this statistical method in our model performance evaluation, we treat the prediction of the multi-task feature system as the independent prediction of variable A , since the output is produced without the variable dependencies. On the other hand, we think that the energy-based model gives prediction similar to prediction of variable $A|B$, where B acts as variable associations. We calculate the reduction in error by counting the absolute number of errors for each slot and the joint slot set of both our systems. The test result is reported in Table 6.

Dataset	Slot				Joint
	food	price	area	type	
DSTC2	0.27	0.13	0.04	-	0.18
DSTC3	0.27	0.10	0.16	0.05	0.19

Table 6: Proportional reduction in error of the energy-based system for each slot and the joint goals.

The analysis shows that for more challenging slots such as *food*, the energy-based model reduces the error rate significantly. In both DSTC 2 & 3 a quarter of errors for *food* are corrected, subsequently the errors in joint goals are reduced by nearly 20%.

6 Conclusion

In this paper our contributions were two-fold. We demonstrated, through a number of statistical tests performed on dialogue data and an empirical analysis on variable associations presented in dialogue domains, that dependencies between variables exist and taking them into account improves system performance. We also demonstrated how variable dependencies can be addressed in dialogue state tracking through a structured prediction methodology, and verified our model with respect to the second and third DSTC datasets. While our results do not directly improve on the state of the art, we showed a significant improvement over a non-trivial baseline. We therefore argue that the methodology is promising, and if applied to what is already a state-of-the-art methodology, may help to improve existing systems beyond the state-of-the-art.

There are a number of elements of this work that we are looking to improve. At a fine level we are looking at refinements of the energy-based deep learning architecture and are considering in particular variations on our selected oracle and objective functions that would be better aligned with the multi-categorical nature of our target variables. Meanwhile, at a higher level we want to generalise and further substantiate our investigation by applying the energy-based tracking methodology to tracking architectures that already show state-of-the-art or very near state-of-the-art performance. Finally, we note that a key benefit of this structured methodology is that it allows a more holistic tracking process for the user to be considered where tracking aspects of personality and preference can be neatly integrated alongside the tracking of fine-grained dialogue state. Our longer term goal is thus to apply the structured learning approach in the context of user intent and preference tracking.

Acknowledgments

This research was supported by the ADAPT Centre for Digital Content Technology which is funded under the SFI Research Centres Programme (Grant 13/RC/2106) and is co-funded under the European Regional Development Fund.

References

- David Belanger and Andrew McCallum. 2016. [Structured Prediction Energy Networks](#). In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48.
- David DeVault. 2008. *Contribution tracking: Participating in task-oriented dialogue under uncertainty*. Phd dissertation, State University of New Jersey.
- Layla El Asri, Hannes Schulz, Shikhar Sharma, Jeremie Zumer, Justin Harris, Emery Fine, Rahul Mehrotra, and Kaheer Suleman. 2017. [Frames: A Corpus for Adding Memory to Goal-Oriented Dialogue Systems](#). In *Proceedings of the SIGDIAL 2017 Conference*, pages 207–219.
- James R. Glass. 1999. [Challenges For Spoken Dialogue Systems](#). Technical report, Massachusetts Institute of Technology.
- Michael Gygli, Mohammad Norouzi, and Anelia Angelova. 2017. [Deep Value Networks Learn to Evaluate and Iteratively Refine Structured Outputs](#). In *Proceedings of the 34th International Conference on Machine Learning*.
- Matthew Henderson, Blaise Thomson, and Jason D. Williams. 2014a. The Second Dialog State Tracking Challenge. In *Proceedings of the SIGDIAL 2014 Conference*, pages 263–272.
- Matthew Henderson, Blaise Thomson, and Jason D. Williams. 2014b. The Third Dialog State Tracking Challenge. In *Proceedings of 2014 IEEE Workshop on Spoken Language Technology*, pages 324–329.
- Matthew Henderson, Blaise Thomson, and Steve Young. 2014c. Robust dialog state tracking using delexicalised recurrent neural networks and unsupervised adaptation. In *Proceedings of 2014 IEEE Workshop on Spoken Language Technology*, pages 360–365.
- Matthew Henderson, Blaise Thomson, and Steve Young. 2014d. Word-Based Dialog State Tracking with Recurrent Neural Networks. In *Proceedings of the SIGDIAL 2014 Conference*, pages 292–299.
- Sepp Hochreiter and Jurgen Schmidhuber. 1997. [Long Short-Term Memory](#). *Neural Computation*, 9(8):1735–1780.
- Takaaki Hori, Hai Wang, Chiori Hori, Shinji Watanabe, Bret Harsham, Jonathan Le Roux, John R. Hershey, Yusuke Koji, Yi Jing, Zhaocheng Zhu, and Takeyuki Aikawa. 2016. [Dialog State Tracking With Attention-Based Sequence-To-Sequence Learning](#). In *Proceedings of 2016 IEEE Workshop on Spoken Language Technology*, pages 552–558.
- Youngsoo Jang, Jiyeon Ham, Byung-Jun Lee, Youngjae Chang, and Kee-eung Kim. 2016. [Neural Dialog State Tracker for Large Ontologies by Attention Mechanism](#). In *Proceedings of 2016 IEEE Workshop on Spoken Language Technology*, pages 531–537.
- Frédéric Landragin. 2013. *Man-Machine Dialogue: Design and Challenges*. ISTE Ltd and John Wiley & Sons, Inc.
- Yann LeCun, Sumit Chopra, Raia Hadsell, Marc’ Aurelio Ranzato, and Fu Jie Huang. 2006. A Tutorial on Energy-Based Learning. *Predicting Structured Data*.
- Bing Liu, Gokhan Tur, Dilek Hakkani-Tur, Pararth Shah, and Larry Heck. 2018. [Dialogue Learning with Human Teaching and Feedback in End-to-End Trainable Task-Oriented Dialogue Systems](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2060–2069.
- Nikola Mrksic, Diarmuid O’Seaghdha, Blaise Thomson, Milica Gasic, Pei-Hao Su, David Vandyke, Tsung-Hsien Wen, and Steve Young. 2015. [Multi-domain Dialog State Tracking using Recurrent Neural Networks](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics*, pages 794–799.
- Nikola Mrksic, Diarmuid O’Seaghdha, Tsung-Hsien Wen, Blaise Thomson, and Steve Young. 2017. [Neural Belief Tracker: Data-Driven Dialogue State Tracking](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*.
- Tim Paek and Eric J. Horvitz. 2000. [Conversation as Action Under Uncertainty](#). In *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence*, pages 455–464.
- Julien Perez and Fei Liu. 2017. [Dialog state tracking, a machine reading approach using Memory Network](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, volume 1, pages 305–314.
- Abhinav Rastogi, Dilek Hakkani-Tur, and Larry Heck. 2017. [Scalable Multi-Domain Dialogue State Tracking](#). In *Proceedings of 2017 IEEE Automatic Speech Recognition and Understanding Workshop, ASRU*, pages 561–568.

- Antoine Raux, Brian Langner, Dan Bohus, Alan W Black, and Maxine Eskenazi. 2005. [Let's Go Public! Taking a Spoken Dialog System to the Real World](#). In *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, pages 885–888.
- Hang Ren, Weiqun Xu, Yan Zhang, and Yonghong Yan. 2013. Dialog State Tracking using Conditional Random Fields. In *Proceedings of the SIGDIAL 2013 Conference*, pages 457–461.
- Liliang Ren, Kaige Xie, Lu Chen, and Kai Yu. 2018. [Towards Universal Dialogue State Tracking](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2780–2786.
- Robert J. Ross and John Bateman. 2009. [Daisie: Information State Dialogues for Situated Systems](#). In *Proceedings of International Conference on Text, Speech and Dialogue, TSD 2009*, pages 379–386.
- Noah A. Smith. 2009. Structured Prediction for Natural Language Processing. In *The 26th International Conference on Machine Learning, ICML. Tutorial*.
- Anh Duong Trinh, Robert J. Ross, and John D. Kelleher. 2017. Incremental Joint Modelling for Dialogue State Tracking. In *Proceedings of the 21st Workshop on the Semantics and Pragmatics of Dialogue*, pages 176–177.
- Anh Duong Trinh, Robert J. Ross, and John D. Kelleher. 2018. A Multi-Task Approach to Incremental Dialogue State Tracking. In *Proceedings of The 22nd workshop on the Semantics and Pragmatics of Dialogue, SEMDIAL*, pages 132–145.
- Miroslav Vodolan, Rudolf Kadlec, and Jan Kleindienst. 2015. [Hybrid Dialog State Tracker](#). In *Proceedings of the Machine Learning for SLU & Interaction NIPS 2015 Workshop*.
- Miroslav Vodolan, Rudolf Kadlec, and Jan Kleindienst. 2017. [Hybrid Dialog State Tracker with ASR Features](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, EACL*, volume 2, pages 205–210.
- Nigel G. Ward and David DeVault. 2015. Ten Challenges in Highly-Interactive Dialog Systems. In *AAAI Spring Symposium on Turn-Taking and Coordination in Human-Machine Interaction*, pages 104–107.
- Jason D. Williams. 2014. Web-style ranking and SLU combination for dialog state tracking. In *Proceedings of the SIGDIAL 2014 Conference*, pages 282–291.
- Jason D. Williams, Antoine Raux, and Matthew Henderson. 2016. [The Dialog State Tracking Challenge Series: A Review](#). *Dialogue & Discourse*, 7(3):4–33.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2018. [Global-Locally Self-Attentive Dialogue State Tracker](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, pages 1458–1467.

Leveraging Non-Conversational Tasks for Low Resource Slot Filling: Does it help?

Samuel Louvan
University of Trento
Fondazione Bruno Kessler
slouvan@fbk.eu

Bernardo Magnini
Fondazione Bruno Kessler
magnini@fbk.eu

Abstract

Slot filling is a core operation for utterance understanding in task-oriented dialogue systems. Slots are typically domain-specific, and adding new domains to a dialogue system involves data and time-intensive processes. A popular technique to address the problem is transfer learning, where it is assumed the availability of a large slot filling dataset for the source domain, to be used to help slot filling on the target domain, with fewer data. In this work, instead, we propose to leverage source tasks based on semantically related non-conversational resources (e.g., semantic sequence tagging datasets), as they are both cheaper to obtain and reusable to several slot filling domains. We show that using auxiliary non-conversational tasks in a multi-task learning setup consistently improves low resource slot filling performance.

1 Introduction

Language understanding in task-oriented dialogue systems involves recognizing information (i.e., *slot filling*) expressed in an utterance to accomplish a particular dialogue task. For example, in a flight booking scenario, the utterance “*show me all Delta flights from Milan to New York*” contains information belonging to slots in the flight domain, namely *airline_name* (*Delta*), *origin* (*Milan*), and *destination* (*New York*). Slots are usually predefined and domain-specific, e.g. in a hotel domain slots can be different, such as *room_type*, *length_of_stay* etc. Although recent neural based models (Goo et al., 2018; Wang et al., 2018; Liu and Lane, 2016) have shown remarkable performance in slot filling, they are still based on large labeled data, which means that training a separate model for each domain involves a resource intensive process. Thus, as more domains are added to the system, methods that can

generalize slot filling to new domains with *limited labeled data* (i.e., low-resource settings) are preferable.

Existing works in low resource slot filling are mostly based on transfer learning (Mou et al., 2016), whose aim is to leverage relatively large resources in a source domain (\mathcal{D}_S) for a source task (\mathcal{T}_S), to help a task (\mathcal{T}_T) in a target domain (\mathcal{D}_T), where less data are available. Depending on how the adaptation is performed, there are two notable approaches: data-driven adaptation (Jaech et al., 2016; Goyal et al., 2018; Kim et al., 2016), and model-driven adaptation (Kim et al., 2017; Jha et al., 2018). Essentially, both approaches produce a model on the target domain performing training on the same task (slot filling, in our case), i.e., assuming ($\mathcal{T}_S = \mathcal{T}_T$), although from different domains, i.e. ($\mathcal{D}_S \neq \mathcal{D}_T$). All of these approaches assume that slot filling datasets for the source domain are available, and little effort has been devoted in finding and exploiting cheaper \mathcal{T}_S , which is crucial in a situation where a slot filling dataset in \mathcal{D}_S is not ready yet (*cold-start*).

Accordingly, we attempt to leverage non-conversational source tasks ($\mathcal{T}_S \neq \mathcal{T}_T$) i.e., tasks that use widely available non-conversational resources, to help slot filling. These resources are cheaper to obtain compared to domain-specific slot filling datasets, and many of them are annotated with rich linguistic knowledge, which is potentially useful for slot filling (Chen et al., 2016). Among these resources, we mention PropBank (Palmer et al., 2005) and FrameNet (Baker et al., 1998), which consist of annotated documents with verb and frame-based semantic roles, respectively; CoNLL 2003 (Tjong Kim Sang and De Meulder, 2003) and OntoNotes (Pradhan et al., 2013), which provide named entity information; and Abstract Meaning Representation (AMR) (Banarescu et al., 2013), which provides a graph-based seman-

Sentence	what	is	the	most	expensive	flight	from	boston	to	dallas
ATIS Slot	O	O	O	B-COST_REL	I-COST_REL	O	O	B-FROM_LOC	O	B-TO_LOC
NER	O	O	O	O	O	O	O	B-GPE	O	B-GPE
SemTag	B-QUE	B-ENS	B-DEF	B-TOP	B-IST	B-CON	B-REL	B-GPE	O	B-GPE

Table 1: An example of slot filling annotation from the ATIS (Airline Travel Information System) dataset and author-annotated NER and SemTag in IOB format (Ramshaw and Marcus, 1995). Some ATIS slots correspond to NER or SemTag labels, such as FROM_LOC and TO_LOC with GPE in NER and SemTag. Some slot tags can also be composed of several SemTag labels such as COST_REL which is composed of TOP (*superlative positive*) and IST (*intersective adjective*).

tic formalism.

In this work, we leverage non-conversational tasks as auxiliary tasks in a multi-task learning (MTL) (Caruana, 1997) setup. Given appropriate auxiliary tasks, MTL has shown to be particularly effective in which labeled data is scarce and has been applied to various NLP tasks such as parsing (Søgaard and Goldberg, 2016), POS tagging (Yang et al., 2016), neural machine translation (Luong et al., 2016), and opinion role labeling (Marasovic and Frank, 2018). While there are potentially many non-conversational tasks that we can use as auxiliary tasks, we focus on those that assign semantic class categories to a word, as they are similar in nature to slot filling. In particular, in this work we choose Named Entity Recognition (NER) and the recently introduced Semantic Tagging (SemTag) (Abzianidze and Bos, 2017), motivated by the following rationales:

- Both NER and SemTag are semantically related to slot filling. As illustrated in Table 1, slot labels may correspond to either NER or SemTag labels. In addition, SemTag complements NER as its labels subsume NER labels, and thus could be useful to address linguistic phenomena (e.g. comparative expression, intersective adjective) relevant for slot filling and that are beyond named entities.
- Both NER and SemTag can be re-used in many slot filling domains. Labels in both tasks are typically more general (coarse-grained) compared to labels in slot filling.
- The resources for both tasks are cheaper to obtain compared to domain-specific slot filling datasets, as there have been several initiatives in constructing large datasets for NER and SemTag, for example OntoNotes (Pradhan et al., 2013) and Parallel Meaning Bank (PMB) (Abzianidze et al., 2017) respectively. This is beneficial in a *cold-start* situation in which no slot filling dataset is already available in \mathcal{D}_S .

Although NER has been already used in slot filling models, most of these approaches (Mesnil et al., 2013, 2015; Zhang and Wang, 2016; Gong et al., 2019; Louvan and Magnini, 2018) use and incorporate ground truth NER labels or output of NER systems as features to train a slot filling model, our work differs in the method of learning and leveraging such features from disjoint datasets through MTL and evaluating the performance in low-resource settings.

Our contributions are: (i) we propose to leverage non-conversational tasks, namely NER and SemTag, to improve low resource slot filling through MTL; to our knowledge this MTL combination has not been explored before. (ii) We show that MTL models with NER and SemTag strongly improve single-task slot filling models on three well known datasets. While we focus on using NER and SemTag, our study has shed light on the potential use of non-conversational tasks in general to help low resource slot filling.

2 Approach

Slot filling is often modeled as a sequence labeling problem. Given a sequence of words $\mathbf{x} = (x_1, x_2, \dots, x_n)$ as input, a model \mathcal{M} predicts the corresponding slot labels $\mathbf{y} = (y_1, y_2, \dots, y_n)$ as output.

2.1 Base Model

State-of-the-art models on sequence labeling are typically built based on bi-directional LSTM (bi-LSTM), on top of which there is a CRF model (Lample et al., 2016; Ma and Hovy, 2016). The bi-LSTM takes \mathbf{x} as input and each word x_i is represented as an embedding $\mathbf{e}_i = [\mathbf{w}_i; \mathbf{c}_i]$ composed of the concatenation of a word embedding \mathbf{w}_i and character embeddings \mathbf{c}_i . The bi-LSTM layer produces the forward output state $\vec{\mathbf{h}}_i$ and the backward output state $\overleftarrow{\mathbf{h}}_i$. The concatenation of the output states, $\mathbf{h}_i = [\vec{\mathbf{h}}_i; \overleftarrow{\mathbf{h}}_i]$, is then fed to a

feed-forward (FF) layer, followed by a CRF as the final output layer that predicts a slot label y_i by taking into account the mixture of context information captured by the last FF layer and the slot prediction y_{i-1} from the previous word.

2.2 Multi-task Learning Models

In the context of MTL, models for \mathcal{T}_S , often referred as **auxiliary tasks**, and for \mathcal{T}_T , referred as the **target task**, are simultaneously trained (Yang et al., 2017). In order to perform adaptation, the MTL model \mathcal{M} is partitioned into task-specific parts ($\mathcal{M}_{\mathcal{T}_S}$ and $\mathcal{M}_{\mathcal{T}_T}$) and task-shared-parts ($\mathcal{M}_{\mathcal{T}_S \cap \mathcal{T}_T}$). We use two notable MTL architectures:

- **MTL-Fully Shared Network (MTL-FSN).** The word and character embeddings, and the bi-LSTM layers, are parts of $\mathcal{M}_{\mathcal{T}_S \cap \mathcal{T}_T}$. The hidden state outputs of the bi-LSTM are passed to each of the CRF output layers in $\mathcal{M}_{\mathcal{T}_S}$ and $\mathcal{M}_{\mathcal{T}_T}$. During training a mini-batch of a particular task, the output layers of other tasks are not updated.
- **Hierarchical-MTL (H-MTL).** Inspired by (Søgaard and Goldberg, 2016; Sanh et al., 2019), we introduce a hierarchy of tasks in \mathcal{M} to create different levels of supervision. Instead of placing the output CRF layers for all tasks after the shared bi-LSTM layer, we add a task-specific bi-LSTM in $\mathcal{M}_{\mathcal{T}_T}$ after the shared bi-LSTM and then attach the output layer. In other words, we supervise \mathcal{T}_S , which have coarse-grained labels in the lower level output layer and \mathcal{T}_T , which has more fine-grained labels in the higher level output layer.

3 Experiments

The main objective of our experiments is to validate the hypothesis that using non-conversational tasks as auxiliary tasks in a MTL setup can help low resource slot filling. In our MTL configuration, the **target task** (\mathcal{T}_T) is slot filling, and the **auxiliary tasks** (\mathcal{T}_S) are set to NER or SemTag or both.

Baselines. We compare the two MTL approaches (see §2.2) with the following baselines:

- **Single-Task Learning (STL).** The base model is directly trained and tested on \mathcal{T}_T , without incorporating any information from \mathcal{T}_S . The base model (see §2.1) is a bi-LSTM-CRF which is the core of many models for slot filling (Goo

Dataset	Task	#train	#dev	#test	#label
ATIS	Slot Filling	4478	500	893	79
MIT Restaurant	Slot Filling	6128	1532	3385	8
MIT Movie	Slot Filling	7820	1955	2443	12
OntoNotes 5.0	NER	34970	5896	2327	18
PMB	SemTag	67965	682	650	73

Table 2: Statistics about the datasets, reporting the number of sentences in train/dev/test set, and the number of labels.

et al., 2018; Wang et al., 2018; Liu and Lane, 2016) and sequence tagging tasks in general.

- **STL + Feature Based (STL + FB).** The same model as STL but incorporating the outputs of the independently trained NER and SemTag models as an additional feature in the input embeddings.

Datasets. The language of all the datasets that we use is English. We evaluate our approach on three slot filling datasets, namely ATIS (Price, 1990), MIT Restaurant, and Movie (Liu et al., 2013). ATIS is a widely used dataset for spoken language understanding which contains utterances requesting flight related information. While MIT Restaurant and Movie contain utterances requesting information related to restaurants and movies. For NER, we use the newswire section of OntoNotes 5.0 (Pradhan et al., 2012), which is compiled from English Wall St. Journal. For SemTag, we use Parallel Meaning Bank (PMB) (Abzianidze et al., 2017) 2.2.0. The PMB dataset is constructed from twelve different sources, including OPUS News Commentary (Tiedemann, 2012), Tatoeba¹, Sherlock-Holmes stories, Recognizing Textual Entailment (Giampiccolo et al., 2007), and the bible (Christodoulopoulos and Steedman, 2015). Following the previous publication related to SemTag (Abzianidze and Bos, 2017), we train the SemTag model using the silver data and test on gold data. For all datasets, we use the provided train/dev/test splits. Table 2 shows the overall statistics of each dataset. To simulate the low resource settings, in all experiments we only use 10% training data on \mathcal{T}_T .

Training. We do not tune the hyperparameters² but follow the suggestions and adapt the implementation of Reimers and Gurevych (2017)³. The MTL models are trained in an alternate fashion

¹<https://tatoeba.org/eng/>

²The hyperparameters are listed in Appendix B

³<https://github.com/UKPLab/emnlp2017-bilstm-cnn-crf>

(Jaech et al., 2016) between \mathcal{T}_T and \mathcal{T}_S . Consequently, as the training data size of \mathcal{T}_S is larger than \mathcal{T}_T , the same \mathcal{T}_T data is reused until the whole \mathcal{T}_S is used in the training. We evaluate the performance by computing the F1-score on the test set using the standard CoNLL-2000 evaluation⁴.

4 Results and Discussion

Model	\mathcal{T}_S	\mathcal{T}_T		
		ATIS	MIT-R	MIT-M
STL	-	87.91 _{0.56}	67.37 _{0.26}	80.71 _{0.63}
STL+FB	-	87.79 _{0.67}	67.27 _{0.64}	80.56 _{0.54}
MTL-FSN	<i>N</i>	89.56 _{0.16}	68.82 _{0.18}	80.77 _{0.13}
	<i>S</i>	89.19 _{0.26}	68.21 _{0.71}	80.57 _{0.32}
	<i>N,S</i>	89.10 _{0.41}	68.21 _{0.43}	79.69 _{0.33}
H-MTL	<i>N</i>	89.17 _{0.33}	69.22 _{1.00}	81.79 _{0.26}
	<i>S</i>	88.96 _{0.41}	69.09 _{0.24}	81.59 _{0.17}
	<i>N,S</i>	88.78 _{0.37}	68.96 _{0.50}	81.15 _{0.25}

Table 3: Average F1-score and standard deviation (numbers in subscript) of the performance on the test sets. For the \mathcal{T}_T training split, only 10% data is used. **Bold** indicates the best score for each \mathcal{T}_T . *N* and *S* in \mathcal{T}_S denote NER and SemTag, respectively.

Overall Performance. Table 3 lists the overall performance of the baselines and of the MTL models. We report the average F-1 score and also the standard deviation, as recommended by Reimers and Gurevych (2018), over three runs from different random seeds. For all \mathcal{T}_T , it is evident that the MTL models with NER or SemTag combinations yield the best results compared to STL. MTL models also outperform the STL + FB baseline, indicating that training the model simultaneously with the auxiliary task is better than incorporating the output of the independently trained auxiliary models as features for the slot filling model. In terms of the effectiveness of the auxiliary tasks, using NER produces the best results compared to the other \mathcal{T}_S combinations. The difference between MTL with NER and MTL with SemTag is marginal. Regarding the MTL models, on average, H-MTL yields better scores compared to MTL-FSN in MIT-R and MIT-M, which suggests that supervising tasks with coarse-grained labels and fine-grained labels on different layers is beneficial.

Slot-wise Performance. One of our motivations for using NER and SemTag is that their labels are

\mathcal{T}_T	Concept	Model	
		STL	MTL
ATIS	LOC	94.74 _{0.37}	95.82 _{0.34}
	ORG	92.52 _{0.89}	93.37 _{0.29}
MIT-R	LOC	75.29 _{0.46}	76.02 _{0.39}
MIT-M	PER	85.04 _{0.24}	84.58 _{0.56}

Table 4: Performance on slots related to person (PER), location (LOC), and organization (ORG) concepts. We use the best MTL from Table 3 for each \mathcal{T}_T .

coarse-grained, and that they can be re-used for several slot filling domains. We are interested to see whether MTL improves the performance of slots related to these coarse-grained concepts. In order to do this, we manually created a mapping⁵ from the slots to some coarse-grained entity concepts used by CoNLL-2003 (Tjong Kim Sang and De Meulder, 2003) including Person, Organization, and Location. For example, in ATIS, the slot `airline_name` is mapped to Organization, the slot `fromloc.city_name` is mapped to Location, etc. We perform the analysis on the dev set by re-running the evaluation based on the mapping. Results in Table 4 show that in ATIS and MIT-R, MTL brings improvements on slots related to Location and Organization. However, MTL does not help in slots related to Person names in MIT-M. Based on our observation on the prediction results, most errors gain come from misclassifying DIRECTOR slots as ACTOR slots.




Figure 1: Gain ($\Delta F1$) obtained using MTL over STL on increasing training data. Positive numbers mean MTL is better, negative numbers mean MTL is worse. We use the best MTL from Table 3 for each \mathcal{T}_T .

Performance Gain on Increasing Data Size. We also carried on an experiment by increasing

⁴<https://www.clips.uantwerpen.be/conll2000/>

⁵We provide the mapping in Appendix A

the amount of training data on \mathcal{T}_T , and evaluated the performance on the dev set to understand the usefulness of MTL on varying data size. As shown in Figure 1, as we increase the size of the training data, the gain that we obtain using MTL tends to decrease. The results suggest that MTL is indeed more useful in very low resource scenarios, according to our initial hypothesis. After 40% training data size is used (around 2K utterances), MTL is less useful. We believe that this is because the slot filling datasets are relatively simple, e.g. the texts are short and most of them express a single specific request, thus, it is relatively easy for the model to capture the regularities.

Impact on Auxiliary Tasks Performance. We also perform an analysis to understand the effect of MTL to the model performance for \mathcal{T}_S . The STL performance of OntoNotes and Semantic Tagging are around 89% and 96% respectively in terms of F1-score. With MTL, on average, the \mathcal{T}_S model performance decrease about 0.7 points for OntoNotes and 0.2 points for Semantic Tagging. This suggests that \mathcal{T}_S models do not benefit from the low resource \mathcal{T}_T through the MTL framework and the training mechanism that we use. In general, whether MTL can benefit model performance in a target task given auxiliary tasks (or vice versa) is still a question and beyond the scope of this paper. While there is no exact answer yet for this question, we refer to (Bingel and Søgaard, 2017; Alonso and Plank, 2017) which study the characteristics of auxiliary tasks that is potential to help target task performance (Bingel and Søgaard, 2017; Alonso and Plank, 2017).

5 Conclusions

We proposed to leverage non-conversational tasks, Named Entity Recognition and Semantic Tagging, through multi-task learning to help low resource slot filling. Our experiments demonstrate that: (i) non-conversational tasks are effective to improve slot filling performance, and they are reusable in different slot filling domains; (ii) incorporating a task-hierarchy in the multi-task architecture based on the granularity of the labels is beneficial for the model performance on two out of three datasets.

In the future, we plan to explore other non-conversational resources such as FrameNet (Baker et al., 1998) which provide a repository of event frames and semantic roles that can be relevant for intent classification and slot filling in task-oriented

dialogue systems. Also another direction is to apply fine-tuning with the recently popular pre-trained language model e.g. BERT (Devlin et al., 2018).

References

- Lasha Abzianidze, Johannes Bjerva, Kilian Evang, Hessel Haagsma, Rik van Noord, Pierre Ludmann, Duc-Duy Nguyen, and Johan Bos. 2017. The Parallel Meaning Bank: Towards a Multilingual Corpus of Translations Annotated with Compositional Meaning Representations. In *EACL*.
- Lasha Abzianidze and Johan Bos. 2017. *Towards Universal Semantic Tagging*. In *IWCS 2017 — 12th International Conference on Computational Semantics — Short papers*.
- Hector Martinez Alonso and Barbara Plank. 2017. When is multitask learning effective? semantic sequence prediction under varying data conditions. In *EACL 2017-15th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1–10.
- Collin F Baker, Charles J Fillmore, and John B Lowe. 1998. The Berkeley Framenet Project. In *Proceedings of the 17th international conference on Computational linguistics-Volume 1*, pages 86–90. Association for Computational Linguistics.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract Meaning Representation for Sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186.
- Joachim Bingel and Anders Søgaard. 2017. Identifying beneficial task relations for multi-task learning in deep neural networks. *EACL 2017*, page 164.
- Rich Caruana. 1997. Multitask Learning. *Machine Learning*, 28:41–75.
- Yun-Nung Chen, Dilek Hakanni-Tur, Gökhan Tür, Asli Çelikyilmaz, Jianfeng Gao, and Li Deng. 2016. *Syntax or Semantics? Knowledge-guided Joint Semantic Frame Parsing*. In *2016 IEEE Spoken Language Technology Workshop, SLT 2016, San Diego, CA, USA, December 13-16, 2016*, pages 348–355.
- Christos Christodoulopoulos and Mark Steedman. 2015. A massively parallel corpus: the bible in 100 languages. In *Language Resources and Evaluation*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *CoRR*, abs/1810.04805.

- Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. 2007. [The third PASCAL recognizing textual entailment challenge](#). In *Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, pages 1–9, Prague. Association for Computational Linguistics.
- Yu Gong, Xusheng Luo, Yu Zhu, Wenwu Ou, Zhao Li, Muhua Zhu, Kenny Q. Zhu, Linfu Duan, and Xi Chen. 2019. Deep cascade multi-task learning for slot filling in online shopping assistant. In *AAAI 2019*.
- Chih-Wen Goo, Guang Gao, Yun-Kai Hsu, Chih-Li Huo, Tsung-Chieh Chen, Keng-Wei Hsu, and Yun-Nung Chen. 2018. [Slot-Gated Modeling for Joint Slot Filling and Intent Prediction](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 753–757. Association for Computational Linguistics.
- Anuj Kumar Goyal, Angeliki Metallinou, and Spyros Matsoukas. 2018. [Fast and Scalable Expansion of Natural Language Understanding Functionality for Intelligent Agents](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers)*, pages 145–152. Association for Computational Linguistics.
- Aaron Jaech, Larry P. Heck, and Mari Ostendorf. 2016. Domain Adaptation of Recurrent Neural Networks for Natural Language Understanding. In *INTER-SPEECH*.
- Rahul Jha, Alex Marin, Suvamsh Shivaprasad, and Imed Zitouni. 2018. [Bag of experts architectures for model reuse in conversational language understanding](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers)*, pages 153–161. Association for Computational Linguistics.
- Young-Bum Kim, Karl Stratos, and Dongchan Kim. 2017. [Domain attention with an ensemble of experts](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 643–653. Association for Computational Linguistics.
- Young-Bum Kim, Karl Stratos, and Ruhi Sarikaya. 2016. [Domainless Adaptation by Constrained Decoding on a Schema Lattice](#). In *COLING 2016, 26th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers, December 11-16, 2016, Osaka, Japan*, pages 2051–2060.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. [Neural architectures for named entity recognition](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270. Association for Computational Linguistics.
- Bing Liu and Ian Lane. 2016. [Joint Online Spoken Language Understanding and Language Modeling With Recurrent Neural Networks](#). In *Proceedings of the SIGDIAL 2016 Conference, The 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue, 13-15 September 2016, Los Angeles, CA, USA*, pages 22–30.
- Jingjing Liu, Panupong Pasupat, Scott Cyphers, and Jim Glass. 2013. Asgard: A Portable Architecture for Multilingual Dialogue Systems. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 8386–8390. IEEE.
- Samuel Louvan and Bernardo Magnini. 2018. From General to Specific : Leveraging Named Entity Recognition for Slot Filling in Conversational Language Understanding. In *CLiC-it*.
- Minh-Thang Luong, Quoc V. Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. 2016. Multi-task sequence to sequence learning. *ICLR*, abs/1511.06114.
- Xuezhe Ma and Eduard Hovy. 2016. End-to-end Sequence Labeling via Bi-directional LSTM-CNNs-CRF. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1064–1074.
- Ana Marasovic and Anette Frank. 2018. SRL4ORL: Improving Opinion Role Labelling using Multi-task Learning with Semantic Role Labeling. In *NAACL-HLT*.
- Grégoire Mesnil, Yann Dauphin, Kaisheng Yao, Yoshua Bengio, Li Deng, Dilek Hakkani-Tur, Xiaodong He, Larry Heck, Gokhan Tur, Dong Yu, and Geoffrey Zweig. 2015. [Using recurrent neural networks for slot filling in spoken language understanding](#). *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, 23(3):530–539.
- Grégoire Mesnil, Xiaodong He, Li Deng, and Yoshua Bengio. 2013. Investigation of recurrent-neural-network architectures and learning methods for spoken language understanding. In *INTER-SPEECH*.
- Lili Mou, Zhao Meng, Rui Yan, Ge Li, Yan Xu, Lu Zhang, and Zhi Jin. 2016. [How Transferable are Neural Networks in NLP Applications?](#) In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 479–489. Association for Computational Linguistics.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The Proposition Bank: An Annotated Corpus of Semantic Roles. *Computational linguistics*, 31(1):71–106.

Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Hwee Tou Ng, Anders Björkelund, Olga Uryupina, Yuchen Zhang, and Zhi Zhong. 2013. Towards Robust Linguistic Analysis using OntoNotes. In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 143–152.

Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. CoNLL-2012 Shared Task: Modeling Multilingual Unrestricted Coreference in OntoNotes. In *EMNLP-CoNLL Shared Task*.

Patti J Price. 1990. Evaluation of Spoken Language Systems: The ATIS Domain. In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*.

Lance Ramshaw and Mitch Marcus. 1995. [Text Chunking using Transformation-Based Learning](#). In *Third Workshop on Very Large Corpora*.

Nils Reimers and Iryna Gurevych. 2017. [Reporting Score Distributions Makes a Difference: Performance Study of LSTM-networks for Sequence Tagging](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 338–348, Copenhagen, Denmark.

Nils Reimers and Iryna Gurevych. 2018. Why Comparing Single Performance Scores Does Not Allow to Draw Conclusions About Machine Learning Approaches. *CoRR*, abs/1803.09578.

Victor Sanh, Thomas Wolf, and Sebastian Ruder. 2019. A Hierarchical Multi-task Approach for Learning Embeddings from Semantic Tasks. *AAAI*.

Anders Søgaard and Yoav Goldberg. 2016. Deep Multi-task Learning with Low Level Tasks Supervised at Lower Layers. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 231–235.

Jörg Tiedemann. 2012. Parallel data, tools and interfaces in opus. In *LREC*.

Erik F Tjong Kim Sang and Fien De Meulder. 2003. Introduction to the CoNLL-2003 Shared Task: Language-independent Named Entity Recognition. In *Proceedings of the seventh conference on Natural language learning at HLT-NAACL 2003-Volume 4*, pages 142–147. Association for Computational Linguistics.

Yu Wang, Yilin Shen, and Hongxia Jin. 2018. [A bi-model based rnn semantic frame parsing model for intent detection and slot filling](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 309–314. Association for Computational Linguistics.

Zhilin Yang, Ruslan Salakhutdinov, and William W. Cohen. 2016. Transfer Learning for Sequence Tagging with Hierarchical Recurrent Networks. *CoRR*, abs/1703.06345.

Zhilin Yang, Ruslan Salakhutdinov, and William W. Cohen. 2017. Transfer Learning for Sequence Tagging with Hierarchical Recurrent Networks. In *ICLR*.

Xiaodong Zhang and Houfeng Wang. 2016. A joint model of intent determination and slot filling for spoken language understanding. In *IJCAI*.

A Mapping of entity concepts and slots for each dataset

Concept	ATIS	MIT-R	MIT-M
LOC	fromloc.airport_code fromloc.airport_name fromloc.city_name fromloc.state_code fromloc.state_name stoploc.airport_name stoploc.city_name stoploc.state_code toloc.airport_code toloc.airport_name toloc.city_name toloc.country_name toloc.state_code toloc.state_name	location	-
ORG	airline_name	-	-
PER	-	-	character actor director

Table 5: The mapping of entity concepts, namely Location (LOC), Organization (ORG), and Person (PER) to their corresponding slots in each dataset.

B Hyperparameters

Hyperparameter	Value
LSTM cell size	100
Dropout	0.5
Word embedding dimension	300
Character embedding dimension	100
Mini-batch size	32
Optimizer	Adam
Number of epoch	50
Early stopping	10

Collaborative Multi-Agent Dialogue Model Training Via Reinforcement Learning

Alexandros Papangelis, Yi-Chia Wang, Piero Molino, Gokhan Tur

Uber AI

San Francisco, California

{apapangelis, yichia.wang, piero, gokhan}@uber.com

Abstract

We present the first complete attempt at concurrently training conversational agents that communicate only via self-generated language. Using DSTC2 as seed data, we trained natural language understanding (NLU) and generation (NLG) networks for each agent and let the agents interact online. We model the interaction as a stochastic collaborative game where each agent (player) has a role (“assistant”, “tourist”, “eater”, etc.) and their own objectives, and can only interact via natural language they generate. Each agent, therefore, needs to learn to operate optimally in an environment with multiple sources of uncertainty (its own NLU and NLG, the other agent’s NLU, Policy, and NLG). In our evaluation, we show that the stochastic-game agents outperform deep learning based supervised baselines.

1 Introduction

Machine learning for conversational agents has seen great advances (e.g. Tur and Mori, 2011; Gao et al., 2019; Singh et al., 1999; Young et al., 2013; Oh and Rudnicky, 2000; Zen et al., 2009; Reiter and Dale, 2000; Rieser and Lemon, 2010), especially when adopting deep learning models (Deng and Liu, 2018; Mesnil et al., 2015; Wen et al., 2015, 2017; Su et al., 2017; Papangelis et al., 2018; Liu and Lane, 2018b; Li et al., 2017; Williams et al., 2017; Liu and Lane, 2018a). Most of these works, however, suffer from the lack of data availability as it is very challenging to design sample-efficient learning algorithms for problems as complex as training agents capable of meaningful conversations. Among other simplifications, this results in treating the interaction as a single-agent learning problem, i.e. assuming that from the conversational agent’s perspective the world may be complex but is stationary. In this work,

we model conversational interaction as a stochastic game (e.g. Bowling and Veloso, 2000) and train two conversational agents, each with a different role, which learn by interacting with each other via natural language. We first train Language Understanding (NLU) and Generation (NLG) neural networks for each agent and then use multi-agent reinforcement learning, namely the Win or Lose Fast Policy Hill Climbing (WoLF-PHC) algorithm (Bowling and Veloso, 2001), to learn optimal dialogue policies in the presence of high levels of uncertainty that originate from each agent’s statistical NLU and NLG, and the other agent’s erratic behaviour (as the other agent is learning at the same time). While not completely alleviating the need for seed data needed to train the NLU and NLG components, the multi-agent setup has the effect of augmenting them, allowing us to generate dialogues and behaviours not present in the original data.

Employing a user simulator is an established method for dialogue policy learning (Schatzmann et al., 2007, among others) and end-to-end dialogue training (Asri et al., 2016; Liu and Lane, 2018b). Training two conversational agents concurrently has been proposed by Georgila et al. (2014); training them via natural language communication was partially realized by Liu and Lane (2017), as they train agents that receive text input but generate dialogue acts. However, to the best of our knowledge, this is the first study that allows fully-trained agents to communicate only in natural language, and does not allow any all-seeing critic / discriminator. Inspired by Hakkani-Tür (2018), each agent learns in a decentralized setting, only observing the other agent’s language output and a reward signal. This allows new, untrained agents to directly interact with trained agents and learn without the need for adjusting parameters that can affect the already trained agents.

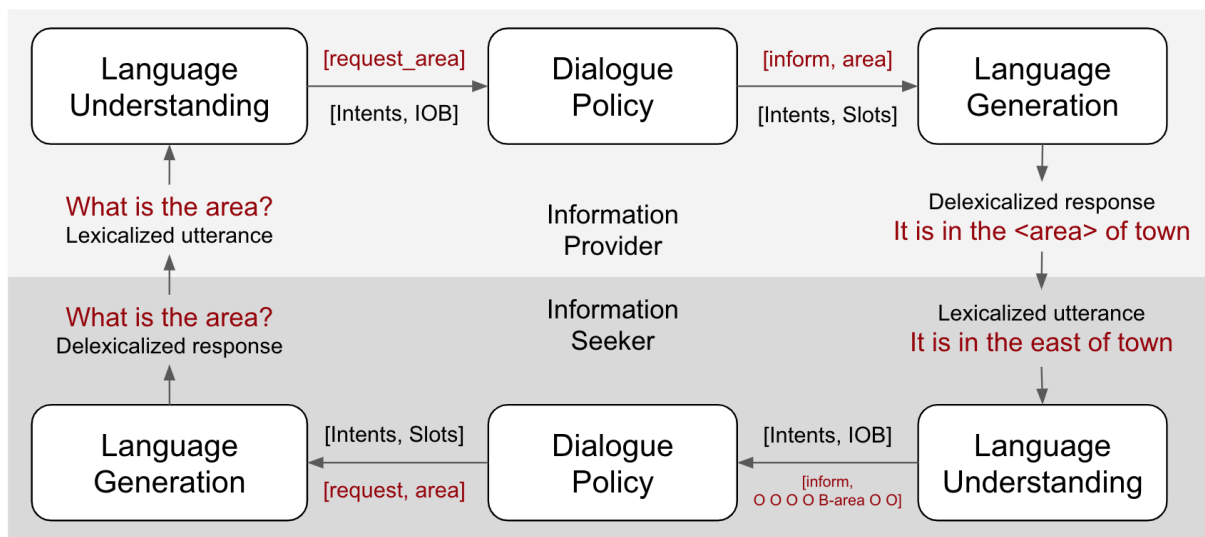


Figure 1: Information flow between two agents on a successful example (shown in red, starting from the Information Seeker’s policy). Where needed, slot values are populated from the tracked dialogue state.

The architecture of each agent is mirrored as shown in Figure 1, so the effort of adding agents with new roles is minimal. As seed data, we use data from DSTC2 (Henderson et al., 2014), which concerns dialogues between humans asking for restaurant information and a machine providing such information. Our contributions are: 1) we propose a method for training fully text-to-text conversational agents from mutually generated data; and 2) we show how agents trained by multi-agent reinforcement learning and minimal seed human-machine data can produce high quality dialogues as compared to single-agent policy models in an empirical evaluation.

1.1 Related Work

Collecting and annotating a big corpus requires significant effort and has the additional challenge that agents trained in a supervised manner with a given corpus cannot easily generalize to unseen / out of domain input. Building a good user simulator to train against can be challenging as well, even equivalent to building a dialogue system in some cases. Directly learning from humans leads to policies of higher quality, but requires thousands of dialogues even for small domains (Gasic et al., 2013). Shah et al. (2018) combine such resources to train dialogue policies. Recently, model-based RL approaches to dialogue policy learning are being revisited (Wu et al., 2018); however, such methods still assume a stationary environment.

Georgila et al. (2014) concurrently learn two negotiator agents’ dialogue policies in a set-

ting where they negotiate allocation of resources. However, their agents do not interact via language, but rather via dialogue acts. They use PHC and WoLF-PHC (Bowling and Veloso, 2001) to train their agents, who use two types of dialogue acts: *accept* and *offer*, each of which takes two numerical arguments. Lewis et al. (2017) train agents on a similar task, but their agents are modelled as end-to-end networks that learn directly from text. However, the authors train their negotiator agent on supervised data and against a fixed supervised agent. Earlier works include English and Heeman (2005), the first to train policies for two conversational agents, but with single-agent RL, and Chandramohan et al. (2014) who applied co-adaptation on single-agent RL, using Inverse RL to infer reward functions from data.

Liu and Lane (2017) train two agents on DSTC2 data, taking text as input and producing dialogue acts that are then fed to template-based language generators. They pre-train their models using the data in a supervised manner and apply reinforcement learning on top. In our setup, information providers and seekers are modeled as active players in a non-stationary environment who interact with each other via language they generate, using statistical language generators. Each agent has their own reward as the objectives are not identical, and their dialogue manager uses a method designed for non-stationary environments. While our setup still needs seed data to ensure linguistic consistency and variability, it augments this data and can train high quality conversational agents.

Goal	Constr(pricerange=cheap), Constr(area=north), Req(addr), Req(phone)
Agent Role	Input / Output
	<i>Example of DM error (Seeker’s policy is also learning):</i>
Prov. NLG Seeker NLU Seeker DM	what part of town do you have in mind? request(area) act_inform food
	<i>Example of NLG error:</i>
Seeker DM Seeker NLG Prov. NLU Prov. DM Prov. NLG Seeker NLU	act_request phone what is the phone request(phone) act_inform phone the post code is c.b 4, 1 u.y . inform(postcode = c.b 4, 1 u.y)
	<i>Example of NLU error:</i>
Provider NLG Seeker NLU	the phone number is 01223 356555 inform(phone= 01223)

Table 1: A failed dialogue between two conversational agents during training. Uncertainty originating from NLU and NLG components on top of the erratic behaviour of each agent’s policy (as they learn concurrently) can have a big impact on the quality of the learned dialogue policies.

Other than the works mentioned above, many approaches have been proposed to train modular or end-to-end dialogue systems. To the best of our knowledge, however, none of them concurrently trains two conversational agents.

2 System Overview

Figure 1 shows the general architecture and information flow of our system, composed of two agents who communicate via written language. Our system operates in the well-known DSTC2 domain (Henderson et al., 2014) which concerns information about restaurants in Cambridge; however, our multi-agent system supports any slot-filling / information-seeking domain. The Language Understanding and Generation components are trained offline as described in the following sections, while the dialogue policies of the agents are trained online during their interaction. Given that our language generation component is model-based rather than retrieval-based or template-based, we believe that the quality of the generated language and dialogues is encouraging (see appendix for some example dialogues).

2.1 Language Understanding

The task of Natural Language Understanding (NLU) consists of mapping a free-form sentence to a meaning representation, usually in the form of a semantic frame. The frame consists of an intent and a set of slots with associated val-

ues. For instance, the semantic frame of the sentence “Book me an Italian restaurant in the south part of the city” can be mapped to the frame “book_restaurant (food: Italian, area: south)” where *book_restaurant* is the intent and *food* and *area* are the slots.

In recent years, deep learning approaches have been adopted for NLU, performing intent classification and slot tagging both independently (Tür et al., 2012; Lee and Deroncourt, 2016; Xu and Sarikaya, 2013; Mesnil et al., 2015; Kurata et al., 2016; Huang et al., 2015) and jointly (Zhang and Wang, 2016; Rojas-Barahona et al., 2016). In Hakkani-Tür et al. (2016), decoders tag each word in the input sentence with a different slot name and concatenate the intent as a tag to the end-of-sentence token, while in Liu and Lane (2016) the encoder is shared, but the two tasks have separate decoders. In most cases, intent detection is treated as a classification problem and the slot name tags for all words are uniquely assigned to the intent detected in the sentence.

In our case, as we decided to use the same NLU model architecture for both agent roles, we could not rely on multi-class classification. In particular, system outputs in DSTC2 often contain multiple acts, so an “information seeker” NLU model has to learn to identify which intents are present in the system utterance as well as to assign slot values to each identified intent. An example of this need is evident in the sentence “There are

no Italian restaurants in the south part of the city, but one is available in the west side” which can be mapped to “ $\{deny(food: Italian, area: south), inform(area: west)\}$ ”. In order to tackle those scenarios, we designed our decoder to predict multiple intents (casting the task as a multi-label classification problem) where each intent is a class and, for the “request” intent, the pair of “request” and all requestable slots are additional classes. This is necessary as the slot values of the request intent are names of slots (e.g. *request(food)*), and they may not be mentioned explicitly in the sentences. Moreover, to account for the multiple intents in the set tagger decoder, we augmented the number of possible tags for each word in the sentence concatenating the name of the intent they are associated with. In the previous example, for instance, the word “south” is assigned a “*deny_area*” tag, while the word “west” is assigned an “*inform_area*” tag, so the name of the intent in the tag identifies which of the multiple intents each slot is assigned to. This increases the number of tags, but allows an unequivocal assignment of the slot values to the intents they belong to.

The whole model, which is composed of a convolutional encoder and the two decoders (one intent multi-label classifier and a slot tagger), is trained end-to-end in a multi-task fashion, with both multi-label intent classification and slot tagging tasks being optimized at the same time. The output set of semantic frames from the NLU is then aggregated over time and passed on to the dialogue policy.

Evaluating NLU Quality Table 2 summarizes the performance (F1 scores) of the trained models, with respect to intent, frame, and slot IOB tags, calculated on the DSTC2 test set. The F1 measure is used instead of accuracy due to the multiple intents, acts and slots in our problem formulation.

Role	Intent F1	Slots F1	Frame F1
Provider	0.929	0.899	0.927
Seeker	0.986	0.995	0.983

Table 2: F1 scores for each agent’s NLU model.

2.2 Dialogue Policy Learning

As already discussed, in this work we train two agents: one seeking restaurant information (“seeker”) and one providing information (“provider”). Each agent’s dialogue policy re-

ceives the tracked dialogue state and outputs a dialogue act. While both agents have the same set of dialogue acts to choose from, they have different arguments to use for these acts (Henderson et al., 2014). Each agent also has a different dialogue state, representing its perception of the world. The seeker’s state models its preferences (goal) and what information the provider has given, while the provider’s state models constraints expressed or information requested by the seeker, as well as attributes of the current item in focus (retrieved from a database) and metrics related to current database results, such as number of items retrieved, slot value entropies, etc. The reward signal is slightly different for each agent, even though the task is collaborative. It assigns a positive value on successful task completion (restaurant provided matches the seeker’s goal, and all seeker’s requests are answered), a negative value otherwise, and a small negative value for each dialogue turn to favor shorter interactions. However, a seeker is penalised for each request in the goal that is not expressed, and a provider is penalised for each request that is unanswered. To train good dialogue policies in this noisy multi-agent environment, we opted for WoLF-PHC as a proof of concept and leave investigation of general-sum and other methods that scale better on richer domains for future work. The dialogue policies that we train operate on the full DSTC2 act and a subset of the slot space. Specifically, not all dialogue acts have slot arguments and we do not allow multiple arguments per act or multiple acts per turn, so the size of our action space is 23. In the input, all policies receive the output of the NLU aggregated over the past dialogue turns (i.e. keeping track of slots mentioned in the past) with - as mentioned above - the state of the seeker including its own goal, and the state of the provider including current database result metrics which are fetched through SQL queries formed using the slot-value pairs in the provider’s state.

2.2.1 WoLF-PHC

A *stochastic game* can be thought of as a *Markov Decision Process* extended to multiple agents. It is defined as a tuple $(n, S, A_{1..n}, T, R_{1..n})$, where n is the number of agents, S is the set of states, A_i is the set of actions available to agent i , $T : S \times A \times S \rightarrow [0, 1]$ is the transition function, and $R_i : S \times A \rightarrow \mathfrak{R}$ is the reward function of agent i .

WoLF-PHC (Bowling and Veloso, 2001) is a PHC algorithm (simple extension to Q-Learning for mixed policies) with variable learning rate and the principle according to which the agent should learn quickly (i.e. with a higher learning rate) when losing and slowly when winning. Briefly, Q is updated as in Q-Learning and an estimate of the average policy is maintained:

$\tilde{\pi}(s, a') \leftarrow \tilde{\pi}(s, a') + \frac{1}{C(s)}(\pi(s, a') - \tilde{\pi}(s, a'))$, where $C(s)$ is the number of times state s has been visited. The policy then is updated as follows:

$$\pi(s, a) \leftarrow \pi(s, a) + \begin{cases} \delta & a = \mathop{\text{amax}}_{a'} Q(s, a') \\ \frac{-\delta}{|A_i|-1} & \text{otherwise} \end{cases}$$

$$\delta = \begin{cases} \delta_w & \sum_a \pi(s, a)Q(s, a) > \sum_a \tilde{\pi}(s, a)Q(s, a) \\ \delta_l & \text{otherwise} \end{cases}$$

where δ_w and δ_l are learning rates.

2.3 Language Generation

Natural language generation (NLG) is a critical module in dialogue systems. It operates in the later phase of the dialogue system, consumes the meaning representation of the intended output provided by the dialogue manager, and converts it to a natural language utterance.

Previous research has approached the NLG problem in various ways (e.g., Langkilde and Knight, 1998; Walker et al., 2002; Oh and Rudnicky, 2000). One common approach is rule-based / template-based generation, which produces utterances from handcrafted rules or templates where slot variables are filled with values from the meaning representation provided by the dialogue manager. This approach has been widely adopted in both industrial and research systems. Although it guarantees high-quality output, it is time-consuming to write templates especially for all possible meaning representations and the generated sentences quickly become repetitive for the users. Moreover, scalability and maintenance of these templates become concerns as we expand the system to deal with more domains or scenarios.

More recently, deep neural networks have been widely adopted in natural language generation because of their effectiveness. Among all types of deep learning architectures, the sequence-to-sequence approach (*seq2seq*) has been most

widely and successfully adopted for language generation in several tasks as machine translation (e.g. Sutskever et al., 2014), question answering (e.g. Yin et al., 2016), text summarization (e.g. Chopra et al., 2016), and conversational models (e.g. Shang et al., 2015; Serban et al., 2016).

Our NLG model is inspired by recent state of the art *seq2seq* models such as Sutskever et al. (2014) and Wen et al. (2015), that transform one sequence of words to another. Our *seq2seq* model was constructed to take a meaning representation string as input and generate the corresponding natural language template as output. Both input and output were delexicalized with slot values replaced by tags, and values are filled in after the template is generated. An example of input and output of the system NLG is shown below:

Input: act_inform <food> act_inform
<pricerange> act_offer <name>
Output: <name> is a great restaurant
serving <food> food and it is in the
<pricerange> price range

Specifically, we implemented our Encoder-Decoder model with Long Short-Term Memory (LSTM) recurrent networks. We employed an attention mechanism (Bahdanau et al., 2015) to emphasize relevant parts of the input sequence at each step when generating the output sequence. We further improved the model by encoding the conversation history as a context vector and concatenating it with the encoded input for output generation. We observed that context not only increases the model performance, but also helps to produce output with more *variation*, which has been considered one of the important factors of a good NLG model (Stent et al., 2005). Both agents' NLG models were built in the same way using the provider- or seeker-side data.

Evaluating NLG Quality BLEU score (Papineni et al., 2002) has been one of the most commonly used metrics for NLG evaluation. Since it is agreed that the existing automatic evaluation metrics for NLG have limitations (Belz and Reiter, 2006), we introduced a modified version of BLEU which attempts to compensate the gap of the current BLEU metric. BLEU, ranging from 0 to 1, is a precision metric that quantifies n-gram overlaps between a generated text and the ground truth text. However, we observed that in the DSTC2 data a meaning representation can map to different templates as the example shown below:

MR: act_inform <pricerange> act_offer
 <name>
T1: the price range at <name> is
 <pricerange>
T2: <name> is in the <pricerange> price
 range

Thus, to compute BLEU of a model-generated template, instead of only comparing it against its corresponding ground truth template, we calculated its BLEU scores with all the possible templates that have the same input meaning representation in the DSTC2 data, and the maximum BLEU score among them is the final BLEU of this generated template. By doing so, the average BLEU scores of the information provider and seeker NLG models on the test set are 0.8625 and 0.5293, respectively. Note that it is not surprising that the seeker model does not perform as well as the provider model because the seeker-side data has many more unique meaning representations and natural language templates, which make the task of building a good seeker model harder.

3 Evaluation

The Plato Research Dialogue System¹ was used to implement, train, and evaluate the agents. To assess the quality of the dialogues our agents are capable of, we compare dialogue success rates, average cumulative rewards, and average dialogue turns along two dimensions: a) access to ground truth labels during training or not; b) stationary or non-stationary environment during training. We therefore train four kinds of conversational agents for each role (eight in total) as shown in Table 3. Due to the nature of our setup, algorithms designed for stationary environments (e.g. DQN) are not considered.

	Stat. Env.	Non-Stat. Env.
Dial. Acts	SuperDAct	WoLF-Dact
Text	Supervised	WoLF-PHC

Table 3: The four conditions under which our conversational agents are trained.

Specifically, the *SuperDAct* agents are modelled as 3-layer Feed Forward Networks (FFN), trained on DSTC2 data using the provided dialogue act annotations. The *Supervised* agents (also 3-layer FFN) are trained on DSTC2 data but

¹The source code for the full dialogue system can be found here <https://github.com/uber-research/plato-research-dialogue-system>

each agent’s policy uses the output of its respective NLU: the provider (dialogue system in the dataset) generates its utterance using its trained NLG with the dialogue acts found in the data as input; the seeker (human caller in the dataset) then uses the provider’s utterance as input to its NLU whose output is then fed to its policy; and the same approach is used for the provider’s side. The *WoLF-DAct* agents are trained concurrently (i.e. in a non-stationary environment) but interacting via dialogue acts, while the *WoLF-PHC* agents are trained concurrently and interacting via generated language, as show in in Figure 1. All of these agents are then evaluated on the full language to language setup². Apart from the above, we trained conversational agents using deep policy gradient algorithms. Their performance could not match the WoLF-PHC or the supervised agents, however, even after alternating the policy gradient agents’ training to account for non-stationarity. This is not unexpected, of course, since those algorithms are designed to learn in a stationary environment. These results therefore are not reported here.

In our evaluation, a dialogue is considered successful if the information seeker’s goal is met by the provider, following the standard definition used for this domain (Su et al., 2017, e.g.). Under this definition, a provider must offer an item that matches the seeker’s constraints and must answer all requests made by the seeker. However, as seen in Table 6, even when the dialogue manager’s output is correct, it can be realized by NLG or understood by NLU erroneously. While none of the models (NLU, DM, NLG) directly optimises this objective, it is a good proxy of overall system performance and allows for direct comparison with prior work. As a reward signal for reinforcement learning we use the standard reward function found in the literature (Gasic et al., 2013; Su et al., 2017, e.g.), tweaked to fit each agent’s perception as described in section 2.2.

Figure 2 shows learning curves with respect to the metrics we use for all conversational agents, where each kind of agent was evaluated against its counterpart (e.g. *Supervised* seeker against *Supervised* provider) on the environment they were trained on. Table 4 shows the main results of

²The *SuperDAct* and *WoLF-DAct* agents achieve 81% and 95% dialogue success rates respectively when evaluated on a dialogue act to dialogue act setup (i.e. without LU/LG) against an agenda-based simulated Seeker. When evaluated against each other (Fig. 2) the performance naturally drops.

Average Dialogue Success			
SuperDAct	Supervised	WoLF-DAct	WoLF-PHC
44.23%	46.30%	52.56%	66.30%
Average Cumulative Rewards			
SuperDAct	Supervised	WoLF-DAct	WoLF-PHC
4.42	6.68	7.84	10.93
Average Dialogue Turns			
SuperDAct	Supervised	WoLF-DAct	WoLF-PHC
10.89	8.65	9.81	9.57

Table 4: Average dialogue success, reward, and number of turns on the agents evaluated, over 3 training/evaluation cycles with goals sampled from the test set of DSTC2. Regardless of training condition, all agents were evaluated in the language to language setting. All differences between SuperDAct - WoLF-DAct, and Supervised - WoLF-PHC are significant with $p < 0.02$.

our evaluation in the language to language setting, where each cell represents the average of 3 train/evaluation cycles of policies trained under the respective conditions for 20,000 dialogues (200 epochs for the supervised agents) and evaluated for 1,000 dialogues. We can see that the WoLF-PHC agents outperform the other conditions in almost every metric, most likely because they model the conversation as a stochastic game and not as a single-agent problem. Comparing Figure 2 with Table 4 we can see that the agents trained on dialogue acts cannot generalise to the language to language setting, even when paired with NLU and NLG models that show strong performance (see previous section). On a similar setup (joint NLU and DM but without statistical NLG), Liu and Lane (2017) report 35.3% dialogue success rate for their supervised baseline and 64.7% for reinforcement learning on top of pre-trained supervised agents.

We attribute the low performance of the supervised policies to a lack of data and context in the DSTC2 dataset. We believe that in the presence of errors from our statistical NLU and NLG, there just are not enough dialogues or information within each dialogue for the supervised policies to learn to associate states with optimal actions. In particular, if one of the NLGs or NLUs (for either agent) makes a mistake, this affects the dialogue state tracking and subsequently the database retrieval, resulting in a state that may not actually be in the dataset. In the presence of this uncertainty we found that seeker and provider do not properly learn how to make requests and address them, respectively and this is the most frequent reason for dialogue task failure in this condition. This is

partly due to the fact that in DSTC2 the provider’s side responds to requests with an offer and an inform, for example a response to a request for phone number would be: offer(name=kymmoy), inform(phone=01223 311911) which may be confusing both models. In light of this, we trained a supervised policy model able to output multiple actions at each dialogue turn. However, this makes the learning problem even harder and we found that in this case such models perform poorly. Overall the two supervised approaches appear to perform similarly on objective dialogue task success but the *Supervised* agents who have seen uncertainty during the training seem to perform better in terms of rewards achieved and number of dialogue turns.

Upon pairing different combinations of the eight agents we trained, we observe that agents who are able to better model the seeker’s behaviour perform best in the joint task. In our case, WoLF-trained agents are able to better model the seeker’s behaviour, which partially explains the higher success rates. However, we note that the *WoLF-DAct* agents do not generalise very well to the much harder language to language environment. Another general trend that we observe is that the WoLF-trained agents seem to take longer number of turns but lead to higher rewards and success rates likely because they persist for more turns before giving up.

It is also worth noting that while we report an objective measure of dialogue success (i.e. if both agents achieved the goal), from each agent’s perspective what is success may be different. For example, if a seeker does not inform about all constraints in the goal but provider respects all con-




Figure 2: Learning curves of the dialogue policies of our conversational agents, each evaluated on the environment it is trained on (see Table 3). Note that the agents are evaluated against each other, not against rational simulators or data.

straints that the provider does mention then the dialogue is successful from the provider’s perspective but failed from the seeker’s perspective. On the other hand, if the seeker provides all constraints and requests but the provider either ignores some constraints, says it cannot help, or does not address some requests then the dialogue is failed from the provider’s perspective but successful from the seeker’s perspective. To test whether optimizing the dialogue policies directly against these subjective measures of task success would lead to better dialogue policies, we performed similar experiments as the ones whose results are reported in Table 4. However, we found that the overall performance was not as good because it would lead to behaviours in which the

agents would not help each other to achieve the objective goal (e.g. the provider would not make many requests, or the seeker would not repeat informs upon wrong offers).

4 Conclusion

We presented the first complete attempt at concurrently training conversational agents that communicate only via self-generated language. Using DSTC2 as seed data, we trained NLU and NLG networks for each agent and let the agents interact and learn online optimal dialogue policies depending on their role (seeker or provider). Future directions include investigating joint optimization of the modules and training the agents online using deep multi-agent RL (e.g. (Foerster et al., 2018)) as well as evaluating our agents on harder environments (e.g. TextWorld (Côté et al., 2018)) and against human players. A natural extension is to train a multi-tasking provider agent that can learn to serve various kinds of seeker agents.

References

- Layla El Asri, Jing He, and Kaheer Suleman. 2016. A sequence-to-sequence model for user simulation in spoken dialogue systems. *INTERSPEECH*, pages 1151–1155.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. *ICLR*.
- Anja Belz and Ehud Reiter. 2006. Comparing automatic and human evaluation of nlg systems. In *11th Conference of the European Chapter of the Association for Computational Linguistics*.
- Michael Bowling and Manuela Veloso. 2000. An analysis of stochastic game theory for multiagent reinforcement learning. Technical report, Carnegie-Mellon Univ Pittsburgh Pa School of Computer Science.
- Michael H. Bowling and Manuela M. Veloso. 2001. Rational and convergent learning in stochastic games. In *IJCAI*, pages 1021–1026. Morgan Kaufmann.
- Senthilkumar Chandramohan, Matthieu Geist, Fabrice Lefèvre, and Olivier Pietquin. 2014. Co-adaptation in spoken dialogue systems. In *Natural Interaction with Robots, Knowbots and Smartphones*, pages 343–353. Springer.
- Sumit Chopra, Michael Auli, and Alexander M Rush. 2016. Abstractive sentence summarization with attentive recurrent neural networks. In *NAACL*, pages 93–98.
- Marc-Alexandre Côté, Ákos Kádár, Xingdi Yuan, Ben Kybartas, Tavian Barnes, Emery Fine, James Moore, Matthew Hausknecht, Layla El Asri, Mahmoud Adada, et al. 2018. Textworld: A learning environment for text-based games. *arXiv preprint arXiv:1806.11532*.

- Li Deng and Yang Liu, editors. 2018. *Deep Learning in Natural Language Processing*. Springer.
- Michael S English and Peter A Heeman. 2005. Learning mixed initiative dialog strategies by using reinforcement learning on both conversants. In *HLT-EMNLP*, pages 1011–1018. Association for Computational Linguistics.
- Jakob N. Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. 2018. Counterfactual multi-agent policy gradients. In *AAAI*, pages 2974–2982. AAAI Press.
- Jianfeng Gao, Michel Galley, and Lihong Li. 2019. [Neural approaches to conversational ai](#). *Foundations and Trends in Information Retrieval*, 13(2-3):127–298.
- Milica Gasic, Catherine Breslin, Matthew Henderson, Dongho Kim, Martin Szummer, Blaise Thomson, Pirros Tsiakoulis, and Steve J. Young. 2013. On-line policy optimisation of bayesian spoken dialogue systems via human interaction. In *ICASSP*, pages 8367–8371. IEEE.
- Kallirroi Georgila, Claire Nelson, and David Traum. 2014. Single-agent vs. multi-agent techniques for concurrent reinforcement learning of negotiation dialogue policies. In *ACL*, volume 1, pages 500–510.
- Dilek Hakkani-Tür. 2018. [Google assistant or my assistant? towards personalized situated conversational agents](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Plenary Talk*.
- Dilek Z. Hakkani-Tür, Gokhan Tur, Asli Celikyilmaz, Yun-Nung Chen, Jianfeng Gao, Li Deng, and Ye-Yi Wang. 2016. Multi-domain joint semantic frame parsing using bi-directional rnn-lstm. In *INTERSPEECH*.
- Matthew Henderson, Blaise Thomson, and Jason D Williams. 2014. The second dialog state tracking challenge. In *SIGDIAL*, pages 263–272.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional lstm-crf models for sequence tagging. *CoRR*, abs/1508.01991.
- Gakuto Kurata, Bing Xiang, Bowen Zhou, and Mo Yu. 2016. Leveraging sentence-level information with encoder lstm for semantic slot filling. In *EMNLP*.
- Irene Langkilde and Kevin Knight. 1998. Generation that exploits corpus-based statistical knowledge. In *COLING-ACL*, pages 704–710.
- Ji Young Lee and Franck Dernoncourt. 2016. Sequential short-text classification with recurrent and convolutional neural networks. In *HLT-NAACL*.
- Mike Lewis, Denis Yarats, Yann Dauphin, Devi Parikh, and Dhruv Batra. 2017. [Deal or no deal? end-to-end learning of negotiation dialogues](#). In *EMNLP*, pages 2443–2453.
- Xiujun Li, Yun-Nung Chen, Lihong Li, Jianfeng Gao, and Asli Celikyilmaz. 2017. End-to-end task-completion neural dialogue systems. In *IJCNLP*, pages 733–743. Asian Federation of Natural Language Processing.
- Bing Liu and Ian Lane. 2016. Attention-based recurrent neural network models for joint intent detection and slot filling. In *INTERSPEECH*.
- Bing Liu and Ian Lane. 2017. Iterative policy learning in end-to-end trainable task-oriented neural dialog models. In *ASRU*, pages 482–489. IEEE.
- Bing Liu and Ian Lane. 2018a. Adversarial learning of task-oriented neural dialog models. *SIGDIAL*, pages 350–359.
- Bing Liu and Ian Lane. 2018b. End-to-end learning of task-oriented dialogs. In *NAACL: Student Research Workshop*, pages 67–73.
- Grégoire Mesnil, Yann Dauphin, Kaisheng Yao, Yoshua Bengio, Li Deng, Dilek Hakkani-Tur, Xiaodong He, Larry Heck, Gokhan Tur, Dong Yu, et al. 2015. Using recurrent neural networks for slot filling in spoken language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(3):530–539.
- Alice H Oh and Alexander I Rudnicky. 2000. Stochastic language generation for spoken dialogue systems. In *Proceedings of the 2000 NAACL Workshop on Conversational systems-Volume 3*, pages 27–32. Association for Computational Linguistics.
- Alexandros Papangelis, Panagiotis Papadakos, Yannis Stylianou, and Yannis Tzitzikas. 2018. Spoken dialogue for information navigation. In *SIGDIAL*, pages 229–234. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.
- Ehud Reiter and Robert Dale. 2000. *Building natural language generation systems*. Cambridge university press.
- Verena Rieser and Oliver Lemon. 2010. Natural language generation as planning under uncertainty for spoken dialogue systems. In *EMNLP*, pages 105–120. Springer.
- Lina Maria Rojas-Barahona, Milica Gasic, Nikola Mrksic, Pei hao Su, Stefan Ultes, Tsung-Hsien Wen, and Steve J. Young. 2016. Exploiting sentence and context representations in deep neural models for spoken language understanding. In *COLING*.
- Jost Schatzmann, Blaise Thomson, Karl Weilhammer, Hui Ye, and Steve Young. 2007. Agenda-based user simulation for bootstrapping a pomdp dialogue system. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Companion Volume, Short Papers*, pages 149–152. Association for Computational Linguistics.
- Iulian Vlad Serban, Alessandro Sordani, Yoshua Bengio, Aaron C Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *AAAI*, volume 16, pages 3776–3784.
- Pararth Shah, Dilek Hakkani-Tur, Bing Liu, and Gokhan Tur. 2018. Bootstrapping a neural conversational agent with dialogue self-play, crowdsourcing and on-line reinforcement learning. In *NAACL, Volume 3 (Industry Papers)*, volume 3, pages 41–51.

- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. *Neural responding machine for short-text conversation*. In *ACL-IJCNLP*, pages 1577–1586.
- Satinder P. Singh, Michael J. Kearns, Diane J. Litman, and Marilyn A. Walker. 1999. Reinforcement learning for spoken dialogue systems. In *NIPS*, pages 956–962. The MIT Press.
- Amanda Stent, Matthew Marge, and Mohit Singhai. 2005. Evaluating evaluation methods for generation in the presence of variation. In *CICLing*, volume 3406 of *Lecture Notes in Computer Science*, pages 341–351. Springer.
- Pei-Hao Su, Pawel Budzianowski, Stefan Ultes, Milica Gasic, and Steve J. Young. 2017. Sample-efficient actor-critic reinforcement learning with supervised data for dialogue management. *SIGDIAL*, pages 147–157.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*, pages 3104–3112.
- G. Tur and R. De Mori, editors. 2011. *Spoken Language Understanding: Systems for Extracting Semantic Information from Speech*. John Wiley and Sons, New York, NY.
- Gökhan Tür, Li Deng, Dilek Hakkani-Tür, and Xiaodong He. 2012. Towards deeper understanding: Deep convex networks for semantic utterance classification. In *ICASSP*, pages 5045–5048. IEEE.
- Marilyn A Walker, Owen C Rambow, and Monica Rogati. 2002. Training a sentence planner for spoken dialogue using boosting. *Computer Speech & Language*, 16(3-4):409–433.
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrksic, Pei-hao Su, David Vandyke, and Steve J. Young. 2015. Semantically conditioned lstm-based natural language generation for spoken dialogue systems. *EMNLP*, pages 1711–1721.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrksic, Milica Gasic, Lina M Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017. A network-based end-to-end trainable task-oriented dialogue system. *EACL*, pages 438–449.
- Jason D Williams, Kavosh Asadi, and Geoffrey Zweig. 2017. Hybrid code networks: practical and efficient end-to-end dialog control with supervised and reinforcement learning. In *ACL (Volume 1: Long Papers)*, volume 1, pages 665–677.
- Yuexin Wu, Xiujun Li, Jingjing Liu, Jianfeng Gao, and Yiming Yang. 2018. Switch-based active deep dyna-q: Efficient adaptive planning for task-completion dialogue policy learning. *arXiv preprint arXiv:1811.07550*.
- Puyang Xu and Ruhi Sarikaya. 2013. Convolutional neural network based triangular CRF for joint intent detection and slot filling. *ASRU*, pages 78–83.
- Jun Yin, Xin Jiang, Zhengdong Lu, Lifeng Shang, Hang Li, and Xiaoming Li. 2016. Neural generative question answering. In *IJCAI’16*, pages 2972–2978. AAAI Press.
- Steve Young, Milica Gašić, Blaise Thomson, and Jason D Williams. 2013. Pomdp-based statistical spoken dialogue systems: A review. *Proceedings of the IEEE*, 101(5):1160–1179.
- Heiga Zen, Keiichi Tokuda, and Alan W. Black. 2009. Statistical parametric speech synthesis. *Speech Communication*, 51(11):1039–1064.
- Xiaodong Zhang and Houfeng Wang. 2016. A joint model of intent determination and slot filling for spoken language understanding. In *IJCAI*.

A Example Dialogues

Here we show two example dialogues generated by the *WoLF-PHC* agents.

Seeker goal	Constr(food=italian), Constr(pricerange=cheap), Req(addr), Req(phone)
Provider	hello , welcome to the cambridge restaurant System?
Seeker	you can ask for restaurants by area , price range or food type .
Provider	cheap restaurant
Seeker	what part of town do you have in mind?
Provider	italian food
Seeker	pizza hut city centre is a nice restaurant in the centre of town serving italian food
Provider	phone number
Seeker	the phone number is 01223 323737 .
Provider	address
Seeker	sure , it is on regent street city centre
Provider	(no output)
Seeker	what part of town do you have in mind?
Provider	good bye
Seeker	good bye

Table 5: A good example dialogue.

Seeker goal	Constr(food=african), Constr(pricerange=expensive), Req(addr)
Provider	hello , welcome to the cambridge restaurant System? how may i help you?
Seeker	expensive restaurant
Provider	what part of town do you have in mind?
Seeker	african food
Provider	bedouin is an expensive restaurant in the centre of town serving african food
Seeker	i do not care
Provider	sorry i am a bit confused ; please tell me again what you are looking for .
Seeker	repeat
Provider	what part of town do you have in mind?
Seeker	address
Provider	sure , it is on 100 mill road city centre
Seeker	good bye
Provider	good bye

Table 6: An imperfect example dialogue.

Scoring Interactional Aspects of Human–Machine Dialog for Language Learning and Assessment using Text Features

Vikram Ramanarayanan[†], Matt Mulholland[‡] and Yao Qian[†]

Educational Testing Service R&D

[†]90 New Montgomery Street, Suite 1500, San Francisco, CA

[‡]660 Rosedale Rd., Princeton, NJ

<vramanarayanan, mmulholland, yqian>@ets.org

Abstract

While there has been much work in the language learning and assessment literature on human and automated scoring of essays and short constructed responses, there is little to no work examining text features for scoring of dialog data, particularly interactional aspects thereof, to assess conversational proficiency over and above constructed response skills. Our work bridges this gap by investigating both human and automated approaches towards scoring human–machine text dialog in the context of a real-world language learning application. We collected conversational data of human learners interacting with a cloud-based standards-compliant dialog system, triple-scored these data along multiple dimensions of conversational proficiency, and then analyzed the performance trends. We further examined two different approaches to automated scoring of such data and show that these approaches are able to perform at or above par with human agreement for a majority of dimensions of the scoring rubric.

Index Terms: dialog systems, computer assisted language learning, conversational assessment, dialog scoring, intelligent tutoring systems.

1 Introduction

Learning and assessment solutions in today’s educational marketplace are placing increasing importance and resources on developing technologies that are dialogic (as opposed to monologic) in nature. Conversational proficiency is a crucial skill for success in today’s workplace (Weldy and Icenogle, 1997; Oliveri and Tannenbaum, 2019), which makes R&D on technologies that help develop and assess this skill important to complement our understanding from sociolinguistics (see for example Young, 2011; Doehler and Pochon-Berger, 2015). Dialog system technologies are

one solution capable of addressing and automating this need by allowing learners to practice and improve their interactional competence at scale (Suendermann-Oeft et al., 2017; Yu et al., 2019). However, such conversational technologies need to be able to provide targeted and actionable feedback to users in order for them to be useful to learners and widely adopted. Automated scoring of multiple aspects of conversational proficiency is one way to address this need.

While the automated scoring of text and speech data has been a well-explored topic for several years, particularly for essays and short constructed responses in the case of the former (Shermis and Burstein, 2013; Burrows et al., 2015; Madnani et al., 2017) and monolog speech for the latter (Neumeyer et al., 2000; Witt and Young, 2000; Xi et al., 2012; Bhat and Yoon, 2015), there has been a relative dearth of work on the *interpretable* automated scoring of dialog. Evanini et al. (2015) examined the automatic scoring of pseudo-dialogues, i.e., there were no branching dialog states; the system’s response was fixed and did not vary based on the learner’s response. Litman et al. (2016) developed a system to predict expert human rater scores based on audio signal and fluency features. Ramanarayanan et al. (2017a) analyzed this scoring problem at the level of each response in the dialog (i.e., each turn) instead of the entire conversation and across multiple dimensions of speaking proficiency. However, no study has performed a comprehensive examination of the automated scoring of *content* of whole dialog responses (with branching) based primarily on text features, based on a comprehensive multidimensional rubric and scoring paradigm designed specifically for dialog data, and interaction aspects in particular.

This study describes our contributions toward (i) developing a comprehensive rubric design

Table 1: *Human scoring rubric for interaction aspects of conversational proficiency. Scores are assigned on a Likert scale from 1-4 ranging from low to high proficiency. A score of 0 is assigned when there were issues with audio quality or system malfunction or off-topic or empty responses.*

Construct	Sub-construct	Description
Interaction	Engagement	Examines the extent to which the user engages with the dialog agent and responds in a thoughtful manner.
	Turn Taking	Examines the extent to which the user takes the floor at appropriate points in the conversation without noticeable interruptions or gaps.
	Repair	Examines the extent to which the user successfully initiates and completes a repair in case of a misunderstanding or error by the dialog agent.
	Appropriateness	Examines the extent to which the user reacts to the dialog agent in a pragmatically appropriate manner.
Overall Holistic Performance	Measures the overall performance.	

specifically tailored to conversational dialog along multiple dimensions, particularly those focused on interaction, (ii) triple-scoring a selection of dialog data based on this rubric, and finally (iii) examining the performance of two methods for automated scoring of such data – the first a state-of-the-art feature engineering method that passes word and character n -grams, length and syntax features into multiple state-of-the-art classifiers, and the second a model engineering method that leverages end-to-end memory networks to model dependencies between turn and prompt histories using memory components – and analyzing this performance vis-a-vis human raters. Note that for the purposes of this paper, while our data is spoken dialog, we will focus on text features derived from transcriptions, and therefore will focus on how they can be used to score various aspects of interaction in an interpretable manner. A subsequent future analysis will comprehensively examine how these can be combined with speech features.

2 Data

2.1 Collection

We crowdsourced, using Amazon Mechanical Turk, the collection of 2288 conversations of non-native speakers interacting with a dialog application designed to test general English speaking competence in workplace scenarios, and pragmatic skills in particular. The application, dubbed “Request Boss” requires participants to interact with their boss and request a meeting with her to review presentation slides using pragmatically appropriate language. To develop and deploy this application, we leveraged HALEF¹, an open-source modular cloud-based dialog system that is compatible with multiple W3C and open industry stan-

dards (Ramanarayanan et al., 2017b). The HALEF dialog system logs speech data collected from participants to a data warehouse, which are then transcribed and scored.

2.2 Human Scoring

In order to understand how well participants performed in our conversational task, we had each of the 2288 dialog responses triple scored by human expert raters on a custom-designed rubric. This rubric was iteratively modified and refined to score constructs specific to dialog data². The final conversational scoring rubric defined 12 sub-constructs under the 3 broad constructs of linguistic control, task fulfillment and interaction, apart from an overall holistic score. However, for purposes of this first study, we will focus on the relatively understudied interaction construct, in particular aspects of engagement, turn-taking, repair and (pragmatic) appropriateness. See Table 1 for more details. We asked expert raters to score each dialog for each rubric dimension on a scale from 1 to 4, and to assign dialogs that contained no or corrupted or significantly off-topic audio responses a score of 0. The expert raters were scoring leaders with significant experience in scoring various spoken and written assessments of English language proficiency. We used an automatic randomized design to assign three (out of eight possible) raters to every dialog such that (i) all raters had a commensurate number of responses to rate, and (ii) the same group of raters did not rate the same set of files (achieved by randomization; this prevents unwitting biases due to individual raters affecting the overall score analysis).

²Three scoring leaders first collaboratively adapted a rubric originally developed to score spoken interaction based on selected benchmark dialog responses. Based on this modified rubric and accompanying scoring notes specific to the task, 8 scoring leaders performed the final round of scoring.

¹<http://halef.org>

Table 2: *c-rater ML features used for machine scoring.*

Feature	Description
Word n -grams	Word n -grams are collected for $n = 1$ to 2. This feature captures patterns about vocabulary usage (key words) in responses.
Character n -grams	Character n -grams (including whitespace) are collected for $n = 2$ to 5. This feature captures patterns that abstract away from grammatical and other language use errors.
Response length	Defined as $\log(chars)$, where $chars$ represents the total number of characters in a response.
Syntactic dependencies	A feature that captures grammatical relationships between individual words in a sentence. This feature captures linguistic information about “who did what to whom” and abstracts away from a simple unordered set of key words.

3 Machine Scoring

This section first lays out our setup for interpretable machine scoring including details of the feature extraction and machine learning methods. We then analyze human performance (by examining inter-rater statistics) and use this to benchmark the performance of machine scoring methods. Following standardized convention in automated scoring, we only consider dialogs with a non-zero score to train scoring models (because a separate filtering mode is typically trained to eliminate “unscorable” responses, which include responses with no, garbled or out-of-topic audio data, see Higgins et al., 2011, for a more detailed motivation and rationale for this approach).

3.1 Feature Engineered Content Scoring

We used a set of features that have been employed in many previously published approaches to building content scoring models (see Madnani et al., 2017, 2018, for instance). We refer to this system as *c-rater ML*; see Table 2 for more details. All of the features are binary (indicating presence or absence) and try to capture how well responses contain (a) the right concepts (approximately captured by words and bigrams), (b) the right syntactic relationships between those concepts (approximately captured by dependency triples), (c) spelling and morphological relations (character n -grams) and (d) length of the response (captured by length features).

We used SKLL,³ an open-source Python package that wraps around the *scikit-learn* package (Pedregosa et al., 2011) to perform machine learning experiments. We experimented with rescaled linear support vector machine (SVM) and multi-layer perceptron (MLP) regressors. The former

allows us to interpret how the algorithm performs, while the latter is used for comparison purposes to understand how deep neural networks might perform on this task given the data we have. In our case, we found that the SVM classifier beat the MLP across the board, possibly because our feature space is sparse and high-dimensional, consisting of binary presence/absence features. We ran 10 fold cross-validation experiments and report the best overall results for the SVM system. We used cross entropy (log-loss) as an objective function for optimizing learner performance. We further tuned and optimized the free parameters of each learner using a grid-search method. We computed both accuracy and quadratic weighted kappa (which takes into account the ordered nature of the categorical labels) as metrics, reported in Table 3.

3.2 End to End Memory Network (MemN2N) architecture

We also investigated the efficacy of the End to End Memory Network (MemN2N) architecture (Sukhbaatar et al., 2015; Chen et al., 2016) adapted to the dialog scoring task. The end to end MemN2N architecture models dependencies in text sequences using a recurrent attention model coupled with a memory component, and is therefore suited to modeling how response and prompt histories contribute to a dialog score. In our case, the MemN2N architecture learns a mapping between an output score and an input tuple consisting of the current response, the response history and the prompt history. See Figure 1. We modified the original MemN2N architecture in Sukhbaatar et al. (2015) in the following ways: (i) instead of the original (*query, fact history, answer*) tuple that is used to train the network in the original paper, we have an (*current response, response his-*

³<https://github.com/EducationalTestingService/skll>

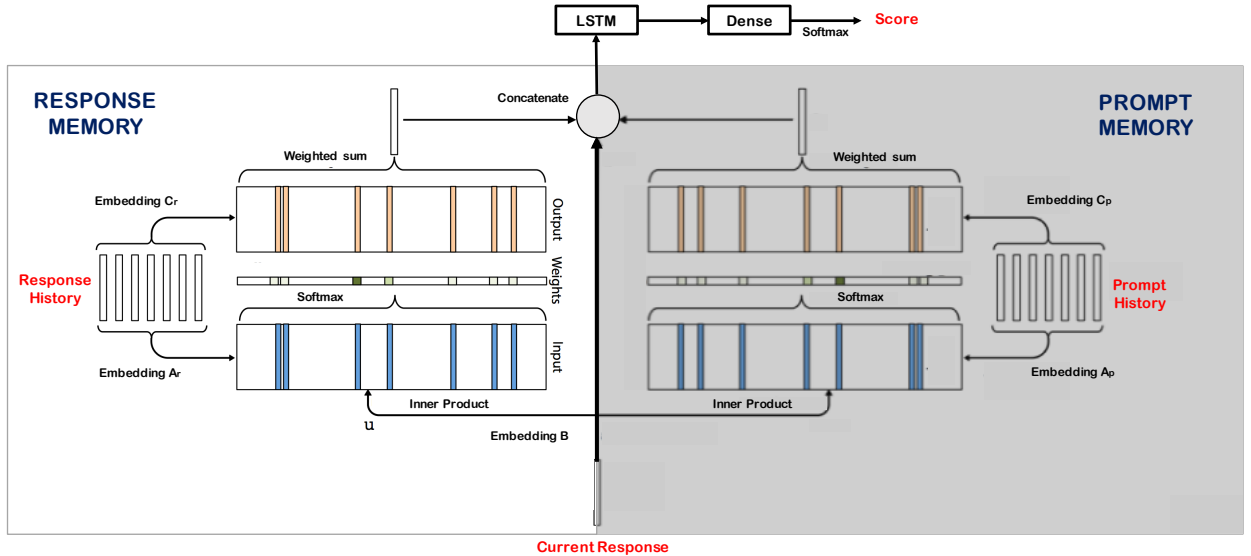


Figure 1: Schematic of a *single hop* module of our modified end-to-end memory network (MemN2N) adapted from Sukhbaatar et al. (2015) for our dialog scoring experiments. Stacking modules on top of each other allows us to model multiple hops.

Table 3: *Human and machine score statistics*

Construct	Sub-construct	c-rater ML		MemN2N		c-rater ML + MemN2N		Human Metrics	
		Accuracy	QW_{κ}	Accuracy	QW_{κ}	Accuracy	QW_{κ}	Conger κ	Krippendorff α
Interaction	Engagement	0.70	0.70	0.65	0.65	0.71	0.72	0.69	0.72
	Turn Taking	0.69	0.67	0.68	0.40	0.71	0.70	0.71	0.74
	Repair	0.66	0.60	0.64	0.58	0.67	0.64	0.73	0.72
	Appropriateness	0.67	0.67	0.62	0.58	0.67	0.67	0.70	0.72
Overall Holistic Performance		0.69	0.72	0.66	0.65	0.70	0.72	0.75	0.75

tory, prompt history, score) tuple in our case. In other words, we not only embed and learn memory representations between the current response and the history of previous responses, but the history of prior system prompts that have been encountered thus far; (ii) we used an LSTM instead of a matrix multiplication at the final step of the network before prediction; and (iii) we experimented with *Google word2vec* (Mikolov et al., 2013) and *GloVe* (Pennington et al., 2014) initializations for word embeddings in addition to experimenting with multiple memory hops. We train the network at the turn level; in other words, for each turn, the training data would consist of an input of (*response for current turn, response history, prompt history*) and an output of the *dialog-level score* (in other words, each turn is assumed to have the same score as that of the full dialog). During testing, we compute the score for each dialog in the test set as the median of scores predicted by the trained network for each turn in that dialog.

We used a similar crossvalidation setup as described in §3.1 with the exact same 10 folds with

experiments optimizing a cross-entropy-based objective function as in the earlier case to enable a fair comparison. We tuned hyperparameters of the network using the *hyperas* toolkit⁴. This included the number of neurons in the *Dense* and *LSTM* layers as well as the addition of *Dropout* layers after each memory component. We experimented with 1, 2 and 3 memory hops and found 2 to be optimal. Interestingly, we also found that initializing the memory embedding matrices with pretrained *Google word2vec* or *GloVe* embeddings worked better than randomly-initialized ones for prompt history encoding as compared to response history encoding.

4 Observations and Results

The final two columns of Table 3 display two interrater agreement statistics – Conger κ and Krippendorff α – for the human expert scores assigned to the data. Recall that each dialog was scored by 3 out of 8 possible raters. We observe a moderate to high agreement between raters for all dimensions

⁴<http://maxpumperla.com/hyperas/>

of the scoring rubric, which is not too surprising given that all our raters had significant experience in rating monologic speech data.

Table 3 also shows the performance of our two different systems in scoring various aspects of interaction at the level of the entire dialog. Observe that fusing the MemN2N with the c-rater ML system leads to a small but significant improvement over either of the systems alone. Additionally, it is interesting to note that the quadratic weighted kappa ($QW\kappa$) of the fusion system is in a similar ballpark as the κ and α metrics for human inter-rater agreement, particularly for engagement and turn-taking subscores. While these measures are *not directly comparable*, this trend is encouraging nonetheless, suggesting that a combination of n -gram, length, syntactic dependency and memory-based attention over embedding representations of words over the entire dialog are useful in capturing at least some aspects of these sub-constructs of interaction. On the other hand, the fusion system performance for repair and appropriateness subscores is still below par, suggesting that more feature engineering and modeling research is required to model these aspects of interaction. These dimensions of interaction are also harder to predict, given that repair and pragmatic appropriateness are more high-level and abstract in nature.

5 Discussion

This paper has examined approaches to both human and machine scoring of text dialogs collected as part of a language learning application, particularly looking at interactional aspects. We observed, through careful design of the human scoring paradigm, a moderate-to-high agreement between the raters. We further examined two methods for automated scoring of such data – the first a feature engineering method that passes word and character n -grams, length and syntax features into an SVM based classifier, and the second a model engineering method that leverages end-to-end memory network (MemN2N) to model dependencies between turn and prompt histories using memory components – and found that a fusion of both methods performs close to or at par with human inter-rater agreement statistics.

While our results are encouraging, there is still much work ahead in understanding and scoring interactional competence. One of the key reasons for this has to do with the fact that the features

were considered were text-based, and it is unclear how some features that don't directly consider information from audio or visual channels are useful in predicting properties related to interaction (engagement, for instance). Repair and appropriateness, and even turn taking to a lesser extent are related to proficiency in language use, and hence it makes sense that features such as n -grams and syntax use might be somewhat useful in predicting these aspects of interaction. However, some of the results might also be explained by some of our examined features being highly correlated with more interpretable/relevant features. For instance, length might be an indication of a more proficient and verbose speaker, which might in turn correlate with a high level of engagement. Nonetheless, an understanding of how meaningful our text-based results are will be incomplete without examining features derived from audio (and visual streams, if available), including non-verbal and prosodic cues.

It is also worth mentioning tangentially related work on dialog interaction quality at this point (see for instance [Schmitt and Ultes, 2015](#); [Stoyanchev et al., 2019](#); [See et al., 2019](#)). While such work primarily focuses on investigating techniques to measure and improve the quality of the overall dialog interaction as opposed to providing targeted assessment and feedback on the quality of spoken language used during interactions, it might nonetheless be useful to take this body of work into account while developing techniques for automated proficiency scoring.

This lays out multiple avenues for future work. First, as mentioned earlier, would be examining both text and speech signals for a more complete examination of the scoring problem. Second, we would like to look at other broad aspects of conversational proficiency, such as delivery (for instance, fluency, intonation, vocabulary and grammar) and topic development (elaboration and task specificity, for example) in addition to building on the interaction aspects described here. Third, we will investigate combining feature-engineering and model-engineering approaches towards developing specific features and model architecture improvements that will help us push the automated scoring performance even higher. These will feed into our ultimate goal of being able to provide language learners with targeted, actionable feedback on different facets of conversational proficiency.

References

- Suma Bhat and Su-Youn Yoon. 2015. Automatic assessment of syntactic complexity for spontaneous speech scoring. *Speech Communication*, 67:42–57.
- Steven Burrows, Iryna Gurevych, and Benno Stein. 2015. The eras and trends of automatic short answer grading. *International Journal of Artificial Intelligence in Education*, 25(1):60–117.
- Yun-Nung Chen, Dilek Hakkani-Tür, Gökhan Tür, Jianfeng Gao, and Li Deng. 2016. End-to-end memory networks with knowledge carryover for multi-turn spoken language understanding. In *Interspeech*, pages 3245–3249.
- S Pekarek Doehler and Evelyne Pochon-Berger. 2015. The development of 12 interactional competence: evidence from turn-taking organization, sequence organization, repair organization and preference organization. *Usage-based perspectives on second language learning*, 30:233.
- Keelan Evanini, Sandeep Singh, Anastassia Loukina, Xinhao Wang, and Chong Min Lee. 2015. Content-based automated assessment of non-native spoken language proficiency in a simulated conversation. In *Proceedings of the Machine Learning for SLU & Interaction NIPS 2015 Workshop*.
- Derrick Higgins, Xiaoming Xi, Klaus Zechner, and David Williamson. 2011. A three-stage approach to the automated scoring of spontaneous spoken responses. *Computer Speech & Language*, 25(2):282–306.
- Diane Litman, Steve Young, Mark Gales, Kate Knill, Karen Ottewell, Rogier van Dalen, and David Vandyke. 2016. Towards using conversations with spoken dialogue systems in the automated assessment of non-native speakers of english. In *17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, page 270.
- Nitin Madnani, Aoife Cahill, Daniel Blanchard, Slava Andreyev, Diane Napolitano, Binod Gyawali, Michael Heilman, Chong Min Lee, Chee Wee Leong, Matthew Mulholland, et al. 2018. A robust microservice architecture for scaling automated scoring applications. *ETS Research Report Series*, 2018(1):1–8.
- Nitin Madnani, Anastassia Loukina, and Aoife Cahill. 2017. A large scale quantitative exploration of modeling strategies for content scoring. In *Proceedings of the 12th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 457–467.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Leonardo Neumeier, Horacio Franco, Vassilios Digalakis, and Mitchel Weintraub. 2000. Automatic scoring of pronunciation quality. *Speech communication*, 30(2):83–93.
- Maria Elena Oliveri and Richard J Tannenbaum. 2019. Are we teaching and assessing the english skills needed to succeed in the global workplace? *The Wiley Handbook of Global Workplace Learning*, pages 343–354.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Vikram Ramanarayanan, Patrick L Lange, Keelan Evanini, Hillary R Molloy, and David Suendermann-Oeft. 2017a. Human and automated scoring of fluency, pronunciation and intonation during human-machine spoken dialog interactions. In *INTERSPEECH*, pages 1711–1715.
- Vikram Ramanarayanan, David Suendermann-Oeft, Patrick Lange, Robert Mundkowsky, Alexei V Ivanov, Zhou Yu, Yao Qian, and Keelan Evanini. 2017b. Assembling the Jigsaw: How Multiple Open Standards Are Synergistically Combined in the HALEF Multimodal Dialog System. In *Multimodal Interaction with W3C Standards*, pages 295–310. Springer.
- Alexander Schmitt and Stefan Ultes. 2015. Interaction quality: assessing the quality of ongoing spoken dialog interaction by experts and how it relates to user satisfaction. *Speech Communication*, 74:12–36.
- Abigail See, Stephen Roller, Douwe Kiela, and Jason Weston. 2019. What makes a good conversation? how controllable attributes affect human judgments. *arXiv preprint arXiv:1902.08654*.
- Mark D Shermis and Jill Burstein. 2013. *Handbook of automated essay evaluation: Current applications and new directions*. Routledge.
- Svetlana Stoyanchev, Soumi Maiti, and Srinivas Bangalore. 2019. Predicting interaction quality in customer service dialogs. In *Advanced Social Interaction with Agents*, pages 149–159. Springer.
- David Suendermann-Oeft, Vikram Ramanarayanan, Zhou Yu, Yao Qian, Keelan Evanini, Patrick Lange, Xinhao Wang, and Klaus Zechner. 2017. A multimodal dialog system for language assessment: Current state and future directions. *ETS Research Report Series*, 2017(1):1–7.

- Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. 2015. End-to-end memory networks. In *Advances in neural information processing systems*, pages 2440–2448.
- Teresa G Weldy and Marjorie L Icenogle. 1997. A managerial perspective: Oral communication competency is most important for business students in the workplace. In *The Journal of Business Communication*, 34(1):67–80.
- Silke M Witt and Steve J Young. 2000. Phone-level pronunciation scoring and assessment for interactive language learning. *Speech communication*, 30(2):95–108.
- Xiaoming Xi, Derrick Higgins, Klaus Zechner, and David Williamson. 2012. A comparison of two scoring methods for an automated speech scoring system. *Language Testing*, 29(3):371–394.
- Richard F Young. 2011. Interactional competence in language learning, teaching, and testing. *Handbook of research in second language teaching and learning*, 2(426-443).
- Zhou Yu, Vikram Ramanarayanan, Patrick Lange, and David Suendermann-Oeft. 2019. An open-source dialog system with real-time engagement tracking for job interview training applications. In *Advanced Social Interaction with Agents*, pages 199–207. Springer.

Spoken Conversational Search for General Knowledge

Lina M. Rojas-Barahona, Pascal Bellec, Benoit Besset, Martinho Dos-Santos, Johannes Heinecke, Munshi Asadullah, Olivier Le-Blouch, Jean Y. Lancien, Géraldine Damnati, Emmanuel Mory and Frédéric Herledan

Orange Labs, 2 Avenue de Pierre Marzin, Lannion, France

{linamaria.rojasbarahona,pascal.bellec,benoit.besset,martinho.dossantos,
johannes.heinecke,munshi.asadullah,olivier.leblouch,jeanyves.lancien
geraldine.damnati,emmanuel.mory,frederic.herledan}@orange.com

Abstract

We present a spoken conversational question answering proof of concept that is able to answer questions about general knowledge from Wikidata¹. The dialogue component does not only orchestrate various components but also solve coreferences and ellipsis.

1 Introduction

Conversational question answering is an open research problem. It studies the integration of *question answering* (QA) systems in a *dialogue system* (DS). Not long ago, each of these research subjects were studied separately; only very recently has studying the intersection between them gained increasing interest (Reddy et al., 2018; Choi et al., 2018).

We present a spoken conversational question answering system that is able to answer questions about general knowledge in French by calling two distinct QA systems. It solves coreference and ellipsis by modelling context. Furthermore, it is extensible, thus other components such as neural approaches for question-answering can be easily integrated. It is also possible to collect a dialogue corpus from its iterations.

In contrast to most conversational systems which support only speech, two input and output modalities are supported speech and text. Thus it is possible to let the user check the answers by either asking relevant Wikipedia excerpts or by navigating through the retrieved name entities or by exploring the answer details of the QA components: the confidence score as well as the set of explored triplets. Therefore, the user has the final word to consider the answer as correct or incorrect and to

¹<https://www.wikidata.org>

provide a reward, which can be used in the future for training reinforcement learning algorithms.

2 Architectural Description

The high-level architecture of the proposed system consists of a speech-processing front-end, an understanding component, a context manager, a generation component, and a synthesis component. The context manager provides contextualised mediation between the dialogue components and several question answering back-ends, which rely on data provided by Wikidata¹. Interaction with a human user is achieved through a graphical user interface (GUI). Figure 1 depicts the components together with their interactions.

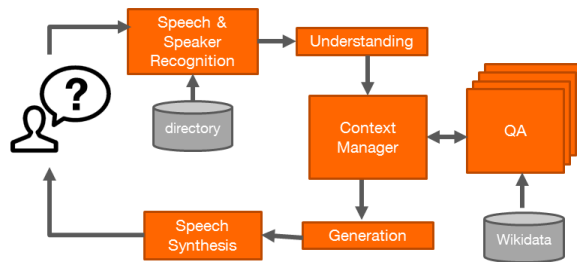


Figure 1: High-level depiction of the proposed spoken conversation question answering system. Arrows indicate data flow and direction.

In the remainder of this section, we explain the components of our system.

2.1 Speech and Speaker Recognition

The user vocally asks her question which is recorded through a microphone driven by the GUI. The audio chunks are then processed in parallel by a speech recognition component and a speaker recognition component.

Speech Recognition The Speech Recognition component enables the translation of

speech into text. Cobalt Speech Recognition for French is a Kaldi-based speech-to-text decoder using a TDNN (Povey et al., 2016) acoustic model; trained on more than 2000 hours of clean and noisy speech, a 1.7-million-word lexicon, and a 5-gram language model trained on 3 billion words.

Speaker Recognition The Speaker Recognition component answers the question “Who is speaking?”. This component is based on deep neural network speaker embeddings called “x-vectors” (Snyder et al., 2018). Our team participated to the NIST SRE18 challenge (Sadjadi et al., 2019), reaching the 11th position among 48 participants.

Once identified, it is possible to access the information of the speaker by accessing a speaker database which includes attributes such as nationality. This is a key module for personalising the behaviour of the system, for instance, by supporting questions such as “Who is the president of the country where I was born?”.

2.2 The Dialogue System

The transcribed utterance and the speaker information are passed to the dialogue system. This system contains an **understanding** component, a **context manager**, and a **generation** component (Figure 2).

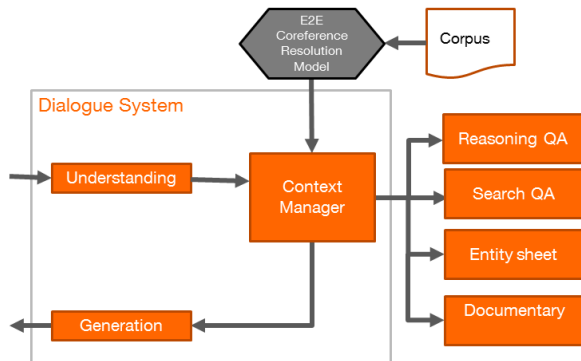


Figure 2: Internal structure of the proposed dialogue system, with emphasis placed on the interactions of the context manager.

Understanding The understanding component relies on a **linguistic module** to parser the user’s inputs. The linguistic module supports part-of-the-speech (POS) tagging, lemmatisation, dependency syntax and semantics provided by an adapted version

	Train	Dev	Test
words	208 245	45 001	89 330
sentences	10 166	2 976	4 853
mentions	15 013	3 008	6 232
incl. prons.	1 465	280	538
chains	3 793	901	1 533

Table 1: Subset of the corpus CALOR used for training, developing and testing of the coreference resolution module. Note that the values given for the mentions include pronouns.

of UDpipe (Straka and Straková, 2017), extended with a French full-form lexicon. UDpipe was trained on the French GSD treebank version 2.3². Since the syntax of questions in French differs from that of declaratives, we annotated manually about 500 questions to be merged into the UD treebank (which originally did not contain questions). Tests show that the labelled attached score (LAS) is thereby increased by 10% absolute, to 92%.

Context Manager The Context Manager component is able to solve coreferences by using an adaptation of the end-to-end model presented in (Lee et al., 2017), that we trained for French by using fasttext multilingual character embeddings (Bojanowski et al., 2017). The data used to train the coreference resolution model is a subset of the corpus CALOR (Marzinotto et al., 2018) (Table 1), which has been manually annotated with coreferences. This corpus contains coreference chains of named entities, nouns and pronouns (such as “the president” – “JFK” – “he” – “his”).

The dependency tree and semantic frames provided by the linguistic module are used to solve ellipsis by taking into account the syntactic and semantic structure of the previous question. Once the question has been resolved, it calls the QA systems and passes their results to the generation module.

Generation The generation component either returns the short answer provided by QA systems or relies on an external generation module that uses dependency grammar templates to generate more elaborated answers.

2.3 QA Systems

Two complementary question answering components were integrated into the system: the Reasoning QA and Search QA. Each of these

²<http://universaldependencies.org/>

QA systems computes a confidence score for every answer by using icsiboost (Favre et al., 2007), an Adaboost-based classifier trained on a corpus of around 21 000 questions. The Context Manager takes into account these scores to pick the higher-confidence of the two answers.

Besides the QA components, there are two other components that are able to provide complementary information about the Wikidata’s entities under discussion: Documentary and Entity Sheet.

Reasoning QA The Reasoning QA system first parses the question by using a Prolog definite clause grammar (DCG), extended with word-embeddings to support variability in the vocabulary. Then it explores a graph containing logical patterns that are used to produce requests in SPARQL³ that agree with the question.

Search QA The Search QA system uses an internal knowledge base, which finely indexes data using Elasticsearch. It is powered by Wikidata and enriched by Wikipedia, especially to calculate a Page-Rank (Page et al., 1997) on each entity. This QA system first determines the potential named entities in the question (i.e. subjects, predicates, and types of subjects). Second, it constructs a correlation matrix by looking for the triplets in Wikidata that link these entities. This matrix is filtered according the coverage of the question and the relevance of each entity in order to find the best answer.

Documentary The documentary component is able to extract pertinent excerpts of Wikipedia. It uses an internal documentary base, which indexes Wikipedia’s paragraphs by incorporating the Wikidata entity’s IDs into elasticsearch indexes. Thus, it is possible to find paragraphs (ranked by elasticsearch) illustrating the answer to the given question by taking into account the entities detected in the question and in the answer.

Entity Sheet The entity sheet component summarises an entity in Wikidata returning the description, the picture and the type of the entity.

³<https://www.w3.org/TR/sparql11-query/>




Figure 3: Distribution of question topics used to evaluate system performance on out-of-context questions.

2.4 Speech Synthesis

Finally, the generated response is passed to the GUI, which in turn passes it to the Voxygen synthesis solution.

3 Evaluation

The evaluation of the individual components of the proposed system was performed outside the scope of this work. We evaluated out-of-context questions, as well as the coreference resolution module.

Performance on out-of-context questions was evaluated on Bench’It, a dataset containing 150 open ended questions about general knowledge in French (Figure 3)⁴. The system reached a macro precision, recall and F-1 of 64.14%, 64.33% and 63.46% respectively⁵.

We also evaluated the coreference resolution model on the test-set of CALOR (Table 1), obtaining an average precision, recall and F-1 of 65.59%, 48.86% and 55.77% respectively. The same model reached a average F-1 of 68.8% for English (Lee et al., 2017). Comparable measurements are not available for French. F-1 scores for French are believed to be lower because of the lower amount of annotated data.

4 Examples

On the one hand, the system is able to answer complex out-of-context questions such as “What are the capitals of the countries of the Iberian Peninsula?”, by correctly answering the list of capitals: “Andorra la Vella, Gibraltar, Lisbon, Madrid”.

⁴Publicly available in <https://github.com/lmrojasb/benchit.git>

⁵Following the metrics of the task-4 of QALD-7 <https://project-hobbit.eu/challenges/qald2017/>

U: Who is Michael Jackson ?
 S: Michael Jackson is an American author,composer,
 singer and dancer
 U: What is his father's name?
 S: Joseph Jackson
 U: and his mother's?
 S: Katherine Jackson
 U: and his brothers' and sisters' ?
 S: Tito Jackson, Rebbie Jackson, Randy Jackson,
 Jackie Jackson, Marlon Jackson, La Toya Jackson,
 Jermaine Jackson, Janet Jackson

Figure 4: English translation of French conversation involving in-context questions.

On the other hand, consider the dialogue presented in Figure 4, in which the user asks several related questions about Michael Jackson. First she asks “Who is Michael Jackson?” and the system correctly answers “Michael Jackson is an American author, composer, singer and dancer”, note that this is the generated long answer.

The subsequent questions are related to the names of his family members. In order to correctly answer these questions, the resolution of coreferences is necessary to solve the possessive pronouns, which in French agree in gender and number with the noun they introduce. In this specific example, while in English “his” is used in all the cases, in French it changes to: *son père* (father), *sa mère* (mother), *ses frères* (brothers). This example also illustrates resolution of elliptical questions in the context, by solving the question “and his mother’s” as “What is the name of his mother”.

5 Conclusion and Future Work

We have presented a spoken conversational question answering system, in French. The DS orchestrates different QA systems and returns the response with the higher confidence score. The system contains modules specifically designed for dealing with common spoken conversation phenomena such as coreference and ellipsis.

We will soon integrate a state-of-the art reading comprehension approach, support English language and improve the coreference resolution module. We are also interested in exploring policy learning, thus the system will be able to find the best criterion to chose the answer or to ask for clarification in the case of ambiguity and uncertainty.

References

- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the ACL*, 5:135–146.
- Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wen-tau Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. **QuAC: Question answering in context**. In *Proceedings of EMNLP 2018*, pages 2174–2184, Brussels, Belgium.
- Benoit Favre, Dilek Hakkani-Tür, and Sebastien Cuendet. 2007. Icsiboost. <https://github.com/benob/icsiboost>.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. **End-to-end neural coreference resolution**. In *Proceedings of the 2017 EMNLP*, pages 188–197, Copenhagen, Denmark. Association for Computational Linguistics.
- Gabriel Marzinotto, Jeremy Auguste, Frédéric Béchet, Géraldine Damnati, and Alexis Nasr. 2018. **Semantic frame parsing for information extraction: The CALOR corpus**. In *LREC*, Miyazaki, Japan. ELRA.
- Larry Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1997. Pagerank: Bringing order to the web.
- Daniel Povey, Vijayaditya Peddinti, Daniel Galvez, Pegah Ghahramani, Vimal Manohar, Xingyu Na, Yiming Wang, and Sanjeev Khudanpur. 2016. Purely sequence-trained neural networks for asr based on lattice-free mmi. In *Proceedings of INTERSPEECH, 2016*.
- Siva Reddy, Danqi Chen, and Christopher D Manning. 2018. Coqa: A conversational question answering challenge. *arXiv preprint arXiv:1808.07042*.
- Seyed Omid Sadjadi, Craig S. Greenberg, Douglas A. Reynolds, Elliot Singer, Lisa P. Mason, , and Jaime Hernandez-Cordero. 2019. The 2018 nist speaker recognition evaluation. In *Proceedings of INTERSPEECH (submitted), Graz, Austria, August 2019*.
- David Snyder, Daniel Garcia-Romero, Gregory Sell, Daniel Povey, and Sanjeev Khudanpur. 2018. X-vectors: Robust dnn embeddings for speaker recognition. In *Proceedings of IEEE ICASSP, April 2018*.
- Milan Straka and Jana Straková. 2017. **Tokenizing, POS Tagging, Lemmatizing and Parsing UD 2.0 with UDPipe**. In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 88–99, Vancouver, Canada. ACL.

Graph2Bots, Unsupervised Assistance for Designing Chatbots

Jean-Léon Bouraoui, Sonia Le Meitour, Romain Carbou,
Lina M. Rojas-Barahona, Vincent Lemaire

Orange, 2 Avenue Pierre Marzin,
22300 Lannion, France

{jeanleon.bouraoui, sonia.lemeitour, romain.carbou,
linamaria.rojasbarahona, vincent.lemaire@orange.com}

Abstract

We present Graph2Bots, a tool for assisting conversational agent designers. It extracts a graph representation from human-human conversations by using unsupervised learning. The generated graph contains the main stages of the dialogue and their inner transitions. The graphical user interface (GUI) then allows graph editing.

1 Introduction

In the field of artificial intelligence, dialogue systems are gaining popularity, especially as they benefit from advances in the understanding of conversational contents and contexts. Mobile and home applications such as Siri (Apple), Google Assistant (Google), Cortana (Microsoft) or Alexa (Amazon) are the most popular. To quantify this growing interest in human-machine interfaces, and dialogue systems in particular, let us cite the studies by the analyst firm Gartner¹. They place dialogue systems among the 10 strategic technologies from 2018 and for the coming years.

One of the current trends is to propose software tools to assist the design of dialogue systems. These tools are customized according to the designer's needs, and the domain of application (e.g. trips reservation). Some solutions allow designing the dialogue architecture through a GUI. However, designers still have to perform this task manually, based on their domain knowledge and eventually the analysis of human conversations on a similar task. Most of the existing solutions do not provide a robust possibility to quickly set up an automated dialogue from human-human conversations.

In this context, we present an unsupervised assistant for the creation and adaptation of dialogue

systems when the designer has a corpus of human-human interactions. This is our main contribution : this Proof of Concept can be applied for any application domain, without any prior knowledge. Thus the user will have a first version of the dialogue architecture ; that he can refine it with the GUI.

The remainder of the paper is organized as follows: first we describe the motivation by detailing the problem in Section 2; in Section 3 we present the prototype and its main innovative features. Later in Section 4 we explain the method of unsupervised learning used to build the graph. Section 5 shows the GUI. Finally, we present the conclusions and future work.

2 Rationale

Throughout this document a dialogue is an exchange of information between two speakers (a human or a machine). We are interested in task-oriented dialogues: the speakers will collaborate to achieve a common goal.

Modeling a dialogue agent specialized in a given domain is mostly done manually: either a priori, from the knowledge that the designer has on the task; or a posteriori, from the consultation of existing corpora as shown in the figure 1 below; in both cases, the process is time-consuming.




Figure 1: A posteriori Use Case workflow

We work in the “a posteriori” use case. Our goal

¹<https://www.gartner.com/smarterwithgartner/gartner-top-10-strategic-technology-trends-for-2019>

is to automatically obtain the dialog graph from the corpus.

We call “dialogue corpus” a set of n dialogues related to a particular domain e.g. of transcripts of train reservation dialogues. Each dialogue is composed of t speech turns; a speech turn corresponds to what is said by one of the speakers without any interruption.

We want to automatically determine from the corpus: (i) the different phases of the dialogue (including expressed intentions - hereafter referred to as “themes”. This term corresponds to either generic themes or sub-goals of the dialogue); (ii) transitions between phases. The goal is to obtain an oriented graph showing the main transitions between themes. Our assumption is that, depending on the position in the dialogue, a given turn is more likely to belong to a given phase than another; this information is therefore taken into account during the process.

The obtained graph can be exploited in multiple ways. For instance, it can be used to initialize a dialogue agent. That is to say, it can serve as a basis for modeling a dialogue agent specialized in a given domain, facilitating its design. In addition, the graph, as well as the steps taken to obtain it, will allow the designer, without prior knowledge of the application domain, to have a first understanding of the topics of the dialogues, their structure, and more generally their knowledge. Thus he can quickly get the most relevant information for the conception of the dialogue agent.

3 The Architecture




Figure 2: Graph2Bots architecture

Imagine that a designer wants to set up a conversational agent for a specific application domain. He owns a domain-related corpus of dialogues. First, he will use the unsupervised approach offered by our tool (i.e. the back-end) to identify the

underlying dialogue phases, and their transitions. He will then edit the obtained graph by using our GUI. We describe these different steps in the following sections. By this way the role of the designer is changed (from the Figure 1 to the Figure 2) from a specialized dialogue analyst to a dialog graph editor.

4 Unsupervised exploitation of large dialogue corpus

In this section we will describe the three steps necessary to generate the dialogue graph from the corpus: the pre-processing to normalize and prepare the data, then the co-clustering to group speech turns in clusters corresponding to the dialogue phases and finally the graph generation.

4.1 Preprocessing

The corpus contains text documents with one speech turn per line identified by: the dialogue it belongs to, its position in the dialogue and the speaker. We begin with an anonymization process to replace the named entities (like customer name, phone number, address, etc.) by a tag. Then, the corpus is filtered in order to remove the “stopwords”. These are words that do not convey semantic information (e.g. prepositions, articles, etc.). We used the list of “stopwords” provided by the NLTK library². Finally we calculate the frequency of the words in the corpus.

4.2 Co-clustering

To identify the dialogue phases, a co-clustering technique (Guigourès, 2013) is used to obtain clusters of speech turns considering the words they contain. We consider that each speech turn corresponds to a text document, being composed of words. This can be represented by a matrix. We use the technique of co-clustering to discover the best reordering and grouping of lines and columns. The method used is based on the Minimum Optimized Description Length (MODL) approach described in (Boullé, 2011) (a tutorial of this approach may be found in (Boullé et al., 2013)). Each speech turn is assigned to one and only one cluster but a cluster can group distinct dialogue turns. This method finds by itself the right number K of clusters and does not require any parameter, which is useful for non-experts. The algorithm maximizes the mutual information between

²http://www.nltk.org/nltk_data

the two clusterings (one partition corresponds to the group of speech turns, and the other to the group of words). A post-processing, an Agglomerative Hierarchical Clustering (Guigourès, 2013), allows to reduce the number of clusters and simplify the result.

It then remains to determine the transitions between the dialogue phases.

4.3 Graph Generation

The desired graph has to represent the architecture of the dialogue, that is to say, the succession of the phases from the beginning to the end of the dialog and the various possible paths. The figure 3 illustrates a part of such a graph.




Figure 3: Graph for a train booking application

In our approach, a phase corresponds to a cluster of turns of speech. Ideally, these speech turns are homogeneous in relation to a given theme. After the discovery of the different phases (the clusters) the connections are computed according to the frequency of phase successions in the corpus: the dialogue turns are "projected" onto the clusters and the number of transitions between the clusters along the dialogues are counted.

This process is illustrated in the figure 4; 'User' represents a human customer speech turn, 'Agent' represents a human agent or a bot and T_i corresponds to the cluster identifier. When a speech turn from cluster T_1 is followed by another from cluster T_3 in a dialogue, it results in a transition from T_1 to T_3 in the graph.

The resulting representation is an oriented graph, whose vertices are the clusters, and whose weighted edges are transitions between clusters. For the interested reader more details of this phase are given in (Bouraoui and Lemaire, 2017).

5 Graphical User Interface

The GUI proposes to visualize interactively in real-time the data processed in the back-end, in the form of graphs. We describe the main features below.

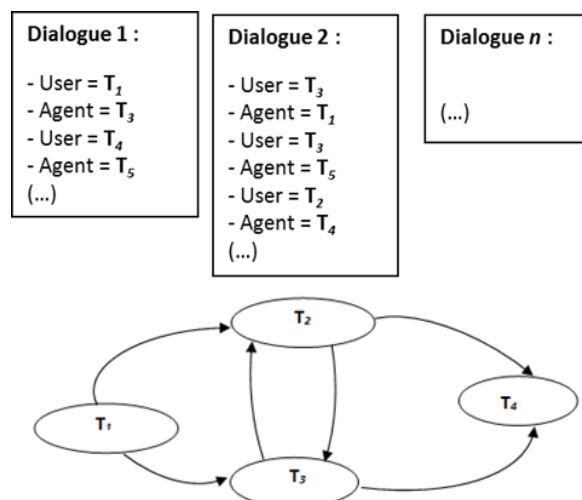


Figure 4: Graph generation

5.1 Real Time and Interactive Display

Here is an example of the list of real time and interactive display functionalities :

- Choice of the granularity of display, according to the size of clusters or the frequency of relations;
- Use of the mouse pointer to "pull" a cluster away from others, to select one or more clusters, to zoom in and out, and so on;
- Ability to rename clusters, which are automatically named with the two most representative words and the speaker; and to browse through their contents: the most representative words, and the corresponding speech turns.

5.2 Dialogue Graph Edition

The designer of the interacting agent can adapt and refine the architecture according to his needs. It is possible to modify:

- The contents of a given cluster by deleting one or several speech turns;
- The architecture itself. Two main features are available. One is the **fusion** of two clusters (if they are thematically similar and therefore redundant). The other is the **split** of a cluster in two groups when speech turns are semantically similar, but heterogeneous with respect to the main theme expressed in the cluster. If any of these features are used, the display of the number of clusters and their connections is updated dynamically.

Figure 5 shows the current version of the prototype GUI. The graph was obtained on an extract from the Datcha corpus (Damnati et al.,




Figure 5: A chatbot architecture displayed by our solution

2016). The extract contains 4,000 chats between customers and call-center agents, restricted to conversations dealing with set-top box problems. This corresponds to 95,000 speech turns. The max number of clusters obtained is 497 at the most detailed level in the hierarchy ; we have empirically chosen a level with 49 clusters. The clusters (nodes of the graph) are displayed in two different colors corresponding to the two kinds of speakers. The user has filtered the display by setting the number of transitions. The cluster "error message" is selected; on the right hand the user can consult the most relevant words for the group and the speech turns belonging to the cluster. A video is available at the url: <https://www.dailymotion.com/video/x6j6xi9>.

We used this prototype for the usecase of a bot specialized in online assistance to the users of a phone company. An other application field was also experimented: a personal home assistant (Bouraoui and Lemaire, 2017).

Once the designer has refined the architecture, he can use it to feed the dialog flow in a chatbot creation tool. The speech turns may constitute examples for the intent detection; the most representative words may correspond to entities.

6 Conclusion and future work

We presented Graph2Bots, an unsupervised assistant for dialogue designing. It is able to extract a graph representation from a corpus of conversations by using unsupervised learning, namely co-

clustering, and to allow graph editing. The graph displays the main stages of the dialogues and their transitions. Our approach is independent of the domain and the language.

In the future, we would like to improve temporality management in the co-clustering (taking speech turn indexes into account); to simplify the graph when necessary (loops and other specific sub-graphs removal); to enrich its state model (e.g. detection of non-communicative actions entwined in the dialogue); and to instantiate a chatbot automatically via an interoperable format of the graph generated.

References

- M. Boullé. 2011. Data grid models for preparation and modeling in supervised learning. In *Hands-On Pattern Recognition: Challenges in Machine Learning*, volume 1, pages 99–130. Guyon, I. and Cawley, G. and Dror, G. and Saffari, A.
- M. Boullé, D. Gay, and A. Bondu. 2013. The many faces of data grid models. <https://bit.ly/2Ef6dDb>.
- J. L. Bouraoui and V. Lemaire. 2017. Cluster-based graphs for conceiving dialog systems. In *Workshop DMNLP at European Conference on Machine Learning (ECML)*.
- Géraldine Damnati, Aleksandra Guerraz, and Delphine Charlet. 2016. Web chat conversations from contact centers: a descriptive study. In *LREC*.
- Romain Guigourès. 2013. *The Application of Co-clustering in Exploratory Data Analysis*. Ph.D. thesis, Université Panthéon-Sorbonne - Paris I.

On a Chatbot Conducting Dialogue-in-Dialogue

Boris Galitsky¹, Dmitry Ilvovsky², and Elizaveta Goncharova²

¹Oracle Inc. Redwood Shores CA

²National Research University Higher School of Economics

boris.galitsky@oracle.com; dilvovsky@hse.ru; egoncharova@hse.ru

Abstract

We demo a chatbot that delivers content in the form of virtual dialogues automatically produced from plain texts extracted and selected from documents. This virtual dialogue content is provided in the form of answers derived from the found and selected documents split into fragments, and questions are automatically generated for these answers.

1 Introduction

Presentation of knowledge in dialogue format is a popular way to communicate information effectively. It has been demonstrated in games, news, commercials, and educational entertainment. Usability studies have shown that for information acquirers dialogues often communicate information more effectively and persuade stronger than a monologue most of times (Cox et al., 1999, Craig et al., 2000).

We demo a chatbot that delivers content in the form of virtual dialogues automatically produced from plain texts extracted and selected from documents. Given an initial query, this chatbot finds documents, extracts topics from them, organizes these topics in clusters according to conflicting viewpoints, receives users clarification on which cluster is most relevant to them, and provides the content for this cluster. This content is presented in the form of a virtual dialogue where the answers are derived from the found and selected documents split into fragments, and questions are automatically generated for these answers.

A virtual dialogue is defined as a multi-turn dialogue between imaginary agents obtained as a result of content transformation. It is designed with the goal of effective information representation and is intended to look as close as possible to a genuine dialogue. Virtual dialogues as search results turn out to be more effective means of information access in comparison with original documents provided by a conventional chatbot or a search engine.

2 Dialogue Construction from Plain Text

To form a dialogue from text sharing information or explaining how to do things, we need to split it into parts which will serve as answers. Then for each answer a question needs to be formed. The cohesiveness of the resultant dialogue should be assured by the integrity of the original text; the questions are designed to “interrupt” the speaker similar to how journalists do interviews.

We employ a general mechanism of conversion of conversion a text paragraph of various styles and genres into a dialogue form. The paragraph is split into text fragments serving as a set of answers, and questions are automatically formed for some of these text fragments. The problem of building dialogue from text T is formulated as splitting it into a sequence of answers $A = [A_1 \dots A_n]$ to form a dialogue $[A_1, \langle Q_1, A_2 \rangle, \dots, \langle Q_{n-1}, A_n \rangle]$, where A_i answers Q_{i-1} and possibly previous question, and $\cup A_i = T$. Q_{i-1} needs to be derived from the whole or a part of A_i by linguistic means and generalization; also some inventiveness may be required to make these questions sound natural. To achieve it, we try to find a semantically similar phrase on the web and merge it with the candidate question.

The main foundation of our dialogue construction algorithm is Rhetorical Structure Theory (RST, Mann and Thompson, 1988). RST represents the flow of entities in text via Discourse Tree – a hierarchical structure that sets inter-relations between text fragments (Elementary Discourse Units, EDU): what elaborates on what, what explains what, what is attributed to what, what contradicts what, etc.

Rhetorical relations between the EDUs are usually binary and anti-symmetric, which defines the main unites (nucleus) and the subordinate ones (satellite). Thus, once we split a text into EDUs, we know which text fragments will serve as answers to questions: satellites of all relations. *Elaboration* rhetorical relation is default and *What*-question to a verb phrase is formed. *Background* relation yields another *What*-question for the satellite ‘...as *<predicate>*-*<subject>*’. Finally, *Attribution* relation is a basis of “What/who is source” question.

A trivial approach to question generation is simple conversion of a satellite EDU into a question. But it would make it too specific and unnatural, such as ‘*the linchpin of its strategy handled just a small fraction of the tests then sold to whom?*’. Instead, a natural dialogue should be formed with more general questions like ‘*What does its strategy handle?*’.

An example of converting a text into a virtual dialogue is shown in Figure 1. First, the text is split into EDUs. They act as answers in the virtual dialogue. The questions generated on their basis are shown in angle brackets and bolded. Each leave of the discourse tree determining an EDU starts with ‘TEXT’. Rhetorical relations (in italics) are followed by the tags ‘LeftToRight’ or ‘RightToLeft’ specifying dependency direction between the units, or which of the following unit is a nucleus and a satellite.

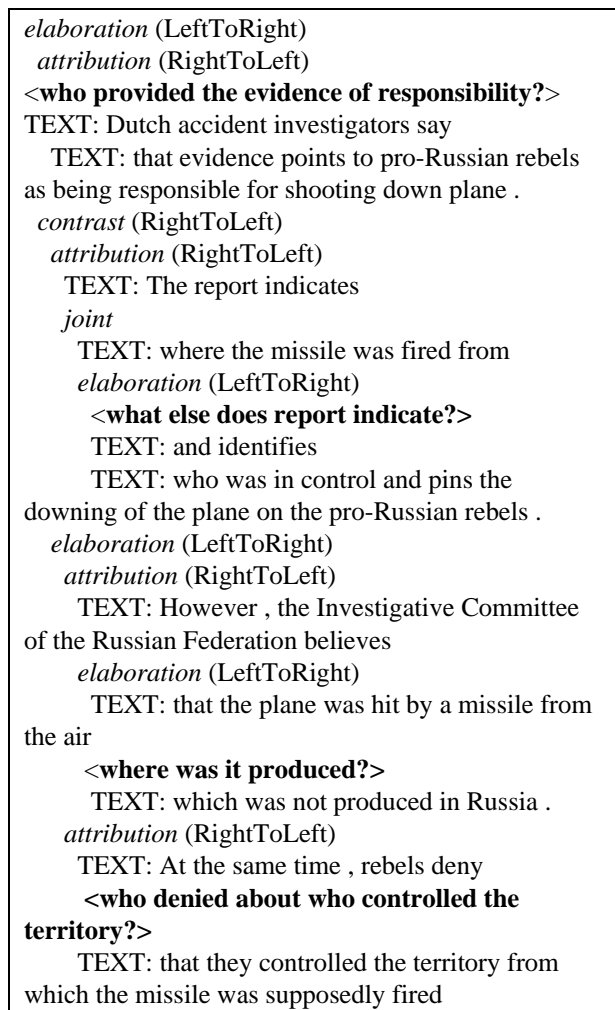


Figure 1: A discourse tree for a text paragraph with questions formulated for satellite EDUs as answers

The scheme of building a dialogue from text process is shown in Figure 2. Each paragraph of a document is converted into a dialogue via building a communicative discourse tree for it and then generating questions from its Satellite Elementary Discourse Units. Current

chatbot is development of the previously built tool that conducted task-oriented conventional dialogues (Galitsky et al., 2017).

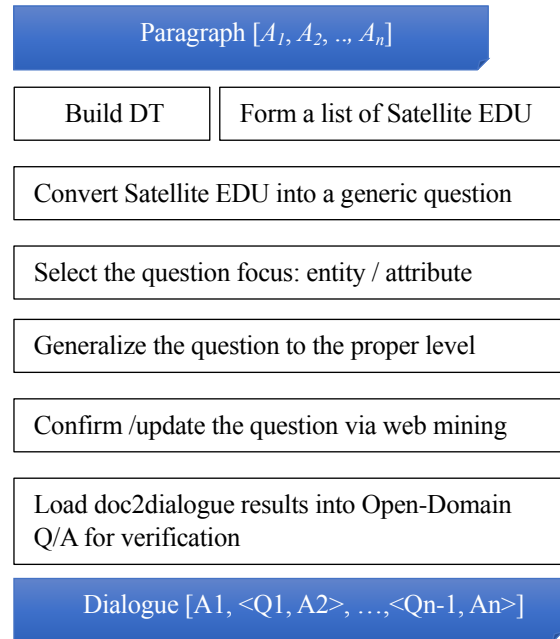


Figure 2: Scheme of dialog building process

3 Evaluation of Effectiveness

Evaluating the effectiveness of information delivery via virtual dialogues, we compare the conventional chatbot sessions where users were given plain-text answers, and the ones where users were given a content via virtual dialogues.

Table 1. Evaluation of comparative effectiveness of conventional and virtual dialogues

	Conventional dialogues			Virtual dialogues		
	# of iterations till found	# of iterations till decision	Coverage of exploration # of entities	# of iterations till found	# of iterations till decision	Coverage of exploration # of entities
Conv. only	4.6	6.3	10.8	-	-	-
Virtual only	-	-	-	4.1	6.0	13.7
Conventional followed by virtual	4.0	5.7	7.6	6.1	11.3	15.1
Virtual followed by convent.	5.6	7.1	12.3	3.7	7.0	11.5

The results on comparative usability of conventional dialogue and virtual dialogue are given in Table 1. We assess dialogues with respect to following usability properties averaged over the number of experiments:

The speed of arriving to the sought piece of information (first column). It is measured as a number

of iteration (a number of user utterances) preceding the final reply of the chatbot provided an answer wanted by the user. We measure the number of steps only if the user confirms that she accepts the answer.

The speed of arriving to decision to commit a transaction, such as purchase or reservation, or product selection (second column). A user is expected to accumulate sufficient information, and this information, such as reviews, should be convincing enough for making such decision. The less these values are the more relevant information was delivered via the dialogue.

We also measure **how many entities** (in linguistic sense) were explored during a session with the chatbot (third column). We are interested in how thorough and comprehensive the chatbot session is, how much a user actually learns from it. This assessment is sometimes opposite to the above two measures but is nevertheless important for understanding the overall usability of various conversational modes.

We do not compare precision and recall of search sessions with either dialogue mode since the same information is delivered, but in distinct modes.

In the first and second rows, we assess the stand-alone systems. One can observe that virtual dialogues take less iteration on average for information access (4.1 compared to 4.6) and a little less number of iterations for decisions than conventional dialogues do (6.0 and 6.3 respectively).

In the bottom two rows, we observe the usability of the hybrid system. Notice that the bottom row corresponds to the inverse architecture, where virtual dialog is followed by the conventional one. This scenario proceed from right to left, so, the first step's results are shown in three last columns of the table, then the values of the first three columns are calculated. When a conventional dialogue is followed by a virtual one, a lower portion of users is satisfied by the first step in comparison to the inverse architecture. Thus, the latter accounts for much less iteration required by user to be satisfied with the answer and make a final decision.

4 Sample ChatBot session

We present an exploratory session that combines information delivery in both the traditional textual answers (conventional dialogue) and a virtual dialogue form. The chatbot session is shown in Figure 3.

The dialogue starts from the user question, '*advantages and new features of 5G*'. The chatbot consults the sources (e.g. public URLs) and extracts the content from each page (or documents) expected to be relevant for the query. In this example seven URLs were processed, from domain-specific to general knowledge portals like Quora.com. Then the chatbot forms the list of topics extracted from these search results so that the user might select one of his interest.

Once the chatbot forms the topics for clarification of the user search intent, it shows them as a list. In Fig. 3 the list of topics proposed by the chatbot is underlined, the topics are numbered from 1 to 5. The user selects his topic of interest and requests a specific answer via the topic number or the topic expression ('*next stage in technology*' or '[5]'). Once the answer is read, there are multiple options (*yes/more/ ... / virtual dialogue*):

- navigate to the next answer from the chatbot list;
- navigate to a specific answer from the chatbot list;
- reject this answer and attempt to reformulate the query;
- reduce search to a specified web domain (such as quota.com, for example);
- proceed in the same direction to more search results in the form of a virtual dialogue;
- accept the answer and conclude the session.

```

Bot > ask a new question
User > advantages and new features of 5G
Bot > exploring....

  I believe these are the main topics of your query: is
that what you meant? Please select
demonstrating the benefits of the technology [1].
wide range of people from student' [2].
5g new radio nr
specification' [3]. ... next generation mobile networks
alliance [4]. the next stage in mobile technology [5].

User > next stage in technology (or [5])
Bot > Put simply, it's the next stage in mobile
technology. It follows 4G and 4G LTE and builds on
what they offer, delivering everything at a faster speed.

  Are you OK with this answer? yes / more / no /
specify [different topic] / reduce search to web domain
/ virtual dialogue
User > virtual dialogue
Bot > This is what other people are chatting about your
topic:

```

```

User1 > What Are the Benefits of Technology?
Bot1 > If you hopped over to Ericsson,
you 'd find a "5G avenue" with multiple exhibits
demonstrating the benefits of the technology ...
User1 > How are these benefits leveraged by new
tablets?
Bot1 > Samsung kicked off its press conference not
with its new tablets, but 5G is the next wireless
network in telecom industry ...
User2 > Are the features right for them?
Bot1 > It is important for people to start learning
about 5G technology now so that they can decide if
the features are right for them ...
User2 > Which are the dual 4G VoLTE sim
smartphones in India?
Bot1 > The current 4G VoLTE smartphones with dual
SIM slots that are available in the Indian markets may
support 4G on single or both SIM slots.

```

Figure 3: Chatbot session with conventional and virtual dialogue

In the example the user selects the last option and the chatbot builds a virtual dialogue. It is a conversation among imaginary people whereas the topic stays the same, matching the original query. The virtual dialog is shown in the bottom frame (Fig. 3). As long as an imaginary chatbot responds to the same person, the dialog is intended to stay cohesive; coreferences in the follow-up questions are maintained. The main dialogue can be viewed as a one in the meta-level, and the object-level dialogue is naturally embedded into the meta-level one.

Now the user can either browse the built virtual dialogue or search it to find a fragment of conversation which is relevant to the user current exploration intent. If the user types the query ‘*Are the features right for me?*’, he is directed to the virtual dialogue fragment where some other users are discussing if the technology is ‘*right for them*’. The search matches the query either against the fragments of an original text, generated questions, or both.

5 Related Work and Conclusions

(Piwek et al 2007) were pioneers of automated construction of dialogues, proposing Text2Dialogue system. The authors provided a theoretical foundation of the mapping that the system performs from RST structures to Dialogue representation structures. The authors introduced a number of requirements for a dialogue generation system (robustness, extensibility, and variation and control) and reported on the evaluation of the mapping rules.

An important body of work concerns tutorial dialogue systems. Some of the work in that area focuses on authoring tools for generating questions, hints, and prompts. Typically, these are, however, single utterances by a single interlocutor, rather than an entire conversation between two agents. Some researchers have concentrated on generating questions together with possible answers such as multiple choice test items, but this work is restricted to a very specific type of Q/A pairs (Mitkov et al 2006).

Dialogue acts are an important source which differentiates between a plain text and a dialogue. Proposed algorithm of virtual dialogues can assist with building domain-specific chatbot training datasets. Recently released dataset, DailyDialog (Li et al., 2017), is the only dataset that has utterances annotated with dialogue acts and is large enough for learning conversation models.

We proposed a novel mode of chatbot interaction via virtual dialogue. It addresses sparseness of dialogue data on the one hand and convincingness, perceived authenticity of information presented via dialogues on the other hand. We quantitatively evaluated improvement of user satisfaction with virtual dialogue in comparison to regular chatbot replies and confirmed the strong points of the former. We conclude that virtual

dialogue is an important feature related to social search to be leveraged by a chatbot.

Chatbot demo videos (please, check 10 min video) and instructions on how to use it are available at <https://github.com/bgalitsky/relevance-based-on-parse-trees> in the “**What is new?**” section.

References

- Mann, William and Sandra Thompson. 1988. *Rhetorical structure theory: Towards a functional theory of text organization*. *Text - Interdisciplinary Journal for the Study of Discourse*, 8(3):243–281.
- Joty, Shafiq R, Giuseppe Carenini, Raymond T Ng, and Yashar Mehdad. 2013. Combining intra-and multi- sentential rhetorical parsing for document-level discourse analysis. In *ACL (1)*, pages 486–496.
- Galitsky, B, Ilvovsky, D. and Kuznetsov SO. 2015. Rhetoric Map of an Answer to Compound Queries. *ACL-2*, 681–686.
- Kipper, K. Korhonen, A., Ryant, N. and Palmer, M. 2008. A large-scale classification of English verbs. *Language Resources and Evaluation Journal*, 42, pp. 21-40.
- Cox, R J. McKendree, R. Tobin, J. Lee, and T. Mayes. Vicarious learning from dialogue and discourse: A controlled comparison. *Instructional Science*, 27:431– 458, 1999.
- Craig, S, B. Gholson, M. Ventura, A. Graesser, and the Tutoring Research Group. Overhearing dialogues and monologues in virtual tutoring sessions: Effects on questioning and vicarious learning. *International Journal of Artificial Intelligence in Education*, 11:242–253, 2000.
- Yi, L., Ji Y, and Mari Ostendorf. 2016. LSTM based conversation models. arXiv preprint arXiv:1603.09457 .
- Li Y, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. Dailydialog: A manually labelled multi-turn dialogue dataset. arXiv preprint arXiv:1710.03957.
- Piwek, Paul; Hernault, Hugo; Prendinger, Helmut and Ishizuka, Mitsuru (2007). T2D: Generating Dialogues Between Virtual Agents Automatically from Text. *Lecture Notes in Artificial Intelligence*, Springer, Berlin Heidelberg, pp. 161–174.
- Mitkov R, L. A. Ha, and N. Karamanis. A computer-aided environment for generating multiple-choice test items. *Natural Language Engineering: Special Issue on using NLP for Educational Applications*, 12(2):177–194, 2006.
- Boris Galitsky and Dmitry Ilvovsky. 2017. Chatbot with a discourse structure-driven dialogue management. *EACL System Demonstrations*.

DEEPCOPY: Grounded Response Generation with Hierarchical Pointer Networks

Semih Yavuz*

University of California, Santa Barbara
syavuz@cs.ucsb.edu

Abhinav Rastogi

Google AI
abhirast@google.com

Guan-Lin Chao

Carnegie Mellon University
guanlinchao@cmu.edu

Dilek Hakkani-Tür

Amazon Alexa AI
dilek@ieee.org

Abstract

Recent advances in neural sequence-to-sequence models have led to promising results for several language generation-based tasks, including dialogue response generation, summarization, and machine translation. However, these models are known to have several problems, especially in the context of chit-chat based dialogue systems: they tend to generate short and dull responses that are often too generic. Furthermore, these models do not ground conversational responses on knowledge and facts, resulting in turns that are not accurate, informative and engaging for the users. In this paper, we propose and experiment with a series of response generation models that aim to serve in the general scenario where in addition to the dialogue context, relevant unstructured external knowledge in the form of text is also assumed to be available for models to harness. Our proposed approach extends pointer-generator networks (See et al., 2017) by allowing the decoder to hierarchically attend and copy from external knowledge in addition to the dialogue context. We empirically show the effectiveness of the proposed model compared to several baselines including (Ghazvininejad et al., 2018; Zhang et al., 2018) through both automatic evaluation metrics and human evaluation on CONVA12 dataset.

1 Introduction

Recently, deep neural networks have achieved state-of-the-art results in various tasks including computer vision, natural language and speech processing. Specifically, neural sequence-to-sequence models (Sutskever et al., 2014; Bahdanau et al., 2015) have led to great progress in important downstream NLP tasks like text summarization (Rush et al., 2015; Nallapati et al., 2016; See et al., 2017;

Tan et al., 2017; Yavuz et al., 2018), machine translation (Cho et al., 2014; Sutskever et al., 2014; Luong et al., 2015; Bahdanau et al., 2015), and reading comprehension (Xiong et al., 2017). However, achieving satisfactory performance on dialogue still remains an open problem. This is because dialogues can have multiple valid responses with varying semantic content. This is vastly different from the aforementioned tasks, where the generation is more conveniently and uniquely constrained by the input source.

Although neural models appear to generate meaningful responses when trained with sufficiently large datasets in the chit-chat setting, such generic chit-chat models reveal several weaknesses that were reported by previous research (Serban et al., 2016; Vinyals and Le, 2015). Most common problems include inconsistency in personality, dull and generic responses, and unawareness of long-term dialogue context. To alleviate these limitations, we turn our focus on a different problem setting for dialogue response generation where the model is provided a set of relevant textual facts (speaker persona descriptions) and is allowed to harness this knowledge when generating responses in a multi-turn dialogue. To handle the personality inconsistency issue, we ground our dialogue generation model on external knowledge facts which are a list of persona descriptions in our application (Li et al., 2016a; Zhang et al., 2018). We explicitly use the dialogue history as memory for the model to condition on which potentially encourages a more natural conversation flow. Towards encouraging generation of more specific and appropriate responses while avoiding generic and dull ones, we use a hierarchical pointer network in our model such that it can copy content from two sources: current dialogue history and persona descriptions.

In this work, we propose a novel and general ar-

*Work done while interning at Google AI.

chitecture DEEPCOPY that extends the attentional sequence-to-sequence model with a hierarchical pointer network that enables the decoder to jointly attend and copy tokens from any of the facts available as external knowledge in addition to the dialogue context (encoder input). This is achieved entirely in an end-to-end fashion through factoring the whole copy mechanism into following three hierarchies/components: (i) a token-level attention mechanism over the dialogue context to determine the probability of copying a token from the dialogue context, (ii) A hierarchical pointer network to determine the probability of copying a token from each fact, and (iii) An inter-source meta attention over the input sources *dialogue context* and *external knowledge*, which combines the two copying probabilities. Using these components, a single copying probability distribution over the unique tokens appearing in the model input is computed exploiting the well-defined hierarchy among them. In addition, the model is equipped with a soft switch mechanism between *copying* and *generation* modes similar to (See et al., 2017), which allows us to softly combine the *copying probabilities* with the decoder’s *generation probabilities* over a fixed vocabulary into a final output probability distribution over an extended vocabulary. We empirically show the effectiveness of the proposed DEEPCOPY model compared to several baselines including (Ghazvininejad et al., 2018; Zhang et al., 2018) on CONVAI2 challenge.

2 Related Work

Earlier work on data-driven, end-to-end approaches to conversational response generation treated the task as statistical machine translation, where the goal is to generate a response given the previous dialogue turn (Ritter et al., 2011; Vinyals and Le, 2015). While these studies resulted in a paradigm change compared to earlier work, they do not include mechanisms to represent conversation context. To tackle this problem and have a better representation of conversation context as input to generation, (Serban et al., 2016) proposed hierarchical recurrent encoder-decoder (HRED) networks. HRED combines two RNNs, one at the token level, modeling individual turns, and one at the dialogue level, inputting turn representations from the token-level RNNs. However, utterances generated by such neural response generation systems are often generic and contentless (Vinyals and Le, 2015). To improve the diversity and content of generated re-

sponses, HRED was later extended with a latent variable that aims to model the higher level aspects (such as topic) of the generated responses, resulting in the VHRED approach (Serban et al., 2017).

Another challenge for dialogue response generation is the integration of knowledge into the generated responses. (Liu et al., 2018) extracted facts relevant to a dialogue from knowledge using string matching, named entity recognition and linking, found additional entities from knowledge that are most relevant to the facts by a neural similarity scorer, and used these as input context features for the dialogue generation RNN. (Ghazvininejad et al., 2018) used end-to-end memory networks to base the generated responses on knowledge, where an attention over the knowledge relevant to the conversation context is estimated, and multiple knowledge representations are included as input during the decoding of responses. In this work, we use end-to-end memory networks as a baseline.

Although much research has focused on response generation in a chit-chat setting, models trained on large datasets of human-human interactions of diverse speaker characteristics often tend to generate responses which are too vague and generic (common for most speakers) or inconsistent in personality (switching between different speakers’ characteristics). Recently, (Zhang et al., 2018) presented the CONVAI2 challenge containing persona descriptions and over 10K real human chit-chats where speakers were required to converse based on their assigned persona. (Li et al., 2016a) learned speaker persona embeddings from a single-speaker setting (e.g. Twitter posts) or a speaker-address style (human-human conversations) to generate personalized responses given a single utterance input. Another related work (Raghu et al., 2018) applies hierarchical memory network for task oriented dialog problem. In this work, we compare our model with (Zhang et al., 2018) which uses a memory-augmented sequence-to-sequence response generator grounded on the dialogue history and persona.

3 Model

In this section, we first set up the problem, and then briefly revisit the baseline models using memory networks (Sukhbaatar et al., 2014) and pointer-generator networks (See et al., 2017). Subsequently, we introduce the proposed DEEPCOPY model with a hierarchical pointer network and our training process.

3.1 Problem Setup

Let $\mathbf{x} = (x_1, x_2, \dots, x_n)$ denote the tokens in the dialogue history. The dialogue is accompanied by a set of K relevant supporting facts, where $\mathbf{f}^{(i)} = (f_1^{(i)}, f_2^{(i)}, \dots, f_{n_i}^{(i)})$ is the list of tokens in the i -th fact. Our goal is to generate the response as a sequence of tokens $\mathbf{y} = (y_1, y_2, \dots, y_m)$ using the dialogue history and supporting facts. Note here that we are not interested in retrieval/ranking based models (Weston et al., 2018) which rely on a set of candidate responses. Generative models are essential for this problem because we want to incorporate content from new facts during inference which may not be present in the training set. Hence, using a predefined set of candidates may not ensure high coverage.

3.2 Baseline Models

In this section, we describe several baseline response generation models including the ones from existing work (Ghazvininejad et al., 2018; Zhang et al., 2018) and the in-house ones we propose as additional baselines.

3.2.1 Seq2Seq

In a sequence-to-sequence model with attention (Bahdanau et al., 2015), a sequence of input tokens is encoded using an LSTM encoder. At decoder step t , the decoder state h_t , a context vector c_t and the previous decoder output y_{t-1} are used together to output a distribution over a fixed vocabulary of tokens obtained from the training set using a non-linear function. The context vector c_t is an attention-weighted combination of the encoder outputs. In the following baseline models, we use different features as inputs to the encoder. The underlying model remains the same.

SEQ2SEQ + NOFACT. Only the dialogue context tokens \mathbf{x} are used as input to the encoder.

SEQ2SEQ + BESTFACTCONTEXT. We select the fact $\mathbf{f}^{(c)}$ whose tokens have highest unigram *tf-idf* similarity to the dialogue context tokens. $[\mathbf{x}||\mathbf{f}^{(c)}]$ is then used as input to the encoder, where $||$ denotes concatenation.

SEQ2SEQ + BESTFACTRESPONSE. We select the fact $\mathbf{f}^{(r)}$ whose tokens have highest unigram *tf-idf* similarity to the ground truth response. $[\mathbf{x}||\mathbf{f}^{(r)}]$ is used as input to the encoder. The aim of this experiment is to have a better understanding of the effect of fact selection on response generation, since using the ground truth for fact selection is not fair.

3.2.2 Memory Network

Our variations of Seq2Seq models described in Section 3.2.1 incorporate facts by concatenating them to the dialogue context. Memory networks (Ghazvininejad et al., 2018; Zhang et al., 2018) are a more principled approach to incorporating external facts. Similar to (Ghazvininejad et al., 2018), we use a context encoder to embed the context tokens \mathbf{x} and obtain a list of outputs and final hidden state $u \in \mathbb{R}^d$. As outlined in (Ghazvininejad et al., 2018), a fact $\mathbf{f}^{(i)}$ is embedded into key and value vectors k_i and m_i , respectively. A summary $o \in \mathbb{R}^d$ of facts is then computed as an attention weighted combination of (m_1, m_2, \dots, m_K) by conditioning on u and (k_1, k_2, \dots, k_K) . We then combine the two summaries into $\hat{u} = u + o$, and use it to initialize the decoder state. We report results on the following variants:

MEMNET. This is equivalent to the model used in (Ghazvininejad et al., 2018), described above. This is essentially a sequence to sequence model without attention at every decoder step, except using the combined summary \hat{u} to initialize the decoder.

MEMNET+CONTEXTATTENTION. At each decoder step, the decoder state attends over the encoder outputs and obtains a context vector $c_t^{(c)}$. This is equivalent to SEQ2SEQ + NOFACT model from Section 3.2.1, except using the fact summary \hat{u} to initialize the decoder state.

MEMNET+FACTATTENTION. At each decoder step, we use the decoder state to attend over the value embeddings (m_1, m_2, \dots, m_K) corresponding to facts, and obtain a context vector $c_t^{(f)}$. This model is similar to the *generative profile memory network* (Zhang et al., 2018), where we apply attention only on facts, and we set the decoder’s initial state to the combined summary \hat{u} .

MEMNET+FULLATTENTION. This model employs attention over both facts and dialogue context at each decoder step. The two attention modules are combined by concatenating $c_t^{(c)}$ and $c_t^{(f)}$ (Zoph and Knight, 2016).

3.2.3 Seq2Seq with Copy Mechanism

Seq2seq models can only generate tokens present in a fixed vocabulary obtained from the training set. Pointer-generator network (See et al., 2017) extends the attentional sequence-to-sequence model (Bahdanau et al., 2015) by employing a pointer network (Vinyals et al., 2015). It has two decoding modes, copying and generating, which are com-




Figure 1: Overview of our proposed approach as described in Section 3.3. The decoder state d_t is used to attend over dialogue context and knowledge source to generate distributions for copying tokens from these sources. The decoder outputs a distribution over a fixed vocabulary. The three distributions are combined to yield the final distribution over tokens at each step t .

binning via a soft switch mechanism, allowing it to copy tokens from source in addition to generating from vocabulary. We report the results for the following additional baselines obtained by equipping the corresponding Seq2Seq model in Section 3.2.1 with copy mechanism: SEQ2SEQ + NOFACT + COPY, SEQ2SEQ + BESTFACTCONTEXT + COPY, SEQ2SEQ + BESTFACTRESPONSE + COPY.

3.3 DeepCopy with Hierarchical Pointer Networks

Pointer-generator network (See et al., 2017) can only copy tokens from the encoder input. In this section, we present our proposed DEEPCOPY model that extends pointer-generator network (See et al., 2017) using a novel hierarchical pointer network. Our model allows copying tokens from multiple input sources (facts $\mathbf{f}^{(i)}$, $1 \leq i \leq K$), besides the encoder input (dialogue context x).

A high-level overview of the proposed approach is illustrated in Figure 1. At decoder step t , the decoder state h_t is used to attend over the dialogue context tokens and fact tokens to give a distribution over the tokens present in context and facts respectively. These distributions are then combined with the distribution output by the decoder over the fixed vocabulary to obtain the overall distribution.

Encoding a sequence. Let $\mathbf{w} = (w_1, w_2, \dots, w_n)$ be a sequence of tokens. We first obtain a trainable embedded representation of each token in the sequence and then use a LSTM cell to encode the sequence of embedding vectors. We define $e, \mathbf{s} = \text{Encode}(\mathbf{w})$, where e denotes the final state of the LSTM and $\mathbf{s} = (s_1, s_2, \dots, s_n)$ denotes the outputs of the LSTM cell at all steps.

Attention. Let $\mathbf{u} = (u_1, u_2, \dots, u_n)$ be a sequence of vectors where $u_i \in \mathbb{R}^p$, $1 \leq i \leq n$ and $v \in \mathbb{R}^q$ be a conditioning vector. The attention module generates a linear combination c of elements in \mathbf{u} by conditioning them on v as defined by the equations below. We define $\alpha, c = \text{Attention}(\mathbf{u}, v)$, where $\alpha_i \in \mathbb{R}^n$ is the weight assigned to u_i , and $c \in \mathbb{R}^p$ is a vector representation of the sequence \mathbf{u} conditioned on v . In the equations below, w_1 and W_2 are parameters of appropriate dimension. In our setup, we use $p = q$, $w_1 \in \mathbb{R}^p$, and $W_2 \in \mathbb{R}^{p \times 2p}$.

$$e_i = w_1^T \tanh(W_2[u_i; v]) \quad (1)$$

$$\alpha_i = \frac{\exp(e_i)}{\sum_{j=1}^n \exp(e_j)} \quad (2)$$

$$c = \sum_{i=1}^n \alpha_i u_i \quad (3)$$

Copying from Dialogue Context. Similar to our baseline models, we encode the dialogue context tokens \mathbf{x} (Equation 4) and apply attention to the encoder outputs at a decoder step t (Equation 5). This outputs attention weights $\alpha_t^{(x)}$ and a representation of the entire context $c_t^{(x)}$. The attention weights are aggregated to obtain the distribution over context tokens $p_t^{(x)}(w)$ (Equation 6),

$$e^{(x)}, \mathbf{s}^{(x)} = \text{Encode}(\mathbf{x}) \quad (4)$$

$$\alpha_t^{(x)}, c_t^{(x)} = \text{Attention}(\mathbf{s}^{(x)}, h_t) \quad (5)$$

$$p_t^{(x)}(w) = \sum_{\{i: x_i=w\}} \alpha_{t,i}^{(x)} \quad (6)$$

Copying from Facts: Hierarchical Pointer Network. We introduce the hierarchical pointer net-

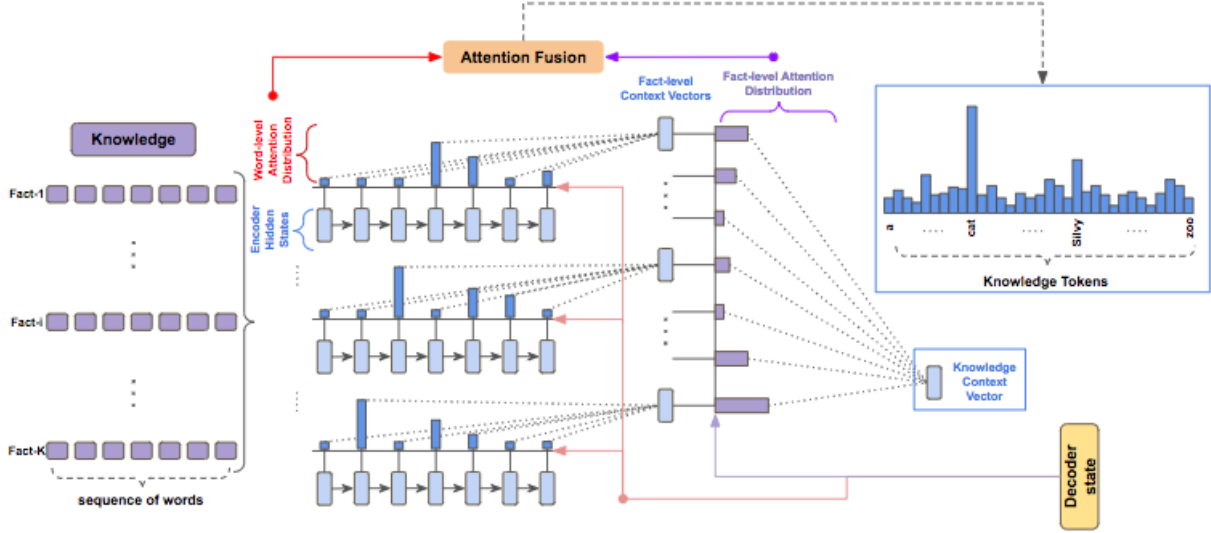


Figure 2: Illustration of hierarchical pointer network. The decoder state d_t is used to attend over tokens for each fact and also over the fact-level context vectors obtained by weighted average of token-level representations (w.r.t token-level attention weights) for each fact. The token-level attention weights are then combined with the attention distribution over facts (Equation 11) to generate the probability of copying each token in all the facts.

work (Figure 2) as a general methodology for enabling token-level copy mechanism from multiple input sequences or facts. Each fact $\mathbf{f}^{(i)}$ is encoded (Equation 7) to obtain token level representations $\mathbf{s}^{(f)(i)}$ and overall representation $e^{(f)(i)}$. The decoder state h_t is used to attend over token level representations (Equation 8) and the overall fact-level representations of each fact (Equation 9) by

$$e^{(f)(i)}, \mathbf{s}^{(f)(i)} = \text{Encode}(\mathbf{f}^{(i)}) \quad (7)$$

$$\alpha_t^{(f)(i)}, c_t^{(f)(i)} = \text{Attention}(\mathbf{s}^{(f)(i)}, h_t) \quad (8)$$

$$\beta_t, c_t^{(f)} = \text{Attention}(\{c_t^{(f)(i)}\}_{i=1}^K, h_t) \quad (9)$$

to compute the probability of copying a word w from facts as

$$\begin{aligned} p_t^{(f)}(w) &= \sum_{j=1}^K p_t^{(f)}(\mathbf{f}^{(j)}) \cdot p_t^{(f)}(w|\mathbf{f}^{(j)}) \\ &= \sum_{j=1}^K \beta_{t,j} \sum_{\{l: f_t^{(j)}=w\}} \alpha_{t,l}^{(f)(j)} \end{aligned} \quad (10)$$

Inter-Source Attention Fusion We now present the mechanism to fuse the two distributions $p_t^{(x)}(w)$ and $p_t^{(f)}(w)$ representing the probabilities of copying tokens from dialogue context and facts respectively. We use the decoder state h_t to attend over dialogue context representation $c_t^{(x)}$ and overall fact representation $c_t^{(f)}$ (Equation 11). The resulting attention weight $\gamma'_t = [\gamma_t, 1 - \gamma_t]$ is used to combine the two copying distributions as shown in

Equation 12.

$$\gamma_t, c_t = \text{Attention}([c_t^{(x)}, c_t^{(f)}], h_t) \quad (11)$$

$$p_t^{\text{copy}}(w) = \gamma_t p_t^{(x)}(w) + (1 - \gamma_t) p_t^{(f)}(w) \quad (12)$$

Similar to Seq2Seq models, the decoder also outputs a distribution p_t^{vocab} over the fixed training vocabulary at each decoder step using the overall context vector c_t and decoder state h_t . Having defined the copy probabilities p_t^{copy} for tokens that appear in the model input, either the dialogue context or the facts in external knowledge source, we combine p_t^{vocab} and p_t^{copy} using the mechanism outlined in (See et al., 2017), except we use c_t defined in Equation 11 as the context vector instead.

To better isolate the effect of copying, a key component of the proposed DEEPCOPY model, we also conduct experiments with MULTISEQ2SEQ model that incorporates the knowledge facts in the same way (by encoding each fact separately with LSTM, and attending on each by the decoder as in (Zoph and Knight, 2016)), but relies completely on *generation probabilities* without a copy mechanism.

3.4 Training

We train all the models described in this section using the same loss function optimization. More precisely, given a model M that produces a probability $p_t(w|y_{<t})$ of generating token w at decoding step t , we train the whole network end-to-end with

the negative log-likelihood loss function of

$$J_{\text{loss}}(\Theta) = -\frac{1}{|\mathbf{y}|} \sum_{t=1}^{|\mathbf{y}|} \log(p_t(y_t | y_{<t}, \mathbf{x}, \{\mathbf{f}^{(i)}\}_{i=1}^K))$$

for a training sample $(\mathbf{x}, \mathbf{y}, \{\mathbf{f}^{(i)}\}_{i=1}^K)$ where Θ denotes all the learnable model parameters.

4 Experiments

In this section, we describe the details of dataset, training process, evaluation metrics, and the performance results of DEEPCOPY model in comparison to proposed and existing baselines.

4.1 Dataset

We perform experiments for our problem setup on the recently released CONVA12 *conversational AI challenge* dataset, which is an extended version of PERSONACHAT (Zhang et al., 2018). The conversations in CONVA12 are obtained by asking a pair of crowdworkers to chat with each other naturally based on their randomly assigned personas (from a set of 1155 personas) towards getting to know each other. Personas are created by a different set of crowdworkers, and they consist of ~ 5 natural language sentences, each describing an aspect of a person that can range from common hobbies like "I like to play basketball" to very specific facts like "I have a pet parrot named Tasha", reflecting a wide range of different personalities. The dataset contains ~ 11000 dialogues with ~ 160000 utterances, and 2000 dialogues with non-overlapping personas are used for validation and test. For our setting, we use personas as external knowledge sources that models can ground on while generating responses.

4.2 Training and Implementation Details

In all the models explored in this paper, we set the dialogue context to concatenation of the last two dialogue turns separated by a special CONCAT token. The models are supplied with the persona facts of the side generating the response at the current turn, while the persona of the other side is concealed. We use a vocabulary of 18650 most frequent tokens and all the remaining tokens are replaced with a special UNK token. Embeddings of size 100 are randomly initialized and updated during training. We set the size of LSTM hidden layer to 100 for both encoder and decoder. The encoder and decoder vocabularies and embeddings are shared. A shared LSTM encoder is used for encoding both dialogue context and facts of external knowledge source. The model parameters are optimized using

Adam (Kingma and Ba, 2015) with a batch size of 32, a fixed learning rate of 0.001. We apply gradient clipping to 5 when its norm exceeds this value. During inference, we generate responses by employing a beam search of width 4. Our models are implemented in *TensorFlow* (Abadi et al., 2016).

4.3 Main Results

In this section, we present the experimental results in terms of both automatic measures and human evaluation.

4.3.1 Automatic Evaluation

In Table 1, we present our results in comparison with the existing and proposed baseline models. We report the performance of each model across several metrics commonly used for evaluation of text generation models including perplexity, corpus BLEU (Papineni et al., 2002), ROUGE-L (Lin and Och, 2004), CIDEr (Vedantam et al., 2014).

As expected, SEQ2SEQ + BESTFACTRESPONSE model and its +COPY version outperform all the other models across all the evaluation metrics. This model pinpoints the importance of selecting the most suitable fact in the persona for the response to be generated at each turn, justifying our underlying motivation for conducting this experiment as highlighted in Section 3.2.1. However, the most suitable fact for the response is not available in the real application scenario, where the models are responsible for picking the useful pieces of information pertaining to the current dialogue turn to generate meaningful responses. Our proposed SEQ2SEQ + BESTFACTCONTEXT model and its +COPY version, on the other hand, are valid baselines for this scenario where the best fact is selected completely based on the dialogue context without relying on the ground-truth response. This model outperforms the previously proposed memory network based model MEMNET (Ghazvininejad et al., 2018) for knowledge grounded response generation on all the evaluation metrics, demonstrating its effectiveness despite the fact that it does not have access to all the facts unlike (Ghazvininejad et al., 2018). However, this approach has the following potential weaknesses: (i) if the best persona fact selected w.r.t dialogue context is wrong (irrelevant) for the ground-truth response, the generated response might be drastically misinforming, and furthermore it is difficult for model to recover from this error because it has no access to other facts, (ii) selecting the best fact w.r.t dialogue context based

Model	Perplexity	BLEU	ROUGE-L	CIDEr	Appropriateness
[M-1] MEMNET	61.30	3.07	59.10	10.52	3.14 (0.51)
[M-2] MEMNET + CONTEXTATTENTION	57.37	3.24	59.20	11.79	3.41 (0.54)
[M-3] MEMNET + FACTATTENTION	61.50	2.43	59.34	9.65	1.45 (0.25)
[M-4] MEMNET + FULLATTENTION	59.64	3.26	59.18	12.25	3.20 (0.49)
[S2S-1] SEQ2SEQ + NOFACT	60.48	3.38	59.46	11.41	3.12 (0.52)
[S2S-2] SEQ2SEQ + BESTFACTCONTEXT	58.68	3.35	59.13	10.77	3.08 (0.45)
[S2S-3] SEQ2SEQ + BESTFACTRESPONSE*	49.74	4.02	60.04	16.15	2.97 (0.51)
[S2SC-1] SEQ2SEQ + NOFACT + COPY	58.84	3.25	59.18	11.15	3.64 (0.54)
[S2SC-2] SEQ2SEQ + BESTFACTCONTEXT + COPY	60.25	3.17	59.46	11.17	3.60 (0.51)
[S2SC-3] SEQ2SEQ + BESTFACTRESPONSE + COPY*	38.60	4.54	60.96	21.47	3.83 (0.46)
[M-S2S] MULTISEQ2SEQ (no COPY)	57.94	2.88	59.10	10.92	3.32 (0.44)
DEEPCOPY [†]	54.58	4.09	60.30	15.76	3.67 (0.59)
G.TRUTH	N/A	N/A	N/A	N/A	4.40 (0.45)

Table 1: Main results on CONVA12 dataset. Evaluation metrics on last three columns are better the higher. Perplexity is lower the better. The results of the proposed approach are presented in bold. * indicates that the corresponding model should be considered as a kind of **ORACLE** because it has access to the fact that is most relevant to the ground-truth response during the inference/test time as defined in Section 3.2.1. † indicates that the improvement of DEEPCOPY in automatic evaluation metrics over each of the other models (except S2SC-3) is statistically significant with p-value of less than 0.001 on the paired t-test.

on *tf-idf* similarity may result in poor fact selection when the lexical overlap between context and response is small which might be a common case especially for the CONVA12 dataset as the focus of conversation may often change swiftly across the dialogue turns. The latter might be the reason why copying does not help much for this model since it might end up copying irrelevant tokens in the scenario mentioned above.

Our proposed DEEPCOPY model is designed to effectively address the aforementioned issues, where it has access to the entire set of persona facts per dialogue from which it is expected to include the useful pieces of information in the response. DEEPCOPY model outperforms all the models reported in Table 1 except for SEQ2SEQ + BESTCONTEXTRESPONSE models, which we already deem as kind of an upper bound because it has access to the most relevant fact to the response. This justifies the effectiveness of DEEPCOPY model compared to the existing works (Ghazvininejad et al., 2018; Zhang et al., 2018) and the additional baselines we explored in this work. On the other hand, MULTISEQ2SEQ performs considerably worse than the DEEPCOPY model despite the fact they both have access to the entire set of facts and employ the same encoder-decoder architecture except for the copy mechanism. This further justifies the effectiveness of incorporating the proposed hierarchical pointer networks in DEEPCOPY because integrating the external knowledge simply by employing multi-source attention as in (Zoph and Knight, 2016) does not yield to a good solution

with competitive results, performing even worse than SEQ2SEQ + NOFACT on 3 of the metrics.

4.3.2 Human Evaluation

Although automatic metrics provide tangible information regarding the performance of the models, we augment them with human evaluations for a more comprehensive analysis of the resulting model generated responses. Towards this end, we randomly sample 100 examples from test data and ask human raters to evaluate the candidate model generated responses in terms of appropriateness. Each example is rated by 3 raters, who are shown a dialog history along with a set of persona facts (of the person in turn), and asked to rate each response based on its *appropriateness* in the dialogue context with a score from 1 (worst) to 5 (best).

In Table 1, we present the results of human evaluation under the *appropriateness* column. Since each response is rated by 3 different human raters, we report the average rating along with the standard deviation in parenthesis. We observe that DEEPCOPY outperforms both the existing memory-network baselines and the proposed sequence-to-sequence baselines on the appropriateness evaluation. It also achieves a performance that is close to the *oracle* model (S2SC-3), which has a leverage of having an access to the fact that is most relevant to the ground-truth response during the inference time. Overall, human evaluation of the responses in terms of appropriateness further justifies the promise and effectiveness of our proposed DEEPCOPY model.

Model	Diversity	Fact-Inclusion			Agreement
	Distinct-2 / 3 / 4	F.Inc	F.Per	F.Hal	F.Inc / F.Per
M-1	.004 / .006 / .010	0.41	0.01	0.40	0.99 / 0.99
M-2	.010 / .019 / .031	0.43	0.01	0.42	0.97 / 0.99
M-3	.001 / .001 / .002	0.06	0.04	0.02	0.99 / 0.99
M-4	.054 / .010 / .156	0.51	0.09	0.42	0.98 / 0.98
S2S-1	.012 / .022 / .036	N/A	N/A	N/A	N/A / N/A
S2S-2	.012 / .022 / .035	0.54	0.04	0.50	0.97 / 0.99
S2S-3	.026 / .043 / .061	0.79	0.16	0.63	0.97 / 0.97
S2SC-1	.039 / .069 / .104	N/A	N/A	N/A	N/A / N/A
S2SC-2	.035 / .067 / .109	0.73	0.36	0.37	0.99 / 0.99
S2SC-3*	.058 / .111 / .178	0.73	0.55	0.18	0.98 / 0.96
M-S2S	.035 / .065 / .104	0.47	0.05	0.42	0.96 / 0.98
DEEPCOPY	.059 / .121 / .201	0.62	0.23	0.39	0.95 / 0.97
G.TRUTH	0.35 / 0.66 / 0.84	0.76	0.49	0.27	0.93 / 0.96

Table 2: Lexical diversity and fact inclusion analysis results. Model names are abbreviated according to Table 1. **F.Inc** denotes the ratio of responses that include factual information. **F.Per** and **F.Hal** denote the ratio of responses where the included fact is consistent with the persona or a hallucinated one, respectively. **Agreement** column corresponds to Cohen’s κ statistic measuring inter-rater agreement on binary factual evaluation metrics for **F.Inc** and **F.Per**. * indicates the **ORACLE** model.

Appropriateness scores also demonstrate the advantage of incorporating the soft copy mechanism. Comparing S2S (and M-S2S) models to their copy-equipped counterparts (S2SC) (and DEEPCOPY) in Table 1 immediately reveals a significant gain in appropriateness score. Another significant observation to note here is that ground-truth responses obtain an average appropriateness score of 4.4/5, which reflects both the noise in CONVAI2 dataset and the difficulty of generating the perfect response even for humans.

4.4 Further Analysis and Discussion

Lexical Diversity Analysis. In Table 2, we report the lexical diversity results using the distinctness metric introduced in (Li et al., 2016b). *distinct-n* score corresponds to the number of distinct n -grams divided by total number of generated n -grams. We can clearly observe that DEEPCOPY generates the most diverse responses among all the models including the copy-augmented oracle model (S2SC-3). Hence, diversity results further show that our proposed model is promising in addressing the most commonly observed *generic response problem* more effectively than existing models by generating more diverse responses.

Fact Inclusion Analysis. We also conduct an analysis on the kinds of factual information included in the model-generated responses. More precisely, our goal is to understand how often the generated

response includes a factual information (F.Inc), and whether this information is consistent with the persona facts (F.Per) or a hallucinated one (F.Hal). A good model can naturally include available facts from the persona and hallucinate others when the conversation context requires them. Towards this end, we ask 3 human raters to label responses with 1 (or 0) based on whether a fact is included, and if so, whether this fact is a persona-fact or not.

In Table 2, we present an analysis for the kinds of factual information included in model generated responses. As can be seen from this analysis, models that have a copy mechanism include more facts from the persona than the ones that do not. Another important observation is that the ground-truth responses include facts from persona only in 49% of the times, which indicates that the provided persona facts remain insufficient to cover the complexity of the high entropy open-ended person-to-person conversations.

In Table 2, we present Cohen’s κ score for each model and fact analysis metric pair using the scores from 3 raters for each example. We observe for each model and metric pair a κ statistic of greater than 0.9, which indicates a near perfect agreement among raters. Note that the ratio of hallucinated facts (**F.Hal**) is derived directly from human labels for fact inclusion (**F.Inc**) and persona-fact (**F.Per**). That is why, there is no separate labelling process

for hallucinated facts (**F.Hal**). Hence, there is no κ statistic for **F.Hal** in Table 2.

Error Analysis. A deeper analysis of the examples where DEEPCOPY is assigned a worse appropriateness score than the best performing memory-network based baselines (M-2 and M-4) reveals the following further insights: (i) Some of these examples are corresponding to the cases where a generic response (e.g., "I've a dog named radar", one of the frequent generic responses, completely independent of persona facts) is rated much higher (5 to 1) than factual but slightly off (by a single word in this example) responses (e.g., "I have a dog for a living." coming from the persona fact "I walk dogs for a living."), (ii) In another subset of the analyzed examples, DEEPCOPY model generates a response (e.g., "yes, but I want to become a lawyer.") by incorporating a fact that has already been used in the previous turn of the dialog whereas M-2 produces a generic response (e.g., "that's great. do you have any hobbies?", again irrelevant to facts) which is rated higher. (iii) And most of the remaining cases fall into the class of examples where incorporating knowledge facts breaks the conversation flow, which is a crucial observation specific to this dataset that can also be supported by the low persona-fact inclusion ratio (49%) of ground-truth responses.

4.5 Qualitative Observations

In Figure 3, we present an example dialogue where DEEPCOPY model generates a meaningful and fluent response by effectively mixing *copy* and *generate* modes. We can observe that it is able to attend on the right persona fact by taking the dialogue context (especially the question at the end of PERSON2's turn) into consideration. Furthermore, attending to the tokens of this fact, it produces a fluent and valid answer to yes/no question by generating "yes" and copying the rest (and most) of the tokens from the fact. Although it copies most of the tokens from the fact, it is good to observe that it copies exactly the relevant pieces instead of just copying the entire fact. SEQ2SEQ + BESTFACTRESPONSE + COPY model's response is also meaningful and fluent although it may not be as engaging for the continuation of dialog. However, the quality of the response by SEQ2SEQ + BESTFACTRESPONSE quickly degrades compared to its +COPY version. Although the response is still fluent and relevant to the dialogue context, it becomes rather irrelevant to the persona as the model seems to have

Persona Facts	
	1. i'm a clean eater.
	2. i'm a cancer survivor.
	3. my parents were both very athletic.
	4. i got a new job just yesterday to be a life coach.
	5. i love running and preparing for marathons.

Model	Conversation/Response
PERSON1	i really miss it but i have been eating healthy ever since i overcame cancer
PERSON2	omg i am glad you did , do you work now ?
MEMNET*	yes, i do not work, but i do not work.
SEQ2SEQ*	no i ' ve a job at a restaurant
SEQ2SEQ**	i'm a life coach
DEEPCOPY	yes, i just got a new job

Figure 3: Example dialogue where the previous two turns from PERSON1 and PERSON2 along with the responses generated by the models acting as PERSON1 are shown on the right. Persona facts for PERSON1 are provided on the left, among which the one in bold is the best fact w.r.t response. MEMNET*, SEQ2SEQ*, SEQ2SEQ** are abbreviations for MEMNET + FULLATTENTION, SEQ2SEQ + BESTFACTRESPONSE, SEQ2SEQ + BESTFACTRESPONSE + COPY models, respectively.

difficulty of picking the useful information from even the best persona fact it is provided with when the copy mechanism is disabled. Lastly, the response generated by MEMNET+FULLATTENTION model seems to still suffer from repetition, semantic consistency, and relevancy problems that were observed and reported by previous work.

5 Conclusion and Future Work

We propose a hierarchical pointer network for knowledge grounded dialogue response generation. Our approach extends the pointer-generator network to enable the decoder to simultaneously copy tokens from the available set of relevant external knowledge in addition to dialogue context. We demonstrate the effectiveness of our approach through various automatic and human evaluations in comparison with several baselines on the CONVAI2 dataset. Furthermore, we conduct diversity, fact inclusion, and error analysis providing further insights into model behaviors. In the future, we plan to apply our model to datasets of the same fashion where the dialogue is accompanied by a much larger set of knowledge facts (e.g., Wikipedia articles) (Galley et al., 2018). This could be done by adding a retrieval component which identifies a few contextually relevant facts (Ghazvininejad et al., 2018) to be used as input to DEEPCOPY.

References

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations (ICLR)*.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Michel Galley, Chris Brockett, Xiang Gao, Bill Dolan, and Jianfeng Gao. 2018. End-to-end conversation modeling: Moving beyond chitchat. http://workshop.colips.org/dstc7/proposals/DSTC7-MSR_end2end.pdf. Online; accessed 23 October 2018.
- Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih, and Michel Galley. 2018. A knowledge-grounded neural conversation model. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios Spathourakis, Jianfeng Gao, and Bill Dolan. 2016a. [A persona-based neural conversation model](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 994–1003. Association for Computational Linguistics.
- Jiwei Li, Michel Galley, Jianfeng Brockett, Chris ad Gao, and Bill Dolan. 2016b. A diversity-promoting objective function for neural conversation models. In *The North American Chapter of the Association for Computational Linguistics (NAACL)*.
- C.Y. Lin and Franz Josef Och. 2004. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Shuman Liu, Hongshen Chen, Zhaochun Ren, Yang Feng, Qun Liu, and Dawei Yin. 2018. [Knowledge diffusion for neural dialogue generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1489–1498. Association for Computational Linguistics.
- Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. In *Empirical Methods on Natural Language Processing (EMNLP)*.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Caglar Gulcehre, and Bing Xiang. 2016. Abstractive text summarization using sequence-to-sequence rnns and beyond. In *Computational Natural Language Learning (CoNLL)*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: A method for automatic evaluation of machine translation. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Dinesh Raghuram, Nikhil Gupta, and Mausam. 2018. Hierarchical pointer-generator network for task oriented dialog. *arXiv preprint arXiv:1805.01216*.
- Alan Ritter, Colin Cherry, and William B Dolan. 2011. Data-driven response generation in social media. In *Proceedings of the conference on empirical methods in natural language processing*, pages 583–593. Association for Computational Linguistics.
- M. Alexander Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. In *Empirical Methods on Natural Language Processing (EMNLP)*.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Iulian Vlad Serban, Alessandro Sordani, Yoshua Bengio, Aaron C Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- Iulian Vlad Serban, Alessandro Sordani, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron C Courville, and Yoshua Bengio. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- Sainbayar Sukhbaatar, Arthur Szlam, Jason Weston, and Fergus Rob. 2014. End-to-end memory networks. In *Advances in Neural Information Processing Systems (NIPS)*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems (NIPS)*.

- Jiwei Tan, Xiaojun Wan, and Jianguo Xiao. 2017. Abstractive document summarization with a graph-based attentional neural model. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2014. [Cider: Consensus-based image description evaluation](#). *arXiv preprint arXiv:1411.5726*.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. In *Advances in Neural Information Processing Systems (NIPS)*.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Jason Weston, Emily Dinan, and Alexander H. Miller. 2018. Retrieve and refine: Improved sequence generation models for dialogue. *arXiv preprint arXiv:1808.04776v2*.
- Caiming Xiong, Victor Zhong, and Richard Socher. 2017. Dynamic coattention networks for question answering. In *International Conference on Learning Representations (ICLR)*.
- Semih Yavuz, Chung-Cheng Chiu, Patrick Nguyen, and Yonghui Wu. 2018. CaLcs: Continuously approximating longest common subsequence for sequence level optimization. In *Empirical Methods on Natural Language Processing (EMNLP)*.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Barret Zoph and Kevin Knight. 2016. Multi-source neural translation. In *The North American Chapter of the Association for Computational Linguistics (NAACL)*.

Towards End-to-End Learning for Efficient Dialogue Agent by Modeling Looking-ahead Ability

Zhuoxuan Jiang¹, Xian-Ling Mao^{2*}, Ziming Huang³, Jie Ma³, Shaochun Li³

¹IBM Research / Shanghai, China

²Beijing Institute of Technology / Beijing, China

³IBM Research / Beijing, China

{jzxjiang, hzmzi, bjmajie, lishaoc}@cn.ibm.com, maoxl@bit.edu.cn

Abstract

Learning an efficient manager of dialogue agent from data with little manual intervention is important, especially for goal-oriented dialogues. However, existing methods either take too many manual efforts (e.g. reinforcement learning methods) or cannot guarantee the dialogue efficiency (e.g. sequence-to-sequence methods). In this paper, we address this problem by proposing a novel end-to-end learning model to train a dialogue agent that can look ahead for several future turns and generate an optimal response to make the dialogue efficient. Our method is data-driven and does not require too much manual work for intervention during system design. We evaluate our method on two datasets of different scenarios and the experimental results demonstrate the efficiency of our model.

1 Introduction

Research for dialogue system attracts a lot of attentions recently due to its potential huge value of reducing human cost in many commercial domains, such as restaurant reservation (Henderson et al., 2014b), travel planning (Peng et al., 2017) and retail service (Zhang et al., 2019). The majority of dialogue agents with goals are expected to be efficient to complete tasks with as few as possible dialogue turns, in contrast with those chat-counterparts (Ritter et al., 2011). The following two examples show the comparison of efficient and inefficient cases respectively. The scenarios is restaurant reservation and the agent’s goal is to reserve a table at noon.

Efficient example:

“Human: We don’t have empty tables at 11 o’clock tomorrow. All are reserved.”

“Agent: How about 12 o’clock? We are also okay then.”

Inefficient example:

“Human: We don’t have empty tables at 11 o’clock tomorrow. All are reserved.”

“Agent: What else time is available?”

“Human: 12 o’clock is ok.”

“Agent: All right. We want that time.”

For expressing the same opinion, the inefficient example consumes four turns while the efficient example only needs two. As it can be seen, the efficiency is important for goal-oriented dialogue systems to achieve goals in a rapid way.

Usually, a dialogue system consists of a pipeline of natural language understanding (NLU), dialogue management (DM) and natural language generation (NLG), where the DM part is treated as two separate components: dialogue state tracking (DST) and dialogue control (DC, i.e. dialogue policy selection). The DM part is widely considered to be relevant to the dialogue’s efficiency, because it makes decisions on what to say for the next turn. Recently, methods based on reinforcement learning are proposed for the policy selection component to build efficient dialogue systems. However, there are some drawbacks of reinforcement learning based methods. For example, they require lots of human work to design the learning strategy. Also a real-world environment which is essential for the agent to learn from is expensive, such as from domain experts. Moreover, training the dialogue manager as a two separate components could lead to error propagation issue (Rastogi et al., 2018).

In addition to reinforcement learning based methods, sequence-to-sequence based methods are also popular recently, because they can learn a dialogue agent purely from data and almost without too many human efforts. The error propagation issue can also be reduced because they are end-to-end, and they have better scalability for different scenarios. However, it is difficult to build efficient dialogue agents by those methods

*Xian-Ling Mao is the corresponding author.

since their objective functions for training models are usually inclined to general responses, such as *I don't know*, *yes* and *OK*, or often generate the same response for totally different contexts because the contextual information is not well-modeled by those methods (Dodge et al., 2015).

In this paper, we address the problem of learning an efficient dialogue manager from the perspective of reducing manual intervention and error propagation, and propose a new sequence-to-sequence based approach. The proposed end-to-end model contains a novel looking-ahead module for dialogue manager to learn the looking-ahead ability. Our intuition is that by predicting the future several dialogue turns, the agent could make a better decision of what to say for current turn, and therefore goals could be sooner achieved in a long run.

More specifically, our model includes three modules: (1) encoding module, (2) looking-ahead module, and (3) decoding module. At each dialogue turn, three kinds of information, the goals, historical utterances and the current user utterance, are utilized. First they are encoded by three separate Bidirectional Gated Recurrent Units (BiGRU) models. Then the three encoded embeddings are concatenated to one vector, which is then sent to a new bidirectional neural network that can look ahead for several turns. The decoding module will generate utterances for each turn through a learned language model. At last, by considering all the predicted future utterances, a new real system utterance for the next turn is re-generated by using an attention model through the same language model.

Our proposed approach has several advantages. First, it is an end-to-end model and does not take too many human efforts for system design. Although the goals should be handcrafted for specific scenario, the number of goals is small and it is a relatively easy work. Moreover, compared with naive sequence-to-sequence based models, our agent can make the dialogue more efficient by modeling the looking-ahead ability. Experimental results show that our model performs better than baselines on two datasets from different domains, which could suggest that our model is also scalable to various domains.

The contributions in this paper include:

- We identify the problem that how to make dialogues efficient by exploiting as little as pos-

sible manual intervention during system design from the perspective of end-to-end deep learning.

- We propose a novel end-to-end and data-driven model that enables the dialogue agent to learn to look ahead and make efficient decisions of what to say for the next turn.
- Experiments conducted on two datasets demonstrate that our model performs better over baselines and can be applied to different domains.

2 Related Work

In most situations, the dialogue systems require handcrafted definition of dialogue states and dialogue policies (Williams and Young, 2007; Henderson et al., 2014a; Asher et al., 2012; Chen et al., 2017). Those methods make the pipeline of dialogue systems clear to design and easy to maintain, but suffer from the massive expensive human efforts and the error propagation issue (Henderson et al., 2014c; Liu and Lane, 2017).

Reinforcement learning based methods for dialogue policy selection are widely studied recently (Lipton et al., 2018; Dhingra et al., 2017; Zhao and Eskenazi, 2016; Su et al., 2016). These methods only need human to design the learning strategies and do not require massive training data. However, the expensive domain knowledge and human expert efforts for agents to learn from are necessary (Liu et al., 2018; Shah et al., 2018). Therefore, hybrid methods that integrate supervised learning and reinforcement learning are proposed recently (Williams et al., 2017; Williams and Zweig, 2016). Thus, collecting massive training data becomes another manual work.

More recently, end-to-end dialogue systems attract much attention because almost no human efforts are required and they are scalable for different domains (Wen et al., 2017; Li et al., 2017; Lewis et al., 2017; Luo et al., 2019), especially with sequence-to-sequence based models (Sutskever et al., 2014). Although those models have been proved to be effective on chit-chat conversations (Ritter et al., 2011; Li et al., 2016a; Zhang et al., 2018), how to build agents that are goal-oriented with efficient dialogue managers through end-to-end approaches still remains questionable (Bordes et al., 2017; Joshi et al., 2017), and we investigate the question in this paper.

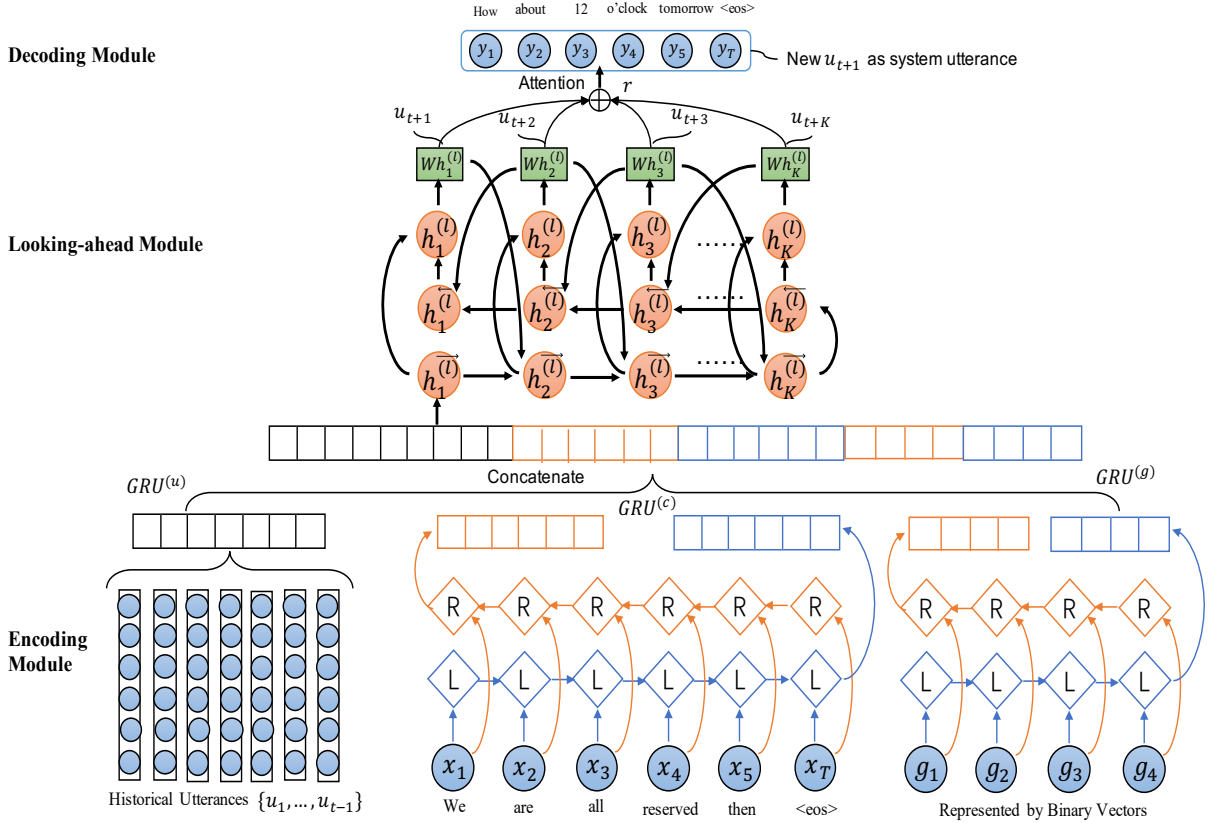


Figure 1: End-to-end model for learning looking-ahead ability.

Our idea of enabling the agent to be efficient by modeling looking-ahead ability is inspired by the AI Planning concept, which is a traditional searching technology in the field of AI, and is suitable for goal-based tasks, such as robotics control (Norvig and Russell, 1995). Recently, the concept is borrowed to dialogue system communities and integrated into deep learning models. For example, a trade-off method for training the agents neither with real human nor with user simulators is proposed, in order to obtain better policy learning results (Peng et al., 2018). In addition, at earlier time, the planning idea has been utilized for improving the dialogue generation task (Stent et al., 2004; Walker et al., 2007).

3 End-to-end Dialogue Model

We propose an end-to-end model that contains three modules: (1) encoding module, (2) looking-ahead module, and (3) decoding module. Figure 1 shows the model architecture. We leverage Bidirectional GRU models (Bahdanau et al., 2014) to encode agent goals, historical and current utterances. Then the obtained representations by encoding goals and utterances are regarded as inputs

of the looking-ahead module, and they are used to predict several future turns. At last the predicted future turns are merged by an attention model and the new real system utterance is generated for the next turn.

Suppose for each dialogue session we have T turns, and we do not distinguish whether it is user's turn or system's turn. If the agent has S goals that are denoted as $g = \{g_1, g_2, \dots, g_S\}$, each goal is formalized as a binary vector. For example in the restaurant reservation scenario, we can define that each variate in the vector $[1, 0]$ corresponds to a yes-no condition, such as the 1 means agent accepts bar table and the 0 means agent does not want to change time. As to the utterance information, imagine at turn $t \in \{1, \dots, T\}$, we denote utterances $\{u_1, \dots, u_{t-1}\} \in \mathbb{U}$ for historical ones and $u_t \in \mathbb{U}$ for current user utterance. Our model predicts the system and user utterances $\{u_{t+1}, u_{t+2}, \dots, u_{t+K}\}$ for the next K turns and then a new u_{t+1} is generated as the system utterance after considering all the predicted turns. The model separates the current user utterance from historical ones in order to highlight the user's current states. In general, the model is end-to-end and

needs little human intervention or domain knowledge.

3.1 Encoding Module

In this module, the agent goals, historical utterances within the dialogue session, and the current user utterance are encoded by using three GRU models which is expected to learn long-range temporal dependencies (Cho et al., 2014). $GRU^{(g)}$ is defined to encode agent’s goals g and the final hidden state $h^{(g)}$ is taken as the representation of goals. The input of $GRU^{(g)}$ is a one-hot binary vector with length S . $GRU^{(u)}$ is used to encode the historical utterances, and $GRU^{(c)}$ is used to encode the current user utterance. $h^{(u)}$ and $h^{(c)}$ are denoted as the final encoded representations of $GRU^{(u)}$ and $GRU^{(c)}$ respectively.

To get the i -th hidden state for the three GRUs, respective inputs include the previous hidden state $h_{i-1}^{(g)}$, $h_{i-1}^{(u)}$ or $h_{i-1}^{(c)}$, and the embeddings of current observations, $E(g_i)$, $E(u_i)$ or $E(x_i)$, where g_i is a goal, u_i is an utterance and x_i is a token. For the textual tokens, we use the Word2vec embeddings as their representations (Mikolov et al., 2013). Then the token embeddings are averaged to represent utterances. The formal denotation of the hidden states for the three GRU models is:

$$h_i^{(g)} = GRU^{(g)}(h_{i-1}^{(g)}, E(g_i)), \quad (1)$$

$$h_i^{(u)} = GRU^{(u)}(h_{i-1}^{(u)}, E(u_i)), \quad (2)$$

$$h_i^{(c)} = GRU^{(c)}(h_{i-1}^{(c)}, E(x_i)), \quad (3)$$

where $E(\cdot)$ represents the embeddings.

The final output of the encoding module is a concatenation of $h^{(g)}$, $h^{(u)}$ and $h^{(c)}$, which is denoted as $h_1^{\vec{r}} = [h^{(g)}, h^{(u)}, h^{(c)}]$. $h_1^{\vec{r}}$ serves as the input of the following looking-ahead module. The right arrow means the initial direction to train the looking-ahead module is from the current to the future.

3.2 Looking-ahead Module

With the input of $h_1^{\vec{r}}$, this module predicts several future dialogue turns. Since the process is sequential, we propose a recurrent neural network to model the process. In order to exploit the predicted information for later generating a real system utterance, another recurrent neural network is used to backtrack the information from future to

current. To reduce the computing cost, the two neural networks share the same parameters, and the whole looking-ahead module looks similar to a bidirectional GRU as shown in Figure 1.

We denote the module as $GRU^{(l)}$. $\{h_k^{(l)} | k > 0\}$ represent the predicted hidden states for future turns. To get $h_k^{(l)}$, the hidden states from two directions, $h_k^{\vec{r}}$ and h_k^{\leftarrow} , are concatenated. To calculate each $h_k^{\vec{r}}$ or h_k^{\leftarrow} , their inputs include the previous hidden state and the previously-predicted hidden state. Formally, suppose we look ahead for K turns, the hidden state of $h_k^{(l)}$ is calculated as following:

$$h_k^{\vec{r}} = GRU^{\vec{r}}(h_{k-1}^{\vec{r}}, Wh_{k-1}^{(l)}), \quad (4)$$

$$h_k^{\leftarrow} = GRU^{\leftarrow}(h_{k+1}^{\leftarrow}, Wh_{k+1}^{(l)}), \quad (5)$$

$$h_k^{(l)} = [h_k^{\vec{r}}, h_k^{\leftarrow}], \quad (6)$$

where W is a weight parameter and $Wh_k^{(l)}$ is the hidden state for predicting future turns. If $K = 1$, it means our model has no looking-ahead ability and it degrades to a naive goal-based sequence-to-sequence model.

3.3 Decoding Module

For generating the real system utterance, as seen in Figure 1, the green hidden states $\{Wh_k^{(l)} | k > 0\}$ are combined through an attention based model (Wang et al., 2016). The formal denotation is:

$$e_k = \tanh(W^{(a)}Wh_k^{(l)}), \quad (7)$$

$$v_k = \frac{\exp(e_k)}{\sum_{k=1}^K \exp(e_k)}, \quad (8)$$

$$r = \sum_{k=1}^K v_k h_k^{(l)}, \quad (9)$$

where $W^{(a)}$ is the attention weight parameter and r is the input representation for generating a new u_{t+1} that is regarded as the real system utterance.

Given the hidden state $Wh_k^{(l)}$, the decoding module can also generate the corresponding utterance for learning the looking-ahead ability. We share the parameters of decoding with those in the

encoding module, in order to reduce the computing cost (Vinyals and Le, 2015). The token sequence in u_{t+k} is generated from left to right by selecting the tokens with the maximum probability distribution through a language model learned by the following equation:

$$p_\theta(y_j^{(t+k)} | y_{1,2,\dots,j-1}^{(t+k)}) \propto \exp(E^T W h_k^{(l)}). \quad (10)$$

3.4 Model Training

To train the proposed model, we define a loss function to maximize three terms: (1) a language model for predicting tokens in language generation, (2) the probability distribution of predicting utterances of future dialogue turns, and (3) a binary classifier to predict if the dialogue will be complete or not. The final joint loss function is formally denoted as:

$$\begin{aligned} L(\theta) = & - \underbrace{\sum_u \sum_i \log p_\theta(x_i | x_{1,\dots,i-1})}_{\text{language model loss}} \\ & - \alpha \underbrace{\sum_{u,g} \sum_k \sum_i \log p_\theta(y_i^{(t+k)} | y_{1,\dots,i-1}^{(t+k)}, u, g)}_{\text{looking ahead prediction loss}} \\ & - \beta \underbrace{\sum_c \log p(z_c | c, u_{t+1})}_{\text{dialog state prediction loss}}, \end{aligned} \quad (11)$$

where

$$u_{t+1} = \arg \max_y p_\theta(y | r), \quad (12)$$

$$\begin{aligned} \log p(z_c | c, u_{t+1}) = & z_c \log(g(c, u_{t+1})) \\ & + (1 - z_c) \log(1 - g(c, u_{t+1})). \end{aligned} \quad (13)$$

$g(\cdot)$ is a sigmoid function and z_c is the label of the dialogue that current user utterance c belongs to, where 1 means the dialogue ends up with goals achieved while 0 means the goals are not achieved. The three terms are weighted with two hyperparameters α and β . We adopt stochastic gradient descent method to minimize $L(\theta)$.

In the looking-ahead module, the hidden state $W h_k^{(l)}$ is used to generate an utterance $y^{(t+k)}$, and is also used to calculate h_{k+1}^{\rightarrow} and h_{k-1}^{\leftarrow} . We design an EM-like algorithm to optimize the loss function, as described in Algorithm 1. Line 3-4 optimize the language model, i.e. the first term of $L(\theta)$. Line 5-16 optimize the looking-ahead module, i.e. the second term, among which Line 7-14 are for E-step and Line 15-16 are for M-step.

In E-step the language model is fixed for updating all the hidden states $h_k^{(l)}$ in looking-ahead module, and in M-step all the hidden states are fixed for updating the language model. Line 17-18 optimize the third term of $L(\theta)$, which is a binary classifier.

Algorithm 1: Learning algorithm for $L(\theta)$

input : Dialogue utterances U , Agent goals g ,
Looking-ahead turns K

output: Agent model θ

- 1 Randomly initializing parameters;
- 2 **for** $c \in U$, g and historical utterances $\{u\}$ **do**
- 3 **for** $x_i \in c$ **do**
- 4 Optimizing $p_\theta(x_i | x_{1,\dots,i-1})$;
- 5 $h_1^{\rightarrow} = [h^{(g)}, h^{(u)}, h^{(c)}]$;
- 6 $u_{t+1} = \arg \max_y p_\theta(y | r)$;
- 7 E-Step: Update $h_k^{(l)}$ with fixed language model;
- 8 **for** $k = 1 : K$ **do**
- 9 $u_{t+k} = \arg \max_y p_\theta(y | h_k^{(l)})$;
- 10 $h_k^{\rightarrow} = [h_{k-1}^{\rightarrow}, W h_{k-1}^{(l)}]$;
- 11 $h_K^{\leftarrow} = h_K^{\rightarrow}$;
- 12 **for** $k = K - 1 : 1$ **do**
- 13 $h_k^{\leftarrow} = [h_{k+1}^{\leftarrow}, W h_{k+1}^{(l)}]$;
- 14 **for** $k = 1 : K$ **do**
- 15 $h_k^{(l)} = [h_k^{\rightarrow}, h_k^{\leftarrow}]$;
- 16 M-Step: Update language model with fixed $h_k^{(l)}$;
- 17 **for** $k = 1 : K$ **do**
- 18 Optimizing $p_\theta(y_i^{(t+k)} | y_{1,\dots,i-1}^{(t+k)})$;
- 19 $u_{t+1} = \arg \max_y p_\theta(y | r)$;
- 20 Optimizing $p(z_c | c, u_{t+1})$;
- 21 **return** θ ;

4 Experiments

4.1 Data Collection

We use two datasets for two different scenarios to evaluate our model. Table 1 shows the statistics of two datasets.

4.1.1 Dataset 1 - Object Division

Dataset 1 contains crowd-sourced dialogues between humans collected from Amazon Mechanical Turk platform (Lewis et al., 2017). The dataset is for *object division task* and both sides have separate goals of each object’s value. We use the textual data and transform their goals to yes-no questions as our binary vectors. The information of each dialogue session’s final state, agree or disagree, is used for training the agent.

4.1.2 Dataset 2 - Restaurant Reservation

To the best of our knowledge, there is no other public dataset for goal-oriented dialogues where

Metric	Dataset 1	Dataset 2
Number of Dialogues	5,808	1,613
Average Turns per Dialogue	6.6	6.3
Average Words per Turn	7.6	8.9
Number of Words	566,779	98,726
% Goal Achieved	80.1%	71.5%

Table 1: Statistic on the two datasets.

the two sides have different goals. To this end, we construct the Dataset 2 to testify the scalability of our model. The common scenario of restaurant table reservation is chosen.

In this dataset, the two agents are expected to have different goals and they dialogue with each other for looking for the intersection of their goals. We denote Agent A as the role of a customer and Agent B as the restaurant server side. At the beginning of each dialogue session, Agent A is given the available time slot, the number of people, and several other constraints (e.g. can sit at bar or not). All the constraints are regarded as its goals represented by a binary vector. Similarly, Agent B has itself constraints (e.g. whether bar tables are available or not), which are also treat as goals represented by a binary vector. We predefine a pool of ‘goals’ and at the beginning of each dialogue session, the goals for two sides are randomly sampled separately from the pool. The two agents cannot see each other’s goals and they dialogue through natural language until a final decision, agreement or disagreement, is reached. In summary, the objective of constructing this dataset is to see if our model can reach the intersection of the two agents’ goals in a more efficient way.

To generate dialogues for Dataset 2, we resort to a rule-based method via AI planning search (Ghahlab et al., 2016; Jiang et al., 2019). Watson AI platform¹ is leveraged for natural language understanding by defining intents and entities with examples. A planner is designed for the dialogue manager by defining several states and actions. The goals are represented as part of the states, and the STRIPS algorithm is used to search the shortest path to goals at each turn and return the first planned action for generating the next response. Each action has several handcrafted utterances since the diversity of utterances is not our focus in this paper. Table 2 shows a sample dialogue.

¹<https://www.ibm.com/watson/ai-assistant/>

<p>Alice: May I reserve a table for 6 people at 17 tomorrow? Bob: Sorry, we don’t have a table at this point. Alice: Can we sit at the bar then? Bob: We don’t have a bar in the restaurant. Alice: Can I have more expensive tables then? Bob: My apologies, we are required not to do that. Alice: In this case, can I reserve a bigger table? Bob: Yes, we have VIP rooms but more expensive. Alice: I want that. Bob: OK. Alice: Bye.</p>

Table 2: Sample of Dataset 2.

4.2 Training Sample Preparation

For each dialogue session with T turns, we reorganize the utterances into T samples. For each turn $t = \{1, 2, \dots, T\}$, we can get the current user utterance c , and a training sample is created with a historical utterance sequence $\{u_1, u_2, \dots, u_{t-1}\}$, and the goals g are consistent with the same dialogue session. The future K turns of utterances $\{u_{t+1}, u_{t+2}, \dots, u_{t+K}\}$ are used as the supervised information. In total, we get 38,333 and 10,162 samples including training set and test set for the two datasets respectively.

4.3 Baselines

Since our model is based on purely data-driven learning, we compare our model with the supervised counterparts. Our baselines include:

- Seq2Seq(goal): This is a naive baseline by adapting the sequence-to-sequence model (Sutskever et al., 2014) and encoding goals, which removes the looking-ahead module and the supervised information of final state prediction from our model.
- Seq2Seq(goal+state): This is a baseline model by removing the looking-ahead module from our proposed model. The parameter α is set to zero.
- Seq2Seq(goal+look): This is a baseline model by removing the supervised information of final state prediction from our model. The parameter β is set to zero.
- Seq2Seq(goal+look+state): This is our proposed model that includes all the modules and supervised information.

4.4 Evaluation Criteria

In a dialogue system, it could be treat as efficient if it obtains more final goal achievement with as few




Figure 2: vs. looking-ahead turns on Dataset 1

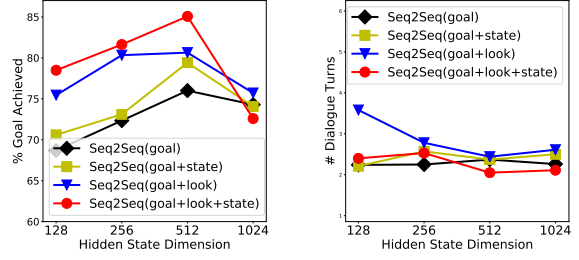


Figure 4: vs. hidden state dimension on Dataset 1




Figure 3: vs. looking-ahead turns on Dataset 2




Figure 5: vs. hidden state dimension on Dataset 2

as possible dialogue turns. Thus we set two criteria for evaluating and comparing models adopted in our experiments: (1) the *goal achievement ratio* that means the ratio of the number of goal achieved dialogue over the number of attempted dialogues, and (2) the *average dialogue turns*.

4.5 Evaluator

Our experiments are to achieve goals through conversations, and it is difficult to directly adopt existing simulators (Asri et al., 2016). We refer to the work (Li et al., 2016b) and fine-tune it to our task. For each dataset, a naive sequence-to-sequence model that encodes goals is regarded as the user simulator. We run 1000 times of dialogue sessions using the simulator.

Apart from using the simulator, we also invite humans to dialogue with the agents for 100 times each person for each dataset and we report the average results.

4.6 Training Settings

All the baselines are implemented by PyTorch. One-hot input tokens are embedded into a 64-dimensional space. The goals are encoded by $GRU^{(g)}$ with a hidden layer of size 64. The sizes of hidden states in input utterance encoder $GRU^{(u)}$, $GRU^{(c)}$ and looking-ahead module $GRU^{(l)}$, $h_k^{(l)}$, are all set to 256. A stochastic gradient descent method is employed to opti-

mize the model with a mini-batch size of 32 for supervised learning, an initial learning rate of 1.0, momentum with $\mu = 0.1$, and clipping gradients 0.5 in L^2 norm. The best model is chosen from the processing of training the model for 400 epochs. After that, the learning rate decays by a factor of 2 for every epoch. The initial hyper-parameters setting in the loss function (Equation (11)) is $\alpha = 0.05$ and $\beta = 1.0$. Words that appear in the training dataset for less than 5 times are replaced with the ‘unknown’ ($\langle unk \rangle$) token. A validation dataset is employed to choose the optimal hyper-parameters.

4.7 Results and Analysis

Table 3 shows the performance of baselines against user simulator and human on the two datasets. Both reveal that models that learn looking-ahead ability can achieve better performance and deliver more efficient dialogues in terms of both goal achievement ratio and dialogue turns. However, in the table, the dialogue turns of Seq2Seq(goal+look+state) are larger than those of Seq2Seq(goal+look), which may suggest that more dialogue turns lead to more achievement. In spite of this, the looking-ahead ability learned by our model is demonstrated to be effective on the two different scenarios. Moreover, the supervised information of final states (the third term of Equation (11)) is also proven effective in deliver-

Model	Dataset 1				Dataset 2			
	vs. <i>Simulator</i>		vs. <i>Human</i>		vs. <i>Simulator</i>		vs. <i>Human</i>	
	% Achieved	# Turns	% Achieved	# Turns	% Achieved	# Turns	% Achieved	# Turns
Seq2Seq(goal)	76.00	4.74	67.74	7.87	67.10	7.38	54.1	7.56
Seq2Seq(goal+state)	79.41	4.74	70.97	6.35	67.37	7.42	58.1	8.04
Seq2Seq(goal+look)	80.64	6.54	74.19	5.41	83.54	5.82	60.3	6.94
Seq2Seq(goal+look+state)	85.07	4.10	77.42	5.02	83.58	6.36	61.2	7.30

Table 3: Performance on two datasets against the user simulator and human.

<i>Seq2Seq(goal) Model:</i>
Alice: i just want the book
Bob: no way i have the book and you can take others
Alice: bye
<Conversation end>
<i>Seq2Seq(goal+look+state) Model:</i>
Alice: i keep the book you can have the hat and two balls
Bob: i will take the book and you can have others
Alice: can we make a deal?
Bob: ok deal
Alice: okay
Bob: bye
<Conversation end>

Table 4: Example that tends to achieve goals on Dataset 1, vs. simulator.

ing more achievement, which can be seen from the second and last rows of Table 3. Compared to the human evaluation, the results with the simulator generally are better. It is probable that human evaluators tend to be more rigorous and more turns are necessary to achieve goals.

We also evaluate the models’ performance with different parameters. Figure 2 and Figure 3 show the results of using different looking-ahead turns K in the looking-ahead module on Dataset 1 and Dataset 2 respectively. In most cases, Seq2Seq(goal+look+state) is better than Seq2Seq(goal+look). Setting the number to 3 is the recommended value, since our scenarios are relatively simple and it is enough for the agents to look ahead for 3 dialogue turns. The results are consistent with those in Table 3.

The performance on different dimension sizes of hidden states used in all the baselines is evaluated and the results are shown in Figure 4 and Figure 5. We find that setting the size to 512 yields the optimal results by balancing the performance and the computing cost in most cases.

In addition, we take some case studies for dialogue examples between the agents trained by Seq2Seq(goal+look+state) (proposed model) and the simulator (Seq2Seq(goal)) for the two datasets. *Alice* denotes the agent and *Bob* denotes the simulator. Table 4 is a case where our model with the looking-ahead ability (Seq2Seq(goal+look+state)) tends to achieve goals while consume more dia-

<i>Seq2Seq(goal) Model:</i>
Alice: can you help me book a table for 2 people at 6pm?
Bob: sorry, we dont have a table at this point
Alice: can i reserve the seats at the bar instead?
Bob: yes you can
Alice: can i book a table for us
Bob: sure i have written down your reservation
Alice: ok
Bob: bye
<Conversation end>
<i>Seq2Seq(goal+look+state) Model:</i>
Alice: can you help me book a table for 2 people at 6pm?
Bob: sorry we dont have a table at this point
Alice: can i reserve the seats at the bar instead?
Bob: sure i have written down your reservation
Alice: bye
<Conversation end>

Table 5: Example that costs less dialogue turns under the same goals on Dataset 2, vs. simulator.

logue turns. Table 5 gives an example where our model not only tends to achieve goals but also takes fewer dialogue turns. Based on our observation, the language models for natural language generation is learned well in spite of that confused utterances are generated sometimes.

5 Conclusion

In this paper, we propose an end-to-end model towards the problem of how to learn an efficient dialogue manager without taking too many manual efforts. We model the looking-ahead ability for foreseeing several turns and then the agent can make a decision of what to say that leads the conversation to achieve goals with as few as possible dialogue turns. Experiments on two datasets from different domains demonstrate that our model is efficient in terms of goal achievement ratio and average dialogue turns. Our method is also scalable and can reduce error propagation due to the nature of end-to-end learning.

For the future work, we expect to investigate whether other kinds of abilities, such as reasoning ability, can be modeled for agent towards the problem. In addition to the efficiency issue, the quality of natural language generation should also be paid attention in order to guarantee the quality of overall dialogue system.

Acknowledgments

The work is partially supported by SFSMBRP (2018YFB1005100), BIGKE (No. 20160754021), NSFC (No. 61772076 and 61751201), NSFB (No. Z181100008918002), CETC (No. w-2018018) and OPBKLICDD (No. ICDD201901). We thank Tian Lan, Henda Xu and Jingyi Lu for experiment preparation. We also thank the three anonymous reviewers for their insightful comments.

References

- Nicholas Asher, Alex Lascarides, Oliver Lemon, Markus Guhe, Verena Rieser, Philippe Muller, Stergos Afantenos, Farah Benamara, Laure Vieu, Pascal Denis, S. Paul, S. Keizer, and C. Degrémont. 2012. Modelling strategic conversation: The stac project. In *SemDial*, page 27.
- Layla El Asri, Jing He, and Kaheer Suleman. 2016. A sequence-to-sequence model for user simulation in spoken dialogue systems. In *INTERSPEECH*, pages 1151–1155.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. ArXiv preprint arXiv:1409.0473.
- Antoine Bordes, Y-Lan Boureau, and Jason Weston. 2017. Learning end-to-end goal-oriented dialog. In *ICLR*.
- Hongshen Chen, Xiaorui Liu, Dawei Yin, and Jiliang Tang. 2017. A survey on dialogue systems- recent advances and new frontiers. *ACM SIGKDD Explorations Newsletter*, 19(2):25–35.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder-decoder approaches. In *SSST-8*, pages 103–114.
- Bhuwan Dhingra, Lihong Li, Xiujun Li, Jianfeng Gao, Yun-Nung Chen, Faisal Ahmed, and Li Deng. 2017. Towards end-to-end reinforcement learning of dialogue agents for information access. In *ACL*, pages 484–495.
- Jesse Dodge, Andreea Gane, Xiang Zhang, Antoine Bordes, Sumit Chopra, Alexander Miller, Arthur Szlam, and Jason Weston. 2015. Evaluating prerequisite qualities for learning end-to-end dialog systems. ArXiv preprint arXiv:1511.06931.
- Malik Ghallab, Dana Nau, and Paolo Traverso. 2016. *Automated Planning and Acting*. Cambridge University Press.
- Matthew Henderson, Blaise Thomson, and Jason D. Williams. 2014a. The second dialog state tracking challenge. In *SIGDIAL*, pages 263–272.
- Matthew Henderson, Blaise Thomson, and Jason D. Williams. 2014b. The third dialog state tracking challenge. In *SLT*, pages 324–329.
- Matthew Henderson, Blaise Thomson, and Steve Young. 2014c. Word-based dialog state tracking with recurrent neural networks. In *SIGDIAL*, pages 292–299.
- Zhuoxuan Jiang, Jie Ma, Jingyi Lu, Guangyuan Yu, Yipeng Yu, and Shaochun Li. 2019. A general planning-based framework for goal-driven conversation assistant. In *AAAI*, pages 9857–9858.
- Chaitanya K. Joshi, Fei Mi, and Boi Faltings. 2017. Personalization in goal-oriented dialog. In *NIPS*.
- Mike Lewis, Denis Yarats, Yann N. Dauphin, Devi Parikh, and Dhruv Batra. 2017. Deal or no deal? end-to-end learning for negotiation dialogues. In *EMNLP*, pages 2443–2453.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. A diversity-promoting objective function for neural conversation models. In *NAACL*, pages 110–119.
- Xiujun Li, Yun-Nung Chen, Lihong Li, Jianfeng Gao, and Asli Celikyilmaz. 2017. End-to-end task-completion neural dialogue systems. In *IJCNLP*, pages 733–743.
- Xiujun Li, Zachary C Lipton, Bhuwan Dhingra, Lihong Li, Jianfeng Gao, and Yun-Nung Chen. 2016b. A user simulator for task-completion dialogues. *arXiv preprint arXiv:1612.05688*.
- Zachary Lipton, Xiujun Li, Jianfeng Gao, Lihong Li, Faisal Ahmed, and Li Deng. 2018. Bbq-networks: Efficient exploration in deep reinforcement learning for task-oriented dialogue systems. In *AAAI*, pages 5237–5244.
- Bing Liu and Ian Lane. 2017. An end-to-end trainable neural network model with belief tracking for task-oriented dialog. In *INTERSPEECH*, pages 2506–2510.
- Bing Liu, Gokhan Tür, Dilek Hakkani-Tür, Pararth Shah, and Larry Heck. 2018. Dialogue learning with human teaching and feedback in end-to-end trainable task-oriented dialogue systems. In *NAACL*, pages 2060–2069.
- Liangchen Luo, Wenhao Huang, Qi Zeng, Zaiqing Nie, and Xu Sun. 2019. Learning personalized end-to-end goal-oriented dialog. In *AAAI*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*, pages 3111–3119.
- Peter Norvig and Stuart J. Russell. 1995. *Artificial Intelligence: A Modern Approach*. Prentice Hall.

- Baolin Peng, Xiujun Li, Jianfeng Gao, Jingjing Liu, and Kam-Fai Wong. 2018. Deep dyna-q: Integrating planning for task-completion dialogue policy learning. In *ACL*, pages 2182–2192.
- Baolin Peng, Xiujun Li, Lihong Li, Jianfeng Gao, Asli Celikyilmaz, Sungjin Lee, and Kam-Fai Wong. 2017. Composite task-completion dialogue policy learning via hierarchical deep reinforcement learning. In *EMNLP*, pages 2231–2240.
- Abhinav Rastogi, Raghav Gupta, and Dilek Hakkani-Tür. 2018. Multi-task learning for joint language understanding and dialogue state tracking. In *SIGDIAL*, pages 376–384.
- Alan Ritter, Colin Cherry, and William B. Dolan. 2011. Data-driven response generation in social media. In *EMNLP*, pages 583–593.
- Pararth Shah, Dilek Hakkani-Tür, Bing Liu, and Gokhan Tür. 2018. Bootstrapping a neural conversational agent with dialogue self-play, crowdsourcing and on-line reinforcement learning. In *NAACL*, pages 41–51.
- Amanda Stent, Rashmi Prasad, and Marilyn Walker. 2004. Trainable sentence planning for complex information presentation in spoken dialog systems. In *ACL*, page 79.
- Pei-Hao Su, Milica Gasic, Nikola Mrksic, Lina Rojas-Barahona, Stefan Ultes, David Vandyke, Tsung-Hsien Wen, and Steve Young. 2016. Continuously learning neural dialogue management. ArXiv preprint arXiv:1606.02689.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*, pages 3104–3112.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. ArXiv preprint arXiv:1506.05869.
- Marilyn Walker, Amanda Stent, François Mairesse, and Rashmi Prasad. 2007. Individual and domain adaptation in sentence planning for dialogue. *Journal of Artificial Intelligence Research*, 30.
- Yequan Wang, Minlie Huang, Li Zhao, and Xiaoyan Zhu. 2016. Attention-based lstm for aspect-level sentiment classification. In *EMNLP*, pages 606–615.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gašić, Lina M. Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017. A network-based end-to-end trainable task-oriented dialogue system. In *EACL*, pages 438–449.
- Jason D. Williams, Kavosh Asadi, and Geoffrey Zweig. 2017. Hybrid code networks: practical and efficient end-to-end dialog control with supervised and reinforcement learning. In *ACL*, pages 665–677.
- Jason D. Williams and Steve Young. 2007. Partially observable markov decision processes for spoken dialogue systems. *Computer Speech & Language*, 21(2):393–422.
- Jason D. Williams and Geoffrey Zweig. 2016. End-to-end lstm-based dialog control optimized with supervised and reinforcement learning. ArXiv preprint arXiv:1606.01269.
- Wei-Nan Zhang, Yiming Cui, Yifa Wang, Qingfu Zhu, Lingzhi Li, Lianqiang Zhou, and Ting Liu. 2018. Context-sensitive generation of open-domain conversational responses. In *COLING*, pages 2437–2447.
- Zheng Zhang, Lizi Liao, Minlie Huang, Xiaoyan Zhu, and Tat-Seng Chua. 2019. Neural multimodal belief tracker with adaptive attention for dialogue systems. In *WWW*, pages 2401–2412.
- Tiancheng Zhao and Maxine Eskenazi. 2016. Towards end-to-end learning for dialog state tracking and management using deep reinforcement learning. In *SIGDIAL*, pages 1–10.

Unsupervised Dialogue Spectrum Generation for Log Dialogue Ranking

Xinnuo Xu[†], Yizhe Zhang[‡], Lars Liden[‡], Sungjin Lee[‡]

[†]The Interaction Lab, Heriot-Watt University, Edinburgh
xx6@hw.ac.uk

[‡]Microsoft Research, Redmond, WA, USA

Yizhe.Zhang, Lars.Liden, sule@microsoft.com

Abstract

Although the data-driven approaches of some recent bot building platforms make it possible for a wide range of users to easily create dialogue systems, those platforms don't offer tools for quickly identifying which log dialogues contain problems. This is important since corrections to log dialogues provide a means to improve performance after deployment. A log dialogue ranker, which ranks problematic dialogues higher, is an essential tool due to the sheer volume of log dialogues that could be generated. However, training a ranker typically requires labelling a substantial amount of data, which is not feasible for most users. In this paper, we present a novel unsupervised approach for dialogue ranking using GANs and release a corpus of labelled dialogues for evaluation and comparison with supervised methods. The evaluation result shows that our method compares favorably to supervised methods without any labelled data.

1 Introduction

Task-oriented dialogue systems provide a natural interface to accomplish various daily-life tasks such as restaurant finding and flight booking. Data-driven approaches offered by common bot building platforms (e.g. Google Dialogflow, Amazon Alexa Skills Kit, Microsoft Bot Framework) make it possible for a wide range of users to easily create dialogue systems with a limited amount of data in their domain of interest. Typically, the development process of a dialogue system based on data-driven approaches (Williams et al., 2017; Bordes et al., 2016) goes around an operational loop in Figure 1: (1) The cycle begins with a developer creating a training dataset with seed dialogues. (2) A dialogue system is trained and deployed. (3) Real users interact with the system and generate log dialogues. (4) The developer reviews the logs to identify which log dialogues




Figure 1: Operational loop of data-driven dialogue system development.

contain problems. (5) The developer updates the training dataset to fix the problems. (6) The cycle repeats from step 2). Of all steps, (4) is the most significant in slowing down the loop, because of the sheer volume of log dialogues that can be generated and the need to manually inspect each. Thus, it is essential to support tools that help developers quickly identify problematic log dialogues. To achieve this goal, we propose a neural dialog ranker whose goal is to place problematic dialogues higher in the rank.

However, training a ranker typically requires labelling a substantial amount of data, which is not feasible for most developers. Furthermore, one might have to repeat this process whenever a significant change is made to the system's behavior. This motivates us to explore a set of unsupervised approaches to reduce the prohibitive cost. The core idea of these methods is that we learn a generative model to produce problematic dialogue examples as positive examples and train a ranker with seed dialogues used as negative examples. Specifically, we propose a novel dialogue generator using Generative Adversarial Networks (GANs) and train the generator with a curriculum learning scheme. Another possible avenue is to leverage off-the-shelf dialogue quality classifiers which are trained on open-domain corpora such as




Figure 2: Overall architecture of our dialogue ranker.

dialogue breakdown detection challenge systems in DSTC6 (Higashinaka et al., 2017). In the experiment on the labelled dialogue corpus that we collected via Amazon Mechanical Turk, we show that our approach outperforms the off-the-shelf model by a significant margin thanks to the capability of generating domain-relevant problematic dialogues. The evaluation result also shows that our unsupervised method compares favorably to supervised methods without any labelled data.

The rest of this paper is organized as follows. In Section 2, we formalize the dialogue ranking task and describe our neural ranking model. In Section 3, we introduce a set of unsupervised methods for problematic dialogue example generation. Section 4 describes the datasets we used for this study. Section 5 explains our experiments. In Section 6, we discuss our experimental results. Section 7 provides a survey of related work. We finish with conclusions and future work in Section 8.

2 Dialogue Ranking

A dialogue ranker aims to assign higher scores to problematic dialogues than normal ones so that developers may quickly identify problematic dialogues in the ranked list of log dialogues. Formally, given a dialogue $\mathbf{D} = \{S_1, U_1, S_2, \dots, U_{n-1}, S_n\}$, a dialogue ranker ψ produces a score of \mathbf{D} being problematic where S_i and U_i are the system and user utterance in i^{th} turn, respectively.¹ To train the dialogue ranker ψ , we formulate the ranking task as binary classifica-

¹One turn consists of a pair of system and user utterances.

tion where problematic and normal dialogues correspond to positive and negative classes, respectively. We optimize the cross-entropy objective:

$$\mathcal{L}_{\text{xent}} = \frac{1}{N} \sum_{i=1}^N [y_i \log \hat{y}_i + (1 - y_i) \log(1 - \hat{y}_i)]$$

where $\hat{y}_i = 1/(1 + \exp(-\psi(\mathbf{D}_i)))$ and $y_i = 1$ for positive class and 0 otherwise.

We design a deep neural ranking model for ψ to automatically learn salient features as illustrated in Figure 2. We first use a bi-directional LSTM (Bi-LSTM) to encode each utterance in dialogue \mathbf{D} :

$$\mathbf{D}^{\text{emb}} = \{S_1^{\text{emb}}, U_1^{\text{emb}}, S_2^{\text{emb}}, \dots, U_{n-1}^{\text{emb}}, S_n^{\text{emb}}\}$$

where $S_i^{\text{emb}} = \text{Bi-LSTM}(S_i)$ and $U_i^{\text{emb}} = \text{Bi-LSTM}(U_i)$. Then, we calculate attention embeddings for each utterance with a multi-head self-attention mechanism (Vaswani et al., 2017):

$$\mathbf{D}^{\text{attn}} = \text{softmax} \left(\frac{\mathbf{D}^{\text{emb}} \mathbf{D}^{\text{emb}\top}}{\sqrt{d}} \right) \mathbf{D}^{\text{emb}}$$

where d is the dimension of the embedding and $\mathbf{D}^{\text{attn}} = \{S_1^{\text{attn}}, U_1^{\text{attn}}, \dots, U_{n-1}^{\text{attn}}, S_n^{\text{attn}}\}$. Now, we apply a feed-forward network (FFN) to the concatenation of utterance embeddings \mathbf{D}^{emb} and their corresponding attentive embeddings \mathbf{D}^{attn} to yield context-sensitive utterance embeddings:

$$\mathbf{D}^{\text{ctxt}} = \{S_1^{\text{ctxt}}, U_1^{\text{ctxt}}, \dots, U_{n-1}^{\text{ctxt}}, S_n^{\text{ctxt}}\}$$

where $S(U)_i^{\text{ctxt}} = \text{FFN}([S(U)_i^{\text{emb}}, S(U)_i^{\text{attn}}])$ and $[\cdot, \cdot]$ denotes a concatenation operator. After that, we apply another FFN followed by a sigmoid activation to each pair of utterances to measure the consistency of adjacency pairs:

$$A_i = \text{sigmoid}(\text{FFN}([X_i^{\text{ctxt}}, Y_i^{\text{ctxt}}]))$$

where $(X_i^{\text{ctxt}}, Y_i^{\text{ctxt}})$ is either $(S_i^{\text{ctxt}}, U_i^{\text{ctxt}})$ or $(U_i^{\text{ctxt}}, S_{i+1}^{\text{ctxt}})$. Finally, the ranker ψ produces a ranking score for the dialogue based on the consistency scores and a set of manually crafted features:

$$\psi(\mathbf{D}) = \text{FFN}([A_1, \dots, A_{n-1}, f_1, \dots, f_m])$$

where f_i denotes a set of manual features. In this study, we use a single manual feature to consider redundant turns:

$$f = \frac{\text{Num}(\text{distinct utterances})}{\text{Num}(\text{all utterances})}$$

Each instance of $FFNs$ has separate parameters and consists of two linear layers with a ReLU activation in between:

$$FFN(x) = \max(0, xW_1 + b_1) W_2 + b_2$$

3 Unsupervised Approach

Training a ranker typically requires labelling a substantial amount of data and one might have to repeat this process whenever a significant change is made to the system’s behavior. This is not feasible for most developers and motivates us to explore a set of unsupervised approaches. The core idea is that we learn a generative user simulator and have it talk with the bot to produce problematic dialogues. We then train a ranker with seed dialogues used as normal examples. A straightforward approach for problematic dialogue generation is to train the generative user simulator on a dialogue corpus collected from a wide range of different domains, dubbed as `MultiDomain`. However, this approach can only produce obvious problematic dialogues where the simulated user mimics users who barely know what the bot is for.

To generate more relevant dialogues, one can fine-tune the `MultiDomain` model on the seed dialogues, dubbed as `FineTune`. But this approach gains an improved relevance at the cost of decreased diversity and it is a formidable task to adjust fine-tuning to strike the right balance between relevance and diversity.

We observe that, in most cases, a dialogue follows a natural course until a problem occurs and the dialogue subsequently gets off track. Table 1 shows a problematic dialogue. To bring this to our problematic dialogue generation, we introduce a novel stepwise fine-tuning approach, called `StepFineTune`. The idea is that we fine-tune the `MultiDomain` model only up to l -th turn to generate dialogues in which it normally unfolds up to l -th turn and starts seeing problems afterward. As we fine-tune the model in this stepwise fashion from $l = 1$ to n , we accumulate all the dialogues that we generate at each step. This allows us to produce a spectrum of diverse problematic dialogues while controlling relevance.

However, it is widely known that the typical MLE training scheme often generates bland and generic responses (Li et al., 2016). To alleviate this problem and generate naturally diverse dialogues, we propose a novel stepwise GAN training scheme, dubbed as `StepGAN`. `StepGAN` differs




Figure 3: The overall pipeline of the StepGAN approach.

from `StepFineTune` in that it conducts GAN training instead of the simple MLE fine-tuning for each step. As we cast the dialogue ranking task as binary classification in Section 2, the dialogue ranking model ψ lends itself well to discriminating real dialogues from simulated ones. In the rest of this section, we describe `StepGAN` in detail.

3.1 StepGAN overview

Figure 3 shows the overall pipeline of the `StepGAN` approach. A dialogue generator consists of a user simulator and the bot, and have them talk with each other. We start off by pre-training a generative user simulator on a large corpus of dialogues collected from multiple domains which teaches the simulator basic language skills and helps learn diverse out-of-domain behavior. We use the pre-trained user simulator to produce problematic dialogues and pre-train a discriminator with seed dialogues used as normal dialogues.

We then begin stepwise GAN training. At each step, starting from the user simulator of turn $l - 1$ and the pre-trained discriminator, we further train them using GAN to make the first l turns of a generated dialogue less distinguishable from the seed dialogues, as listed in Algorithm 1. To achieve this goal, we truncate both seed and generated dialogues up to length l when we train the user simulator and discriminator. Once the GAN training is done, we generate a set of problematic dialogues \mathcal{D}_{pos}^l using the trained user simulator. Note that we don’t truncate these dialogues so that they may follow a normal course up to l -th turn and start seeing problems afterward². When we finish the final step L , we collect all the problematic dialogues generated from each step and construct a

²A dialogue ends either when the system or user terminates or when the pre-set maximum length is reached.

Algorithm 1 GAN training at step l

Require: Problematic dialogues \mathcal{D}_{pos}

- 1: $\mathcal{D}_{seed}^l \leftarrow$ seed dialogues truncated at turn l
 - 2: $g^l \leftarrow$ user simulator from turn $l - 1$
 - 3: $d^l \leftarrow$ pre-trained discriminator
 - 4: **while** not convergent **do**
 - 5: $\mathcal{D}_{gen}^l \leftarrow$ Truncate(Generate(g^l), l)
 - 6: $d^l \leftarrow$ Train_d(\mathcal{D}_{seed}^l , \mathcal{D}_{gen}^l)
 - 7: $g^l \leftarrow$ Train_g(\mathcal{D}_{gen}^l , d^l)
 - 8: **end while**
 - 9: $\mathcal{D}_{pos}^l \leftarrow$ Generate(g^l)
-

accumulated dataset:

$$\mathcal{D}_{pos} = \{\mathcal{D}_{pos}^1, \dots, \mathcal{D}_{pos}^L\}$$

Finally, we train the dialogue ranker ψ on the accumulated problematic data \mathcal{D}_{pos} and seed dialogues \mathcal{D}_{seed} .

3.2 GAN training details

Based on our empirical study, we choose to use a sequence-to-sequence model with attention for our user simulator. With GAN, at the l -th iteration, we optimize the following objective which basically adjusts the user simulator to fool the discriminator:

$$J(\theta) = \mathbb{E}_{p_{\theta}(\mathbf{D})} \left[d^l([\mathbf{D}^{<l}, U_i]) \right] + \lambda \sum_{i=0}^l H_{\theta}(U_i | \mathbf{D}^{<i}) \quad (1)$$

where \mathbf{D} denotes a generated dialogue and $\mathbf{D}^{<i} = [S_1, U_1, \dots, U_{i-1}, S_i]$. $d^l(\cdot)$ denotes the discriminator being trained at the l -th iteration and returns the probability of $\mathbf{D}^{<l}$ being real, as reward for training the generator. θ is the parameters for the user simulator and $H(\cdot)$ is the entropy penalty (Pereyra et al., 2017) for increasing the generation diversity:

$$H_{\theta}(U_i | \mathbf{D}^{<i}) = \sum_{j=0}^{N_u} H_{\theta}(u_j | \mathbf{D}^{<i}, u^{<j}) \quad (2)$$

where,

$$H_{\theta}(u_j | \mathbf{D}^{<i}, u^{<j}) = - \sum_{m=0}^M p_{\theta}(u_j^m | \mathbf{D}^{<i}, u^{<j}) \cdot \log p_{\theta}(u_j^m | \mathbf{D}^{<i}, u^{<j}) \quad (3)$$

In Eq 2 and 3, N_u is the number of tokens in U_i and M is the vocabulary size. $p_{\theta}(u_j^m | \mathbf{D}^{<i}, u^{<j})$ is the conditional distribution over the vocabulary at time step j in the generation of U_i . Since Eq 1 is not differentiable, we adopt the REINFORCE algorithm (Williams, 1992) for gradient updates:

$$\nabla_{\theta} J(\theta) \propto d^l(U_l, \mathbf{D}^{<l}) \sum_{i=0}^l \nabla_{\theta} \log p_{\theta}(U_i | \mathbf{D}^{<i}) + \sum_{i=0}^l \nabla_{\theta} H_{\theta}(U_i | \mathbf{D}^{<i})$$

To stabilize the learning process, we employ two common techniques: 1) a baseline: we take the average of rewards in each training batch 2) teacher forcing: we occasionally draw a random dialogue from the seed dialogues with $d^l(\cdot)$ set to return 1. To increase the diversity of the output of the user simulator, during inference, we combine sampling with beam search. At each time step j , instead of choosing the top *beam_size* terms, we sample *beam_size* terms according to the probability distribution $p_{\theta}(\mathbf{u}_j | \mathbf{D}^{<i}, u^{<j})$.

Since we cast the dialogue ranking task as binary classification, we use the same architecture as the dialogue ranking model in Section 2 to discriminate seed dialogues from simulated ones. The only difference is that seed and generated dialogues now correspond to positive and negative classes, respectively.

4 Datasets

In this work, we build a log dialogue ranker for the restaurant inquiry bot offered by the PyDial platform.³ The task for the bot is to search for restaurants based on user’s requirements in a multi-turn natural language communication. Three main corpora are introduced: (1) log dialogues with labels, (2) seed dialogues for the restaurant domain, (3) a large corpus of dialogues collected from multiple domains (Lee et al., 2019).

Log dialogues with labels

To collect log dialogues, we deployed the Pydial restaurant bot via the Amazon Mechanical Turk (AMT) platform.⁴ We ask turkers to find

³<http://www.camdial.org/pydial/>

⁴We use the data collection toolkit offered by ParlAI http://www.parlai/static/docs/tutorial_mturk.html.

restaurants that satisfy automatically generated requirements, such as food type, location and price range, by chatting with the restaurant bot. To make the conversation natural, we encourage turkers to speak in natural utterances and do not allow any turkers to carry out more than 20 dialogues in total. At the end of each task, turkers are required to answer a questionnaire whether they found restaurants meeting their requirements, and whether they experienced contextually unnatural turns in the conversation. We control the quality of a turker’s judgements by checking if a turker judges correctly for some obvious cases that we can automatically identify.

From the collected dialogues, we label successful dialogues without any contextually unnatural turns as 0 (normal dialogue), and the rest as 1 (problematic dialogue). Table 2 shows the number and average length of log dialogues. Examples are shown in Table 1.

We split the corpus as shown in Table 3.

Note that, the training and validation sets are used only for supervised training, whereas the test set is used for evaluating all approaches.

Seed dialogues

The corpus of seed dialogues has two use cases: 1) we use it to fine-tune the user simulator for the `FineTune` and `StepFineTune` approaches, 2) `StepGAN` takes it as input to the discriminator training and teacher forcing process. Since the restaurant bot does not have associated seed dialogues, we collect 100 seed dialogues by having the bot talk with the agenda-based user simulator that `Pydial` offers.⁵

Multi-domain dialogues

The multi-domain corpus⁶ has two use cases: 1) we use it for training the user simulator for the `MultiDomain` approach, 2) we pretrain the simulator for the `StepFineTune` and `StepGAN` approaches. The multi-domain corpus consists of around 40,000 dialogues with 11 turns on average. Each dialogue is a task-oriented conversational interaction between two real speakers over 51 domains and 242 tasks, collected by crowd-sourcing in which one turker is simulating a user and the

other one is simulating a chatbot. We preprocess dialogues into training pairs for the sequence-to-sequence model learning. A training pair consists of a dialogue context and the corresponding response. We consider three consecutive turns as dialogue context and the following turn as response. The number of dialogues in train/val/test sets are 38000/1000/1000 and the training pairs are 394821/10411/10370.

5 Experiments

We describe a set of experimental settings in Section 5.1 and evaluation metrics in Section 5.2.

5.1 Experimental settings

We present three experiments to demonstrate the efficacy of our approaches.

- *StepGAN vs. other unsupervised approaches:* To compare the `StepGAN` against other unsupervised approaches introduced in Section 3:
 - *MultiDomain:* We train the user simulator on the multi-domain corpus.
 - *FineTune:* We fine-tune the user simulator on the seed dialogues.
 - *StepFineTune:* We fine-tune the user simulator in a stepwise fashion.

We train dialogue rankers for each approach on a balanced dataset, consisting of 1000 problematic dialogues generated by each approach and 1000 normal dialogues oversampled from the seed dialogues to balance the class distribution.

We evaluate two additional models:

- *DBDC3:* We use the off-the-shelf CRF model trained on open-domain dialogue corpora⁷ for the dialogue breakdown detection challenge (Higashinaka et al., 2017). Since the `DBDC3` model operates at the utterance level, we aggregate the utterance scores to obtain a dialogue-level score. Based on our experiments, we choose to use a model trained on the collection of four challenge datasets (`CIC_115`, `IRIS_100`, `TKTK_100`, `YI_100`) and aggregate the utterance scores through summation.
- *PlainMultiDomain:* Instead of training a user simulator, we simply take the entire multi-domain corpus (38,000 dialogues) as problematic dialogues and train the ranker with the same number of normal dialogues oversampled from the seed dialogues.

⁵We collected 100 seed dialogues based on our observations that most developers start training their bots with a seed dialogue corpus on a similar scale.

⁶The multi-domain corpus (`MetaLWOz`) will be made available through a `DSTC8` track (Lee et al., 2019).

⁷<https://github.com/dbd-challenge/dbdc3>

One normal example in the labelled corpus	
SYS: hello, welcome to the restaurant finding system. how may i help you?	
USER: hi i ' m looking for a restaurant.	
SYS: what part of town do you have in mind?	
USER: the east part of town please.	
SYS: what kind of food would you like?	
USER: i'd like chinese food please.	
SYS: yu garden is a nice place, it serves chinese food and it is in the east part of town.	
USER: what is the phone number?	
SYS: yu garden is a nice place their phone number is 0 1223 248882.	
USER: what is the address?	
SYS: yu garden is a nice place their address is 529 newmarket road fen ditton.	
USER: thank you	
SYS: thank you, goodbye.	
One problematic example in the labelled corpus	
SYS: hello, welcome to the restaurant finding system. how may i help you?	
USER: hello i would really like to eat some cheap indian tonight	
SYS: what part of town do you have in mind?	
USER: i would really like to eat in the centre of town	
SYS: i am sorry but there is no place, in the expensive price range, serving corsica food and in the area of centre.	
USER: i would really like so indian thats cheap in the centre of town	
SYS: the gandhi is a nice place, it is in the cheap price range, it is in the centre of town and it serves indian food.	
USER: thats great, have you got a contact number?	
SYS: what part of town do you have in mind?	
USER: the centre of town	
SYS: efes restaurant is a nice place, it is in the centre of town, it is in the moderate price range and it serves turkish food.	
USER: i want to eat at the gandhi have you got the address?	
SYS: the gandhi is a nice place their address is 72 regent street city centre.	
USER: thank you, have you got a phone number for the gandhi?	
SYS: the gandhi is a nice place their phone number is 0 1223 353942.	
USER: thanks	
SYS: chiquito bar is a nice place, it serves mexican food in the expensive price range and it is in the south part of town.	
USER: quit	

Table 1: One normal and one problematic examples in the labelled corpus.

	Normal dials	Problematic dials
Num of dials	1621	1329
Avg length	5.87 turns	8.28 turns

Table 2: The number of dialogues and average dialogue length in the log dialogue corpus.

	Train	Val	Test
Normal dialogues	1321	100	200
Problematic dialogues	1029	100	200

Table 3: Log dialogues are randomly sampled and split into training, validation and test sets.

- *StepGAN vs. supervised approach:* One advantage for the unsupervised approach is to generate a larger training set without human efforts. To compare the best performance of the StepGAN model against the supervised model, we train StepGAN models on three balanced datasets with 2000, 4000 and 6000 dialogues each and compare them to a supervised dialogue ranker trained on a balanced dataset of 2000 labelled dialogues randomly sampled from the training set described in Section 4.
- *Semi-supervised learning attempts:* On top of the labelled data, we can employ unsupervised approaches for data augmentation. For the eval-

uation of such a semi-supervised setting, we compare the performance of supervised models with 500 and 2000 labelled examples and that of their counterparts which leverage additional 6000 examples generated by StepGAN.

Note that, all dialogue rankers are tested on the 400-instance balanced test set described in Table 3. We train 10 models on randomly sampled training sets and report average performance.

5.2 Evaluation metrics

We use ranking metrics for evaluation:

- $P@K$ – Precision at k , corresponds to the number of problematic dialogues in the top k ranked options.
- $R@K$ – Recall at k , corresponds to the number of problematic dialogues in the top k ranked options against the number of all problematic dialogues in the test set (i.e. 200). Note that we modified the standard of Recall at k to get monotonic increase with respect to k .

6 Results and Discussion

In this section, we first present the results for the experimental settings in Section 5.1 that we de-




Figure 4: Precision@k StepGAN vs. other unsupervised approaches.




Figure 5: Precision@k StepGAN with different training size vs. supervised approach.




Figure 6: Precision@k of semi-supervised StepGAN vs. supervised settings.




Figure 7: Average length of generated dialogues in each training step.




Figure 8: Success rates of generated dialogues in each training step.




Figure 9: Average score of generated dialogues in each training step.

Model	DB	PM	MD	FT	SF	SG	Sup
P@10	.700	.750	.911	.767	.844	.889	.850
P@20	.800	.775	.867	.794	.850	.900	.870
P@30	.800	.777	.811	.807	.859	.896	.887
P@40	.825	.763	.792	.814	.867	.886	.878
P@50	.800	.738	.762	.800	.836	.864	.870
P@100	.720	.617	.680	.758	.758	.773	.827
P@150	.713	.521	.641	.734	.738	.749	.799
P@200	.655	.499	.612	.688	.702	.708	.765
R@10	.035	.038	.046	.038	.042	.044	.043
R@20	.080	.077	.087	.079	.085	.090	.087
R@30	.120	.117	.122	.121	.129	.134	.133
R@40	.165	.153	.158	.163	.173	.177	.176
R@50	.200	.185	.191	.200	.209	.216	.218
R@100	.360	.308	.340	.379	.379	.387	.414
R@150	.535	.391	.481	.551	.553	.562	.600
R@200	.655	.499	.612	.688	.702	.708	.765

Table 4: Evaluation results. DB, PM, MD, FT, SF and SG stand for the DBDC3, PlainMultiDomain, MultiDomain, FineTune, StepFineTune and StepGAN approach, respectively. The Sup denotes the supervised approach trained on the balanced labelled dialogues.

signed to study the efficacy of our unsupervised approaches. Then, we analyze the generated dialogues to test if StepGAN can generate reasonably problematic dialogues. Table 4 shows the overall results in Precision@k and Recall@k.

6.1 Comparative results

StepGAN vs. other unsupervised approaches:

Figure 4 shows that StepGAN outperforms

other unsupervised approaches by a large margin. The only exception is MultiDomain when $k = 10$. One noteworthy observation is made by comparing MultiDomain with FineTune – MultiDomain is more effective than FineTune when k is small, less than 30 in this case. This is because most turns are contextually wrong or unnatural when we look into the most problematic dialogues (e.g. $k < 10$) and MultiDomain generates exactly such dialogues. On the contrary, as k increases, generated dialogues gradually resemble normal ones with less wrong turns and FineTune essentially generates this type of dialogues.

This observation indicates that a high-quality model should be able to generate a spectrum of dialogues, ranging from obvious dialogues to subtle ones. That is why we introduced a stepwise training scheme and thus StepFineTune shows a significantly improved performance. Furthermore, StepGAN employs the GAN training procedure to generate more natural and diverse dialogues and almost always outperforms all other unsupervised approaches. The fact that StepGAN outperforms DBDC3 and PlainMultiDomain highlights that the StepGAN’s capability of generating domain-relevant problematic dia-

logues is crucial in obtaining high performance.

StepGAN vs. supervised approach: In Figure 5, *StepGAN2000*, *StepGAN4000* and *StepGAN6000* denote ranking models trained on 2000, 4000, 6000 balanced datasets generated by StepGAN respectively. *Sup* stands for a ranker trained on 2000 balanced labelled dialogues. Interestingly, StepGAN performs even better than the supervised approach when $k < 50$. Even though the supervised approach yields higher performance when k is large, StepGAN still compares favorably and the gap is narrower if more dialogues are generated. Note that having developers review a large number of log dialogues (over 100) induces a significant cognitive load. Thus, the higher performance of StepGAN in the small k regime can offer more practical value.

Semi-supervised learning attempts: In Figure 6, *Sup500* and *Sup2000* denote supervised dialogue rankers trained on randomly sampled 500 and 2000 balanced labelled dialogues, respectively. *StepGAN+Sup500* and *StepGAN+Sup2000* denote semi-supervised approaches trained on the 500 and 2000 labelled datasets plus 6000 simulated dialogues generated by StepGAN, respectively. The higher performance of the semi-supervised approaches compared to the supervised counterparts highlights that our unsupervised approach can bring additional generalization by simulating a wide range of dialogues that are not covered by labelled data. As expected, the performance gain increases as we move to a smaller data regime, e.g. 500 labelled dialogues.

6.2 Analysis on generated dialogues

To investigate how generated dialogues move toward normal dialogues, we examine dialogues generated at each step of StepGAN training in terms of three quantitative metrics: average dialogue length, task success rate and ranking score produced by the supervised ranker. Figure 7, 8 and 9 clearly show that as training progresses, the characteristics of generated dialogues are getting more similar to normal dialogues – average dialogue length of generated dialogues is decreasing while task success rate and average ranking score are increasing. In Table 5, we show example dialogues generated by StepGAN after two different training steps to qualitatively examine generated

dialogues. The generated dialogues clearly show that StepGAN manages to follow a normal course up to the training-step turn and then start to get off track. For the full dialogue examples generated by different generative models, refer to Supplemental Material A.

7 Related Work

In contrast to conventional approaches (Jokinen and McTear, 2009), end-to-end neural models have been recently investigated for task-oriented dialogue systems which allows for directly learning dialogue systems from human-human dialogue data (Wen et al., 2016; Bordes et al., 2016; Williams et al., 2017). For such approaches, our proposed dialogue ranking method can help developers quickly go over log dialogues to improve their dialogue systems. There are prior studies on dialogue quality assessment. Walker et al. (1997); Engelbrecht et al. (2009) investigate automatic ways of dialogue evaluation. Walker et al. (2002); Herm et al. (2008); Kim (2007); Higashinaka et al. (2010, 2016) focus on dialogue breakdown detection during the dialogue. Ultes et al. (2012); Su et al. (2016) estimate dialogue quality to improve dialogue policy. However, none has directly dealt with log dialogue ranking task and offered unsupervised approaches to reduce labelling cost. For automatic error handling, Clark (1996); Bohus and Rudnicky (2005); Higashinaka et al. (2015) examined different types of errors. Whereas, for our work, it suffices to rank problematic dialogues higher than normal ones so that developers may take appropriate measures based on the goal of their bots, e.g., adding new slots or system actions to handle previously uncovered scenarios.

8 Conclusions and Future Work

We proposed unsupervised approaches for the task of *log dialogue ranking* and provided comprehensive experimental results. StepGAN outperforms other unsupervised approaches and compares favorably with the supervised approach. Moreover, StepGAN improves upon the supervised approach in semi-supervised settings. Finally, we release the labelled dialogue corpus and seed dialogues. Future work includes 1) in-depth human evaluation on the quality of generated dialogues, 2) testing StepGAN on a broader range of tasks 3) performing dialogue clustering before ranking to reduce redundant workload for developers.

Dialogues simulated by the StepGAN approach in different training steps

StepGAN training at step 1

SYS: hello, welcome to the restaurant finding system. how may i help you?
USER: hi, i want to book a table
SYS: what part of town do you have in mind?
USER: boston.,
SYS: what part of town do you have in mind?
USER: part is going up in colorado.
SYS: what part of town do you have in mind?
USER: i do not want to book a movie price range and in the usa in the west of town.
SYS: i am sorry but there is no place, serving north american food and in the area of west.
USER: alright, that's understandable, can you tell me more?
SYS: i am sorry but there is no place, in the area of west and serving north american food.
USER: that's sound great. oh that's the limit. that is all i need as well though, can you help
SYS: thank you, goodbye.

StepGAN training at step 3

SYS: hello, welcome to the restaurant finding system. how may i help you?
USER: i need another restaurant.
SYS: what part of town do you have in mind?
USER: i want a venue in the moderate price range and in the north of town.
SYS: what kind of food would you like?
USER: i need a indian venue.
SYS: the nirala is a nice place, it is in the moderate price range, it serves indian food and it is in the north part of town.
USER: i am not looking for this. please change.
SYS: the nirala is a nice place, it is in the moderate price range, it serves indian food and it is in the north part of town.
USER: the account 555 - venue is great. i am in the moderate price range and i am going at town.
SYS: the nirala is a nice place, it serves indian food, it is in the moderate price range and it is in the north part of town.
USER: well, i want a venue in the business site. please do, i am going to get a brief globe.
SYS: the nirala is a nice place, it serves indian food, it is in the moderate price range and it is in the north part of town.
USER: that doesn't matter.
SYS: the nirala is a nice place, it is in the north part of town, it is in the moderate price range and it serves indian food.
USER: okay, thanks for the help me book
SYS: thank you, goodbye.

StepGAN training at step 6

SYS: hello, welcome to the restaurant finding system. how may i help you?
USER: i need another restaurant.
SYS: what part of town do you have in mind?
USER: it doesn't matter.
SYS: what part of town do you have in mind?
USER: it doesn't matter.
SYS: what part of town do you have in mind?
USER: i want a venue in the expensive price range and in the centre of town.
SYS: what kind of food would you like?
USER: i need a european venue.
SYS: michaelhouse cafe is a nice place, it is in the centre of town,
it is in the expensive price range and it serves european food.
USER: what is the phone number of the restaurant?
SYS: michaelhouse cafe is a nice place their phone number is 0 1223 309147.
USER: thanks for your help, goodbye.
SYS: thank you, goodbye.

StepGAN training at step 9

SYS: hello, welcome to the restaurant finding system. how may i help you?
USER: i need a european restaurant.
SYS: what part of town do you have in mind?
USER: i prefer a venue in centre.
SYS: would you like something in the cheap, moderate, or expensive price range?
USER: i want a venue in the expensive price range.
SYS: michaelhouse cafe is a nice place, it is in the expensive price range,
it is in the centre of town and it serves european food.
USER: what is the address of michaelhouse cafe?
SYS: michaelhouse cafe is a nice place their address is st. michael's church trinity street city centre.
USER: thanks for your help, goodbye.
SYS: thank you, goodbye.

Table 5: Dialogues simulated by the StepGAN approach in different training steps.

References

- Dan Bohus and Alexander I Rudnicky. 2005. Sorry, i didn't catch that!-an investigation of non-understanding errors and recovery strategies. In *6th SIGdial workshop on discourse and dialogue*.
- Antoine Bordes, Y-Lan Boureau, and Jason Weston. 2016. Learning end-to-end goal-oriented dialog. *arXiv preprint arXiv:1605.07683*.
- Herbert H Clark. 1996. *Using language*. Cambridge university press.
- Klaus-Peter Engelbrecht, Florian Gödde, Felix Hartard, Hamed Ketabdar, and Sebastian Möller. 2009. Modeling user satisfaction with hidden markov models. In *Proceedings of the SIGDIAL 2009 Conference*, pages 170–177.
- Ota Herm, Alexander Schmitt, and Jackson Liscombe. 2008. When calls go wrong: How to detect problematic calls based on log-files and emotions? In *Ninth Annual Conference of the International Speech Communication Association*.
- Ryuichiro Higashinaka, Kotaro Funakoshi, Masahiro Araki, Hiroshi Tsukahara, Yuka Kobayashi, and Masahiro Mizukami. 2015. Towards taxonomy of errors in chat-oriented dialogue systems. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 87–95.
- Ryuichiro Higashinaka, Kotaro Funakoshi, Michimasa Inaba, Yuiko Tsunomori, Tetsuro Takahashi, and Nobuhiro Kaji. 2017. Overview of dialogue breakdown detection challenge 3. *Proceedings of Dialog System Technology Challenge*, 6.
- Ryuichiro Higashinaka, Kotaro Funakoshi, Yuka Kobayashi, and Michimasa Inaba. 2016. The dialogue breakdown detection challenge: Task description, datasets, and evaluation metrics. In *LREC*.
- Ryuichiro Higashinaka, Yasuhiro Minami, Kohji Dohsaka, and Toyomi Meguro. 2010. Modeling user satisfaction transitions in dialogues from overall ratings. In *Proceedings of the 11th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 18–27. Association for Computational Linguistics.
- Kristiina Jokinen and Michael McTear. 2009. Spoken dialogue systems. *Synthesis Lectures on Human Language Technologies*, 2(1):1–151.
- Woosung Kim. 2007. Online call quality monitoring for automating agent-based call centers. In *Eighth Annual Conference of the International Speech Communication Association*.
- Sungjin Lee, Hannes Schulz, Adam Atkinson, Jianfeng Gao, Kaheer Suleman, Layla El Asri, Mahmoud Adada, Minlie Huang, Shikhar Sharma, Wendy Tay, and Xiujun Li. 2019. **Multi-domain task-completion dialog challenge**.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. **A diversity-promoting objective function for neural conversation models**. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119. Association for Computational Linguistics.
- Gabriel Pereyra, George Tucker, Jan Chorowski, Łukasz Kaiser, and Geoffrey Hinton. 2017. Regularizing neural networks by penalizing confident output distributions. *arXiv preprint arXiv:1701.06548*.
- Pei-Hao Su, Milica Gasic, Nikola Mrksic, Lina Rojas-Barahona, Stefan Ultes, David Vandyke, Tsung-Hsien Wen, and Steve Young. 2016. On-line active reward learning for policy optimisation in spoken dialogue systems. *arXiv preprint arXiv:1605.07669*.
- Stefan Ultes, Alexander Schmitt, and Wolfgang Minker. 2012. Towards quality-adaptive spoken dialogue management. In *NAACL-HLT Workshop on Future Directions and Needs in the Spoken Dialog Community: Tools and Data*, pages 49–52. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Marilyn A Walker, Irene Langkilde-Geary, H Wright Hastie, Jerry Wright, and Allen Gorin. 2002. Automatically training a problematic dialogue predictor for a spoken dialogue system. *Journal of Artificial Intelligence Research*, 16:293–319.
- Marilyn A Walker, Diane J Litman, Candace A Kamm, and Alicia Abella. 1997. Paradise: A framework for evaluating spoken dialogue agents. *arXiv preprint cmp-lg/9704004*.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrksic, Milica Gasic, Lina M Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2016. A network-based end-to-end trainable task-oriented dialogue system. *arXiv preprint arXiv:1604.04562*.
- Jason D Williams, Kavosh Asadi, and Geoffrey Zweig. 2017. Hybrid code networks: practical and efficient end-to-end dialog control with supervised and reinforcement learning. *arXiv preprint arXiv:1702.03274*.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.

Supplemental Material

A Dialogues simulated by unsupervised approaches

Dialogues simulated by the MultiDomain approach
SYS: hello, welcome to the restaurant finding system. how may i help you?
USER: i am looking for cheap french cuisine
SYS: i'm having trouble understanding what you want. could you please rephrase it?
USER: i want a champagne without a different tables
SYS: midsummer house restaurant is nice
USER: where is the restaurant?
SYS: rajmahal is nice
USER: how about champagne's?
SYS: the good luck chinese food takeaway is nice
USER: ohh okay. that's the problem! i really'm craving midsummer vegas.
SYS: what part of town do you have in mind?
USER: the galleria vegas adventure
SYS: galleria is a nice place it is in the centre of town.
USER: awesome, i'll go with that topic to serve, need else
SYS: galleria is a nice place it is in the moderate price range.
USER: , my account is very good i guess
SYS: galleria is a nice place it serves european food.
USER: i will thanks anyways
SYS: thank you, goodbye.

Dialogues simulated by the FineTune approach
SYS: hello, welcome to the restaurant finding system. how may i help you?
USER: hi, i want to book a table
SYS: what part of town do you have in mind?
USER: i prefer a venue in south.
SYS: what kind of food would you like?
USER: i need a chinese venue.
SYS: would you like something in the cheap, moderate, or expensive price range?
USER: it doesn't matter.
SYS: would you like something in the cheap, moderate, or expensive price range?
USER: i want a venue in the expensive price range.
SYS: the good luck chinese food takeaway is a nice place, it is in the south part of town, it is in the expensive price range and it serves chinese food.
USER: thanks for your help, goodbye.
SYS: thank you, goodbye.

Table 6: Dialogues simulated by the MultiDomain approach/FineTune approach.

Dialogues simulated by the StepFineTune approach in different training steps

StepFineTune training at step 1

SYS: hello, welcome to the restaurant finding system. how may i help you?

USER: i need another restaurant.

SYS: what part of town do you have in mind?

USER: i need another restaurant.

SYS: what part of town do you have in mind?

USER: washington dc.

SYS: what part of town do you have in mind?

USER: in the west of town.

SYS: what kind of food would you like?

USER: i want a table

SYS: what kind of food would you like?

USER: i need another restaurant.

SYS: what kind of food would you like?

USER: i want restaurant in centre.

SYS: what kind of food would you like?

USER: i want a restaurant in east.

SYS: what kind of food would you like?

USER: i want a restaurant in the moderate price range.

SYS: what kind of food would you like?

USER: i want a restaurant in the cheap price range and in the centre of town.

SYS: what kind of food would you like?

USER: i want a restaurant in the moderate price range.

SYS: what kind of food would you like?

StepFineTune training at step 3

SYS: hello, welcome to the restaurant finding system. how may i help you?

USER: i want a restaurant in the moderate price range.

SYS: what part of town do you have in mind?

USER: i prefer a venue in west.

SYS: what kind of food would you like?

USER: i need a venue in south.

SYS: what kind of food would you like?

USER: i need a international venue.

SYS: i am sorry but there is no place, serving international food, in the area of south and in the moderate price range.

USER: no, i am not looking for this. please show me another one.

SYS: i am sorry but there is no place, serving international food, in the moderate price range and in the area of south.

USER: no, i need a venue in south.

SYS: i am sorry but there is no place, serving international food, in the area of south and in the moderate price range.

USER: what about the venue in south?

SYS: i am sorry but there is no place, in the moderate price range, in the area of south and serving international food.

USER: thanks for the help, goodbye.

SYS: thank you, goodbye.

StepFineTune training at step 6

SYS: hello, welcome to the restaurant finding system. how may i help you?

USER: hi, i want to book a table

SYS: what part of town do you have in mind?

USER: i want a venue in any part of town.

SYS: what kind of food would you like?

USER: i need a asian oriental venue.

SYS: would you like something in the cheap, moderate, or expensive price range?

USER: i want a venue in the expensive price range.

SYS: saigon city is a nice place, it serves asian oriental food, it is in the north part of town and it is in the expensive price range.

USER: what is the phone number of the restaurant?

SYS: saigon city is a nice place their phone number is 0 1223 356555.

USER: thanks for your help, goodbye.

SYS: thank you, goodbye.

StepFineTune training at step 9

SYS: hello, welcome to the restaurant finding system. how may i help you?

USER: hi, i want to book a table

SYS: what part of town do you have in mind?

USER: i prefer a venue in south.

SYS: what kind of food would you like?

USER: i need a chinese venue.

SYS: would you like something in the cheap, moderate, or expensive price range?

USER: i need a restaurant in the expensive price range.

SYS: the good luck chinese food takeaway is a nice place, it is in the south part of town, it serves chinese food and it is in the expensive price range.

USER: thanks for your help, goodbye.

SYS: thank you, goodbye.

Table 7: Dialogues simulated by the StepFineTune approach in different training steps.

Tree-Structured Semantic Encoder with Knowledge Sharing for Domain Adaptation in Natural Language Generation

Bo-Hsiang Tseng,[†] Paweł Budzianowski,[†] Yen-Chen Wu[†] Milica Gašić[‡]

[†]University of Cambridge [‡] Heinrich Heine University Düsseldorf

bht26@cam.ac.uk, gasic@uni-duesseldorf.de

Abstract

Domain adaptation in natural language generation (NLG) remains challenging because of the high complexity of input semantics across domains and limited data of a target domain. This is particularly the case for dialogue systems, where we want to be able to seamlessly include new domains into the conversation. Therefore, it is crucial for generation models to share knowledge across domains for the effective adaptation from one domain to another. In this study, we exploit a tree-structured semantic encoder to capture the internal structure of complex semantic representations required for multi-domain dialogues in order to facilitate knowledge sharing across domains. In addition, a layer-wise attention mechanism between the tree encoder and the decoder is adopted to further improve the model's capability. The automatic evaluation results show that our model outperforms previous methods in terms of the BLEU score and the slot error rate, in particular when the adaptation data is limited. In subjective evaluation, human judges tend to prefer the sentences generated by our model, rating them more highly on informativeness and naturalness than other systems.

1 Introduction

Building open-domain Spoken Dialogue Systems (SDS) remains challenging. This is partially because of the difficulty of collecting sufficient data for all domains and the high complexity of natural language. Typical SDSs are designed based on a pre-defined ontology (Figure 1) which might cover knowledge spanning over multiple domains and topics (Young et al., 2013).

A crucial component of a Spoken Dialogue System is the Natural Language Generation (NLG) module, which generates the text that is finally presented to the user. NLG is especially challeng-

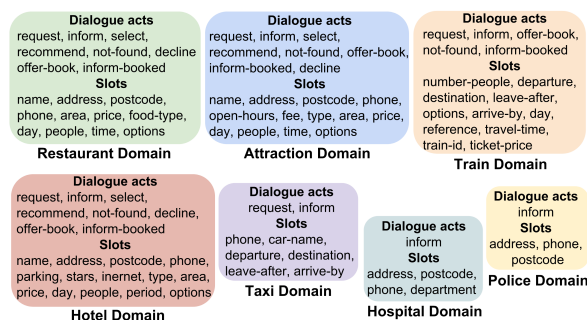


Figure 1: The ontology for multi-domain spoken dialogue systems.

ing when building a multi-domain dialogue systems. Given a semantic representation (SR), the task for NLG is to generate natural language conveying the information encoded in the SR. Typically, an SR is composed of a set of slot-value pairs and a dialogue act consistent with an ontology. A dialogue act represents the intention of the system output and the slots provide domain-dependent information. Figure 2 presents examples of SRs with their corresponding natural language representations in various datasets.

The input semantics has its own hierarchical structure in which there are different sets of slot-value pairs under different dialogue acts across various domains. Modelling the semantic structure might be helpful for sharing information across domains and achieve better performance for domain adaptation. However, prior work encodes semantic representation in a flat way such as using a binary vector (Wen et al., 2015a,b) or using a sequential model such as an LSTM (Dušek and Jurcicek, 2016; Tran and Nguyen, 2017). In that case, the structure of semantics is not fully captured by these encoding methods. This might limit models' performance especially when adapting to a new domain.

This paper investigates the possibility of lever-

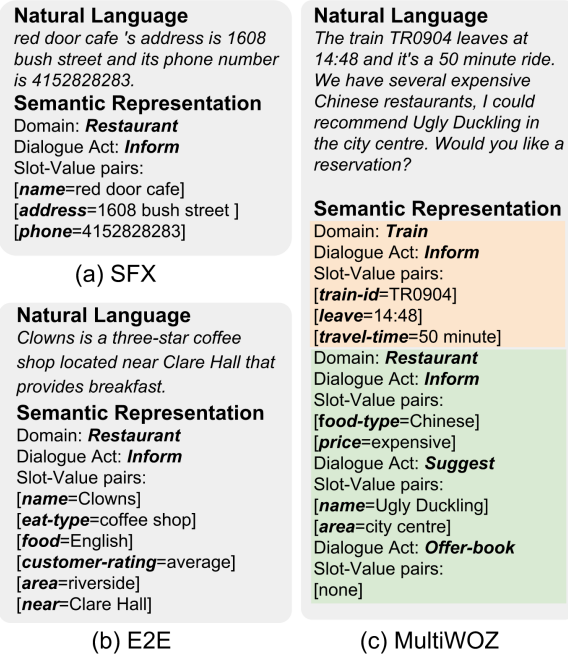


Figure 2: Examples of semantic representations in (a) SFX dataset (Wen et al., 2015b), (b) E2E dataset (Novikova et al., 2017) and (c) MultiWOZ dataset (Budzianowski et al., 2018).

aging the semantic structure for NLG domain adaptation in dialogue systems. We present a generation model with a tree-structured semantic encoder that models the internal structure of the semantic representation to facilitate knowledge sharing across domains. Moreover, we propose a layer-wise attention mechanism to improve the generation performance. We perform experiments on the multi-domain Wizard-of-Oz corpus (MultiWOZ) (Budzianowski et al., 2018) and with human subjects. The results show that the proposed model outperforms previous methods on both automatic metrics and with human evaluation, suggesting that modelling the semantic structure can facilitate domain adaptation. To the best of our knowledge, this work is the first study exploiting the tree LSTM (Tai et al., 2015) to model the input semantics of NLG in spoken dialogue systems.

2 Related Work

Recently, recurrent neural network-based NLG models have shown their powerful capability and flexibility compared to traditional approaches that depend on hand-crafted rules in dialogue systems. A key development was the heuristic gate which turns off the slots that are already generated in the output sentence (Wen et al., 2015a). Subsequently, the semantically conditioned LSTM (SCLSTM)

(Wen et al., 2015b) was proposed with an extra reading gate in the LSTM cell to let the model automatically learn to control the binary representation of the semantics during generation. The sequence-to-sequence (seq2seq) model (Cho et al., 2014; Sutskever et al., 2014) with attention mechanism (Bahdanau et al., 2014) that has achieved huge success in machine translation has also been applied to the NLG task. In (Dušek and Jurcicek, 2016) the slot-value pairs in the semantics were treated as a sequence and encoded by LSTM. Based on the seq2seq model, in (Tran et al., 2017; Tran and Nguyen, 2017) the refinement gate was introduced to modify the input words and hidden states in the decoder by considering the attention result. Different training strategies were studied in prior work. The hierarchical decoding method was proposed by considering the linguistic pattern of the generated sentence (Su et al., 2018). The variational-based model was proposed to learn the latent variable from both natural language and semantics (Tseng et al., 2018). Lampouras and Vlachos (2016) proposed to use imitation learning to train NLG models, where the Locally Optimal Learning to Search framework was adopted to train against non-decomposable loss functions.

Domain adaptation has been widely studied in different areas such as machine translation (Koehn and Schroeder, 2007; Foster et al., 2010), part of speech tagging (Blitzer et al., 2006) and dialogue state tracking (Mrkšić et al., 2015) in spoken dialogue systems. In NLG for spoken dialogue systems, the trainable sentence planner proposed in (Walker et al., 2002; Stent et al., 2004) provides the flexibility of adapting to different domains. Subsequently, generators that can tailor user preferences (Walker et al., 2007) or learn their personality traits (Mairesse and Walker, 2008, 2011; Oraby et al., 2018) were proposed. To achieve multi-domain NLG, exploiting the shared knowledge between domains is important to handle unseen semantics. A multi-step procedure to train a multi-domain NLG model was proposed in (Wen et al., 2016). Adversarial learning is used in (Tran and Nguyen, 2018) in which two critics were introduced during model adaptation.

3 Model

Our generation model is composed of two parts: (a) a tree-structured semantic encoder and (b) an LSTM decoder with additional gates. The tree-

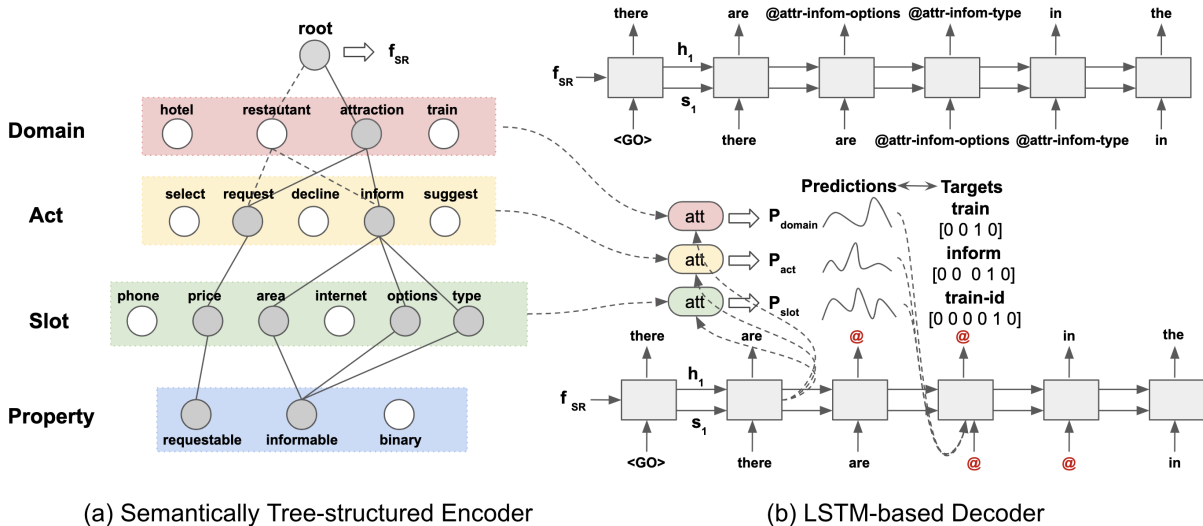


Figure 3: The overview of our generation model. The tree-structured semantic encoder (a) encodes semantic representation to obtain semantic embedding f_{SR} . Each node in the tree denotes a vector representation for that token. Grey node means it is activated during encoding with the corresponding token specified in the semantics. The LSTM-based decoder without layer-wise attention ((b), above) and with layer-wise attention ((b), below) takes f_{SR} as an initial state to generate natural language. The example utterance here is "there are @attraction-inform-options @attraction-inform-type in the @attraction-inform-area, do you have a price range in mind?"

structured semantic encoder extracts a semantic embedding from the semantics in a bottom-up fashion. The obtained embedding is then fed into the decoder as a condition to generate natural language with corresponding delexicalised tokens¹. In addition, we further propose a layer-wise attention mechanism between the tree-structured semantic encoder and the decoder. The proposed attention mechanism further improves the model’s ability to generate the correct information when adapting to a new domain with limited data.

3.1 Tree-Structured Semantic Encoder

There exists a hierarchical relationship between dialogue acts and slot-value pairs within various domains. Inspired by the tree-structured LSTM (Tai et al., 2015) that encodes natural language by capturing its syntactic properties, we propose a tree-structured semantic encoder to encode the semantic representation (SR) by exploiting its internal hierarchy.

3.1.1 Tree Hierarchy

Figure 3 (a) illustrates our tree-structured semantic encoder. The hierarchy of the tree represents

¹Each value in a natural language utterance is replaced by a delexicalised token in the format @domain-act-slot. For instance, the informed restaurant *Golden House* will be replaced by the token @restaurant-inform-name. The mapping from values to delexicalised tokens is called delexicalisation. The inverse process is called lexicalisation.

the ontology with each layer symbolizing a different level of information. At each layer, a node denotes a possible type defined by the ontology. Given an SR, each slot-value pair is associated with a dialogue act (DA) within a domain. This relationship is modelled by the links between different layers in a tree as parents and children. For instance, the node denoting slot name is the child of the node denoting DA *suggest* and DA *suggest* is the child of the node representing domain *restaurant*. In addition, a slot can be *requestable*, *informable* or *binary*. Each of them behaves differently in natural language². Each leaf node denotes a property that describes a slot. As a result, given an SR there is a one-to-one mapping between SR and its corresponding tree and a path from the root to a leaf node describes a slot-value pair along with its domain, DA, slot and property of slot information.

3.1.2 Semantic Representation Encoding

Given a tree representing an SR, each node j of the LSTM contains input, forget and output gates i_j , f_j and o_j respectively to obtain its hidden state and memory cell h_j and c_j . With a set of children $C(j)$, the non-leaf node j has two sources of input:

²For instance, the utterance with a *requestable* slot *area* might be: *Which part of the city you are looking for?*. The utterance with the *informable* slot *area* might be: *There are several restaurants in the @restaurant-inform-area.*

(a) the token embedding e_j ³ and (b) children states h_k, c_k . The transition equations are as following:

$$\begin{aligned}\tilde{h}_j &= \sum_{k \in C(j)} h_k, \\ \tilde{c}_j &= \sum_{k \in C(j)} c_k, \\ i_j &= \sigma(W_E^{(i)} e_j + U_E^{(i)} \tilde{h}_j + b_E^{(i)}), \\ f_j &= \sigma(W_E^{(f)} e_j + U_E^{(f)} \tilde{h}_j + b_E^{(f)}), \\ o_j &= \sigma(W_E^{(o)} e_j + U_E^{(o)} \tilde{h}_j + b_E^{(o)}), \\ g_j &= \tanh(W_E^{(g)} e_j + U_E^{(g)} \tilde{h}_j + b_E^{(g)}), \\ c_j &= i_j \circ g_j + f_j \circ \tilde{c}_j, \\ h_j &= o_j \circ \tanh(c_j),\end{aligned}$$

where k is the children index, \tilde{h}_j and \tilde{c}_j are the sum of children’s hidden states and memory cells respectively.

The semantic embedding is obtained in a bottom-up fashion. Starting from the leaf nodes with their corresponding embeddings, the information is propagated from the property layer through the slot layer, act layer and domain layer to the root. The hidden state at the root is the final semantic embedding f_{SR} for the SR and it will be used to condition the decoder during generation.

During domain adaptation, the model might have seen some semantics in source domain (denoted by dash lines in the tree encoder in Figure 3) that shares a partial tree structure with the semantics in the target domain. For instance, the SR informing about options, type and area in restaurant domain shares partial tree structure with the SR informing about the same information in attraction domain. Modelling semantic structure by the tree encoder benefits knowledge sharing across domains.

3.2 Decoder

Figure 3 (b) presents the LSTM-based decoder with two introduced gates. The representation of the semantics, s_t , is initialised by the semantic embedding f_{SR} and then updated at each time step duration generation. Updating the semantics at each step is crucial to avoiding generating redundant or missing information in the SR. As in standard LSTMs, the transition equations of memory

³All the domains, dialogue acts and slots appearing in an SR are viewed as tokens and encoded in the 1-hot vectors. The 1-hot vectors are then passed through an embedding layer to attain the token embeddings as inputs to the nodes.

cell c_t are as following:

$$\begin{aligned}i_t &= \sigma(W_D^{(i)} x_t + U_D^{(i)} h_{t-1} + b_D^{(i)}), \\ f_t &= \sigma(W_D^{(f)} x_t + U_D^{(f)} h_{t-1} + b_D^{(f)}), \\ o_t &= \sigma(W_D^{(o)} x_t + U_D^{(o)} h_{t-1} + b_D^{(o)}), \\ g_t &= \tanh(W_D^{(g)} x_t + U_D^{(g)} h_{t-1} + b_D^{(g)}), \\ c_t &= i_t \circ g_t + f_t \circ c_{t-1}.\end{aligned}$$

The two introduced gates, reading gate r_t and writing gate w_t , are responsible for updating the semantic state s_t . The reading gate determines what information should be kept from the semantics at previous time step, while the writing gate decides what new information should be added into the current semantic state:

$$\begin{aligned}r_t &= \sigma(W_D^{(r)} x_t + U_D^{(r)} h_{t-1} + V_D^{(r)} s_{t-1} + b_D^{(r)}), \\ w_t &= \sigma(W_D^{(w)} x_t + U_D^{(w)} h_{t-1} + V_D^{(w)} s_{t-1} + b_D^{(w)}), \\ d_t &= \tanh(W_D^{(d)} x_t + U_D^{(d)} h_{t-1} + V_D^{(d)} s_{t-1} + b_D^{(d)}), \\ s_t &= w_t \circ d_t + r_t \circ s_{t-1}.\end{aligned}$$

The hidden state h_t is then defined as the weighted sum of the memory cell and the semantic state with the output gate as weight:

$$h_t = o_t \circ \tanh(c_t) + (1 - o_t) \circ \tanh(s_t).$$

The probability of the word label y_t at each time step t is formed by a applying a softmax classifier that takes the hidden state h_t as input:

$$p(y_t | x_{<t}, f_{SR}) = \text{softmax}(W^{(s)} h_t).$$

The objective function is the standard negative log-likelihood:

$$J(\theta) = - \sum_t \log p(y_t | x_{<t}, f_{SR}). \quad (1)$$

3.3 Layer-wise Attention Mechanism

The semantic embedding obtained from the tree encoder contains high-level information regarding the semantic representation. However, the information in the tree is not fully leveraged during generation. Thanks to the hierarchical structure of a tree encoder with defined meaning for each layer, we can apply an attention mechanism to each layer to let the decoder concentrate on the different levels of information. We expect the decoder to leverage information regarding domain, dialogue act and slot from the hidden states in a tree to influence the generation process.

Whenever the decoder generates the token @⁴, the semantics s_t is used to drive an attention mechanism with hidden states in the different layers of the tree to obtain distributions over domains $p(d_t)$, dialogue acts $p(a_t)$ and slots $p(s_t)$ respectively:

$$p(d_t|x_{<t}, s_t) = \frac{\exp(\text{score}(s_t, h_d))}{\sum_{d' \in D} \exp(\text{score}(s_t, h_{d'}))},$$

$$p(a_t|x_{<t}, s_t) = \frac{\exp(\text{score}(s_t, h_a))}{\sum_{a' \in A} \exp(\text{score}(s_t, h_{a'}))},$$

$$p(s_t|x_{<t}, s_t) = \frac{\exp(\text{score}(s_t, h_s))}{\sum_{s' \in S} \exp(\text{score}(s_t, h_{s'}))},$$

where h_d , h_a and h_s are the hidden states of domain, dialogue act and slots in the tree encoder. D , A and S are the sets of domains, dialogue acts and slots defined in the ontology respectively. The score function used to calculate the similarity between two vectors is defined as following:

$$\text{score}(f, h) = f^T h.$$

The distributions $p(d_t)$, $p(a_t)$ and $p(s_t)$ are then used to predict domain, dialogue act and slot at time step t by taking the argmax operation to form the delexicalised tokens @domain-act-slot back into the generated sentence.

In order to avoid generating redundant or missing information in a given SR, the three predicted distributions are fed into next time step to augment the original input word⁵ to condition the model on what information has already been generated.

During training, the error signals between predicted distributions and the true labels for domain, dialogue act and slot are added to the objective function. The objective function for the generation model with layer-wise attention mechanism is defined as following:

$$J_{att}(\theta) = J(\theta) - \sum_{t'} (\log p(d_{t'}|x_{<t'}, s_{t'})) \\ + \log p(a_{t'}|x_{<t'}, s_{t'}) + \log p(s_{t'}|x_{<t'}, s_{t'}),$$

where $J(\theta)$ is the original objective function in equation 1 and t' is the index for the time step where each token @ is generated.

⁴With the layer-wise attention mechanism, all values in the natural language are replaced by the same delexicalised token @ instead of the tokens in the format @domain-act-slot, and the corresponding information regarding domain, dialogue act and slot will be used as signals to guide the decoder to predict the correct information.

⁵Only at the next time step of generating delexicalised token @ the input is the concatenation of the word vector x_t and three predicted distributions. In any other time steps, the input is the word vector padded with zeros.

Table 1: The data statistics for each domain.

Domain	Restaurant	Hotel	Attraction	Train	Taxi
Examples	8.5k	6.6k	6.4k	11k	3.4k
Distinct SR	346	378	314	338	47
Dialogue acts	8	8	8	5	2
Slots	11	14	13	11	6

4 Experimental Results

4.1 Dataset

We perform our experiments with the Multi-Domain Wizard-of-Oz (MultiWOZ) dataset (Budzianowski et al., 2018) that is a rich dialogue dataset spanning over 7 domains. There are 10438 dialogues and over 115k turns in total. The dataset contains a high level of complexity and naturalness which is suitable for developing multi-domain NLG models. There are multiple utterances in a single turn with an average of 18 words, 1.6 dialogue acts and 2.9 slots per turn. Some turns provide information for more than 1 domain. Comparing with previous NLG datasets which contain only 1 utterance in a turn with 1 dialogue act within 1 domain, the MultiWOZ dataset provides significantly more complexity and makes NLG more challenging. The number of examples, distinct semantic representation (SR) and numbers of dialogue acts and slots are reported in Table 1. The data split for train, dev and test is 3:1:1. The details of the ontology is presented in Figure 1.

4.2 Experimental Setup

The generators are implemented using the Pytorch library (Paszke et al., 2017). Our code is public⁶. The number of hidden units in the LSTMs is 100 with 1 hidden layer. The dropout rate is 0.25 and the Adam optimizer is used. The learning rate is 0.0025 for the models trained from scratch, and 0.001 for the models adapted from one domain to another in adaptation experiments. Beam search is used during decoding with a beam size 10. For automatic metrics, the BLEU scores and the slot error rate (SER) used in (Wen et al., 2015b) are reported. The SER is used to evaluate how accurate a generated sentence is in terms of conveying the desired information in the given semantic representation (SR). The SER is defined as: $(p + q)/N$, where p, q are the numbers of missing and redundant slots

⁶<https://github.com/andy194673/TreeEncoder-NLG-Dialogue>

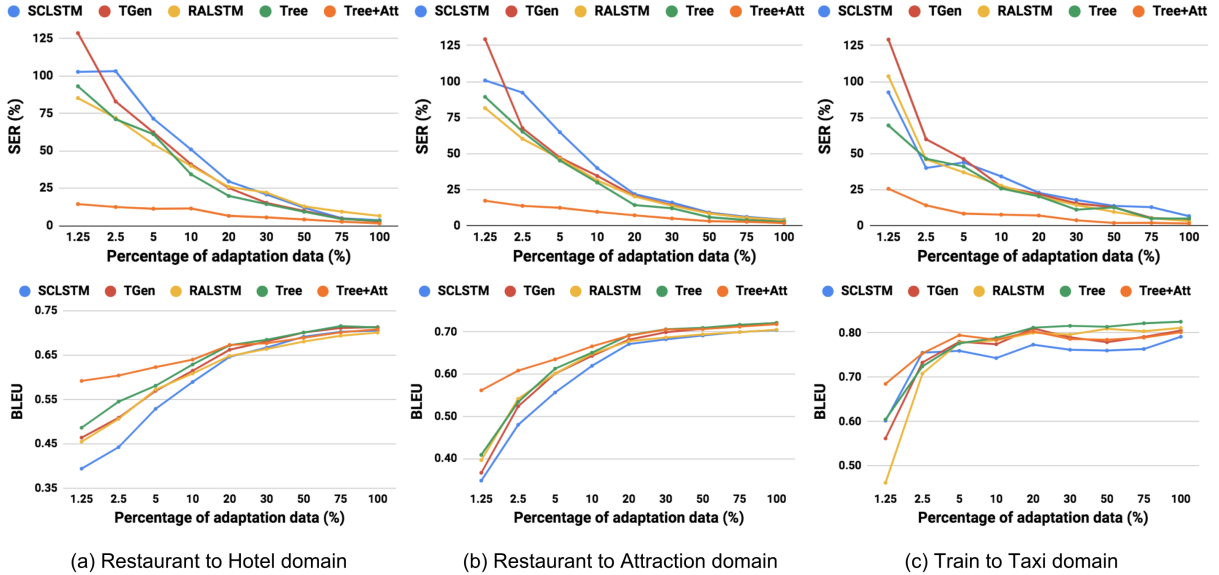


Figure 4: Domain adaptation experiments in three different settings. (a) adapting to hotel from restaurant domain. (b) adapting to attraction from restaurant domain. (c) adapting to taxi from train domain.

in a generated sentence, and N is the number of total slots that a generated sentence should contain. The results are averaged over 10 samples and 5 random initialised seeds. As explained above each delexicalised slot token in an utterance is in the format of @domain-act-slot. When calculating the SER, the predicted slot token is correct only if its domain, dialogue act and slot information are all correct. For example, if there is a desired slot area under dialogue act `inform` within restaurant domain in SR, the model needs to generate the token @restaurant-inform-area.

The tree-structured semantic encoder (Tree) and the variant with attention (Tree+Att) are compared against three baselines: (1) the semantically-conditioned LSTM (SCLSTM) that has an extra gate to update the binary vector of the semantic representation (Wen et al., 2015b); (2) TGen that is a seq2seq model with attention mechanism mapping SR into a word sequence (Dušek and Jurcicek, 2016); (3) a refinement adjustment LSTM (RALSTM) that is an improved seq2seq model with a refinement gate and an adjustment gate in the decoder (Tran and Nguyen, 2017).

As the decoding method is slightly different between our model Tree+Att and baseline models⁷, in order to guarantee the optimised baseline systems, we also trained baseline models in the same

⁷Tree+Att only generates token @ and reply on attention results to form the complete slot token while baseline models directly generate slot tokens.

decoding way as Tree+Att to only predict @ with three additional classifiers for domain, act and slot prediction. However, baseline models obtain better performance by the original decoding method so we keep that in the following experiments. All the models are optimized by selecting the best one based on the validation set result.

4.3 Automatic Evaluation

In order to examine the models' ability to share knowledge between domains, we performed experiments in three domain adaptation scenarios: (a) adapting to hotel from restaurant domain; (b) adapting to attraction from restaurant domain and (c) adapting to taxi from train domain. The adaptation models were fine-tuned with adaptation data based on the models trained on source domain⁸. The SER results are presented in the first row of Figure 4. Generally, our model without attention (Tree) performs similarly with RALSTM but better than TGen and SCLSTM. With the layer-wise attention mechanism, our model (Tree+Att) improves significantly and performs better than baselines at all different levels of adaptation data amount. Especially when the adaptation data used is only 1.25%, the SER is reduced from above 75% to around 25%. We found that this is because baseline models tend to predict the slots with the wrong dialogue act or in the wrong domain as the

⁸All the multi-domain turns are removed in case the model have seen any examples related to target domain before adaptation.

Table 2: Human evaluation for utterance quality in three adaptation settings: Restaurant (Rest.) to Hotel domain; Restaurant to Attraction (Attr.) domain and Train to Taxi domain. Informativeness (Info.) and Naturalness (Nat.) are reported (rating out of 5).

Model	Rest. to Hotel		Rest. to Attr.		Train to Taxi	
	Info.	Nat.	Info.	Nat.	Info.	Nat.
SCLSTM	2.96	3.85	2.81	3.69	3.05	4.26
TGen	2.87	3.33	3.00	3.23	3.42	3.90
RALSTM	2.79	3.48	2.91	3.40	3.48	3.15
Tree	3.08	3.54	3.38	3.41	3.81	3.81
Tree+Att	4.04	4.10	4.30	3.92	4.29	3.78

limited adaptation data makes it difficult to learn the sentence pattern in the target domain. However, with the layer-wise attention mechanism, our model is able to pay attention on the information at different levels in the tree to make the correct predictions. (See more details in section 5 with error analysis and visualisation of attention distributions.) A similar trend can be observed in the BLEU results in the second row of Figure 4.

4.4 Human Evaluation

Because automatic evaluation such as BLEU may not consistently agree with human perception (Stent et al., 2005), we performed human testing via the Amazon Mechanical Turk service. We showed MTurk workers the generated sentences in adaptation experiments with adaptation data from 1.25% to 10% as we focus on the models’ performance with limited adaptation data. Five models were compared together by showing, for each model, the 2 sentences with the highest probabilities out of the 10 generated sentences by beam search. The workers were asked to score each sentence from 1 (bad) to 5 (good) in terms of its informativeness and naturalness. The *informativeness* is defined as the degree to which the generated sentence contains all the information specified in the given semantic representation (SR) without conveying extra information and the *naturalness* is defined as whether the sentence is natural like human language. Ipeirotis et al. (2010) pointed out that malicious workers might take advantage of the difficulty of verifying the results and therefore submit answers with low quality. In order to filter out submissions with bad quality, we also asked them to score the ground truth sentence and an artificial sentence containing irrelevant information to the SR. If the worker gave ground truth sentence a low score (< 3) or gave the artificial sentence a high score (> 3) in terms of informativeness, the

submission was discarded.

The results pertaining to informativeness and naturalness are reported in Table 2 in three adaptation settings: Restaurant (Rest.) to Hotel domain; Restaurant to Attraction (Attr.) domain and Train to Taxi domain. For informativeness, our models (both Tree+Att & Tree) outperform all baseline models in the different settings. This result is consistent with the slot error rate of the automatic evaluation reported in Figure 4 and indicates that the tree-structured semantic encoder does help the model to produce utterances with the correct information. For naturalness, Tree+Att performs the best in two settings, while SCLSTM performs better when adapting to taxi domain. This might be because SCLSTM is good at generating utterances with simple patterns and the taxi domain is relatively easy due to its low number of combinations of SR⁹. When adapting to more complex domains such as hotel or attraction, our models provide both informative and natural utterances. Table 3 presents example semantic representations with corresponding ground truth sentence and the top-1 utterance generated by each model.

5 Error Analysis and Observation

In order to investigate what type of testing data our model performs better on, we divide all test set into two subsets - *seen* and *unseen*. If the semantics of a testing example appear in the training set, the example is defined as *seen*. Otherwise, the example is marked as *unseen*. Table 4 reports the number of seen and unseen examples and the number of wrong utterances (at least 1 missing or redundant slot) generated by each model with different amount of adaptation data when adapting from restaurant to hotel domain. With more adaptation data, more SRs of testing examples appear in the training set. We observe that our model obtains better generalisation ability for unseen SRs. For instance, with 1.25% adaptation data, Tree+Att generates 134 wrong utterances out of 902 unseen semantics (14.8%). However, the baseline models such as SCLSTM produces 729 wrong sentences out of 902 semantics (80.5%). We hypothesize that our model is more capable of learning sentence patterns from source domain and generate correct content for domain adaptation. For example, when adapting from restaurant to hotel domain (see Table 3 - Hotel column), Tree+Att cor-

⁹There are only 2 dialogue acts and 6 slots in taxi domain.

Table 3: Example semantic representations (SR) with ground truth sentences in two adaptation settings with 1.25% adaptation data and the top-1 sentences generated by each model. Both are adapted from restaurant domain. The slot-value pairs are in bold. Each generated sentence is followed by a brief description to explain if the sentence correctly conveys the information in the SR.

Domain	Attraction	Hotel
Semantic Representation	Dialogue Act: Inform Slot-Value pairs: [Area: west] [Options: five] [Type: colleges] Dialogue Act: Request Slot-Value pairs: [Price=?]	Dialogue Act: Inform Slot-Value pairs: [Options= two] Dialogue Act: Select Slot-Value pairs: [type1= guesthouse] [type2= hotel]
Ground Truth	<i>there are five colleges in the west. do you mind paying an entrance fee ?</i>	<i>i have two, would you prefer a guesthouse or hotel ?</i>
SCLSTM	<i>what type of place are you looking for ? (miss 3 slots & request wrong)</i>	<i>what area would you like to stay in ? (miss 3 slots)</i>
TGen	<i>there are located in the . do you have a price range in mind ? (miss 3 slots)</i>	<i>i have found options. would you prefer or ? (miss 3 slots)</i>
RALSTM	<i>we have five colleges in the west area . do you have an attraction type in mind ? (request wrong)</i>	<i>i have two options. do you have a preference ? (miss 2 slots)</i>
Tree	<i>there are five colleges in the west . do you have an area of town you would prefer ? (request wrong)</i>	<i>i have found two options for you. do you have a preference ? (miss 2 slots)</i>
Tree+Att	<i>there are five colleges in the west . do you have a price range in mind ? (correct)</i>	<i>i have two options for you. would you prefer guesthouse or hotel ? (correct)</i>

Table 4: Error analysis - number of examples in the testing set and the number of wrong generated utterances (at least 1 missing or redundant slot) by each model in different adaptation data scenarios. The testing example is defined as seen if its semantics appears in the training set.


Percentage	1.25%		5%		10%		50%	
	seen	unseen	seen	unseen	seen	unseen	seen	unseen
Testing examples	439	902	858	483	1069	272	1330	11
SCLSTM	248	729	307	412	302	190	111	5
TGen	309	741	176	353	178	168	102	6
Tree+Att	10	134	31	103	60	55	76	3

rectly learns to generalize from the training sentence: "i have two options for you, would you prefer American or Chinese" in restaurant domain. However, SCLSTM fails to produce a similar sentence pattern.

Figure 5 shows the example of visualisation of layer-wise attention distributions over domains, acts and slots generated by the Tree+Att model. The model is confident of generating the correct slot tokens with the distinct peaks indicated by the dark red color in the attention distributions even though the adaptation data used is simply 1.25%.

6 Conclusion and Future Work

This paper investigates the possibility of leveraging internal structure of input semantics for NLG domain adaptation in dialogue systems. The proposed tree-structured semantic encoder is able to



Generated utterance: i have @hotel-inform-options (two) options for you. would you prefer @hotel-select-type (guesthouse) or @hotel-select-type (hotel) ?

Figure 5: The visualisation of the layer-wise attention distributions over domains, acts and slots at each time step k when slot token is generated and the generated utterance with lexicalised values in the parentheses. The color shades signify the attention weight.

capture the structure of semantic representations and facilitate knowledge sharing across domains. In addition, we have proposed a layer-wise attention mechanism between the tree-structured semantic encoder and the decoder to enhance the performance. Our proposed model was evaluated on the complex multi-domain MultiWOZ dataset. The automatic evaluation results show that our model is more efficient in terms of adaptation data usage and outperforms previous methods by reducing the slot error rate up to 50% when the adaptation data is limited. What is more, human judges rate our model more highly than previous methods. Future work will explore a tree encoder exploiting both semantic representation and context information in end-to-end dialogue systems.

Acknowledgments

Bo-Hsiang Tseng is supported by Cambridge Trust and the Ministry of Education, Taiwan. This work was partly funded by an Alexander von Humboldt Sofja Kovalevskaja grant.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- John Blitzer, Ryan McDonald, and Fernando Pereira. 2006. Domain adaptation with structural correspondence learning. In *Proceedings of the 2006 conference on empirical methods in natural language processing*, pages 120–128. Association for Computational Linguistics.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. 2018. Multiwoz-a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026.
- Kyunghyun Cho, Bart Merriënboer, Caglar Gulcehre, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *EMNLP*.
- Ondřej Dušek and Filip Jurčicek. 2016. Sequence-to-sequence generation for spoken dialogue via deep syntax trees and strings. In *The 54th Annual Meeting of the Association for Computational Linguistics*, page 45.
- George Foster, Cyril Goutte, and Roland Kuhn. 2010. Discriminative instance weighting for domain adaptation in statistical machine translation. In *Proceedings of the 2010 conference on empirical methods in natural language processing*, pages 451–459. Association for Computational Linguistics.
- Panagiotis G Ipeirotis, Foster Provost, and Jing Wang. 2010. Quality management on amazon mechanical turk. In *Proceedings of the ACM SIGKDD workshop on human computation*, pages 64–67. ACM.
- Philipp Koehn and Josh Schroeder. 2007. Experiments in domain adaptation for statistical machine translation. In *Proceedings of the second workshop on statistical machine translation*, pages 224–227. Association for Computational Linguistics.
- Gerasimos Lampouras and Andreas Vlachos. 2016. Imitation learning for language generation from unaligned data. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1101–1112.
- François Mairesse and Marilyn Walker. 2008. Trainable generation of big-five personality styles through data-driven parameter estimation. *Proceedings of ACL-08: HLT*, pages 165–173.
- François Mairesse and Marilyn A Walker. 2011. Controlling user perceptions of linguistic style: Trainable generation of personality traits. *Computational Linguistics*, 37(3):455–488.
- Nikola Mrkšić, Diarmuid Ó Séaghdha, Blaise Thomson, Milica Gasic, Pei-Hao Su, David Vandyke, Tsung-Hsien Wen, and Steve Young. 2015. Multi-domain dialog state tracking using recurrent neural networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, volume 2, pages 794–799.
- Jekaterina Novikova, Ondřej Dušek, and Verena Rieser. 2017. **The E2E dataset: New challenges for end-to-end generation**. In *Proceedings of the 18th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, Saarbrücken, Germany. ArXiv:1706.09254.
- Shereen Oraby, Lena Reed, Shubhangi Tandon, TS Sharath, Stephanie Lukin, and Marilyn Walker. 2018. Controlling personality-based stylistic variation with neural natural language generators. *arXiv preprint arXiv:1805.08352*.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. In *NIPS-W*.
- Amanda Stent, Matthew Marge, and Mohit Singhai. 2005. Evaluating evaluation methods for generation in the presence of variation. In *International Conference on Intelligent Text Processing and Computational Linguistics*, pages 341–351. Springer.
- Amanda Stent, Rashmi Prasad, and Marilyn Walker. 2004. Trainable sentence planning for complex information presentation in spoken dialog systems. In *Proceedings of the 42nd annual meeting on association for computational linguistics*, page 79. Association for Computational Linguistics.
- Shang-Yu Su, Kai-Ling Lo, Yi-Ting Yeh, and Yun-Nung Chen. 2018. Natural language generation by hierarchical decoding with linguistic patterns. *arXiv preprint arXiv:1808.02747*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Kai Sheng Tai, Richard Socher, and Christopher D Manning. 2015. Improved semantic representations

- from tree-structured long short-term memory networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 1556–1566.
- Van-Khanh Tran and Le-Minh Nguyen. 2017. Natural language generation for spoken dialogue system using rnn encoder-decoder networks. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 442–451.
- Van-Khanh Tran and Le-Minh Nguyen. 2018. Adversarial domain adaptation for variational neural language generation in dialogue systems. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1205–1217.
- Van-Khanh Tran, Le-Minh Nguyen, and Satoshi Tojo. 2017. Neural-based natural language generation in dialogue using rnn encoder-decoder with semantic aggregation. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 231–240.
- Bo-Hsiang Tseng, Florian Kreyszig, Paweł Budzianowski, Iñigo Casanueva, Yen-Chen Wu, Stefan Ultes, and Milica Gasic. 2018. Variational cross-domain natural language generation for spoken dialogue systems. In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue*, pages 338–343.
- Marilyn A Walker, Owen C Rambow, and Monica Rogati. 2002. Training a sentence planner for spoken dialogue using boosting. *Computer Speech & Language*, 16(3-4):409–433.
- Marilyn A Walker, Amanda Stent, François Mairesse, and Rashmi Prasad. 2007. Individual and domain adaptation in sentence planning for dialogue. *Journal of Artificial Intelligence Research*, 30:413–456.
- Tsung-Hsien Wen, Milica Gašić, Dongho Kim, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015a. Stochastic language generation in dialogue using recurrent neural networks with convolutional sentence reranking. In *16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, page 275.
- Tsung-Hsien Wen, Milica Gašić, Nikola Mrkšić, Lina M. Rojas-Barahona, Pei-Hao Su, David Vandyke, and Steve Young. 2016. Multi-domain neural network language generation for spoken dialogue systems. In *Proceedings of the 2016 Conference on North American Chapter of the Association for Computational Linguistics (NAACL)*. Association for Computational Linguistics.
- Tsung-Hsien Wen, Milica Gašić, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015b. Semantically conditioned lstm-based natural language generation for spoken dialogue systems. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- Steve Young, Milica Gašić, Blaise Thomson, and Jason D. Williams. 2013. POMDP-based statistical spoken dialog systems: A review. *Proceedings of the IEEE*, 101(5):1160–1179.

Structured Fusion Networks for Dialog

Shikib Mehri*, Tejas Srinivasan*, and Maxine Eskenazi
Language Technologies Institute, Carnegie Mellon University
{amehri, tsriniva, max+}@cs.cmu.edu

Abstract

Neural dialog models have exhibited strong performance, however their end-to-end nature lacks a representation of the explicit structure of dialog. This results in a loss of generalizability, controllability and a data-hungry nature. Conversely, more traditional dialog systems do have strong models of explicit structure. This paper introduces several approaches for explicitly incorporating structure into neural models of dialog. Structured Fusion Networks first learn neural dialog modules corresponding to the structured components of traditional dialog systems and then incorporate these modules in a higher-level generative model. Structured Fusion Networks obtain strong results on the MultiWOZ dataset, both with and without reinforcement learning. Structured Fusion Networks are shown to have several valuable properties, including better domain generalizability, improved performance in reduced data scenarios and robustness to divergence during reinforcement learning.

1 Introduction

End-to-end neural dialog systems have shown strong performance (Vinyals and Le, 2015; Dinan et al., 2019). However such models suffer from a variety of shortcomings, including: a data-hungry nature (Zhao and Eskenazi, 2018), a tendency to produce generic responses (Li et al., 2016b), an inability to generalize (Mo et al., 2018; Zhao and Eskenazi, 2018), a lack of controllability (Hu et al., 2017), and divergent behavior when tuned with reinforcement learning (Lewis et al., 2017). Traditional dialog systems, which are generally free of these problems, consist of three distinct components: the natural language understanding (NLU), which produces a structured representation of an




Figure 1: A traditional dialog system consisting of a natural language understanding (NLU), dialog manager (DM) and natural language generation (NLG).

input (e.g., a belief state); the natural language generation (NLG), which produces output in natural language conditioned on an internal state (e.g. dialog acts); and the dialog manager (DM) (Bohus and Rudnicky, 2009), which describes a policy that combines an input representation (e.g., a belief state) and information from some database to determine the desired continuation of the dialog (e.g., dialog acts). A traditional dialog system, consisting of an NLU, DM and NLG, is pictured in Figure 1.

The structured components of traditional dialog systems facilitate effective generalizability, interpretability, and controllability. The structured output of each component allows for straightforward modification, understanding and tuning of the system. On the other hand, end-to-end neural models of dialog lack an explicit structure and are treated as a black box. To this end, we explore several methods of incorporating the structure of traditional dialog systems into neural dialog models.

First, several neural *dialog modules* are constructed to serve the role of the NLU, the DM and the NLG. Next, a number of methods are proposed for incorporating these dialog modules into end-to-end dialog systems, including Naïve Fusion, Multitask Fusion and Structured Fusion Networks (SFNs). This paper will show that SFNs

* Equal contribution.

obtain strong results on the MultiWOZ dataset (Budzianowski et al., 2018) both with and without the use of reinforcement learning. Due to the explicit structure of the model, SFNs are shown to exhibit several valuable properties including improved performance in reduced data scenarios, better domain generalizability and robustness to divergence during reinforcement learning (Lewis et al., 2017).

2 Related Work

2.1 Generation Methods

Vinyals and Le (2015) used a sequence-to-sequence network (Sutskever et al., 2014) for dialog by encoding the conversational context and subsequently generating the reply. They trained and evaluated their model on the OpenSubtitles dataset (Tiedemann, 2009), which contains conversations from movies, with a total of 62M training sentences.

Most research on generative models of dialog has built on the baseline introduced by Vinyals and Le (2015) by incorporating various forms of inductive bias (Mitchell, 1980) into their models, whether it be through the training procedure, the data or through the model architecture. Li et al. (2015) use Maximum Mutual Information (MMI) as the objective function, as a way of encouraging informative agent responses. Serban et al. (2016) proposes to better capture the semantics of dialog with the use of a hierarchical encoder decoder (HRED), comprised of an utterance encoder, a conversational context encoder, and a decoder. Li et al. (2016b) incorporate a number of heuristics into the reward function, to encourage useful conversational properties such as informativity, coherence and forward-looking. Li et al. (2016a) encodes a speaker’s persona as a distributed embedding and uses it to improve dialog generation. Liu and Lane (2016) simultaneously learn intent modelling, slot filling and language modelling. Zhao et al. (2017) enables task-oriented systems to make slot-value-independent decisions and improves out-of-domain recovery through the use of entity indexing and delexicalization. Wu et al. (2017) present Recurrent Entity Networks which use action templates and reasons about abstract entities in an end-to-end manner. Zhao and Eskenazi (2018) present the Action Matching algorithm, which maps utterances to a cross-domain embedding space to improve zero-shot generaliz-

ability. Mehri et al. (2019) explore several dialog specific pre-training objectives that improve performance on downstream dialog tasks, including generation. Chen et al. (2019) present a hierarchical self-attention network, conditioned on graph structured dialog acts and pre-trained with BERT (Devlin et al., 2018).

2.2 Generation Problems

Despite their relative success, end-to-end neural dialog systems have been shown to suffer from a number of shortcomings. (Li et al., 2016b) introduced the dull response problem, which describes how neural dialog systems tend to produce generic and uninformative responses (e.g., “I don’t know”). Zhao and Eskenazi (2018) describe generative dialog models as being data-hungry, and difficult to train in low-resource environments. Mo et al. (2018); Zhao and Eskenazi (2018) both demonstrate that dialog systems have difficulty generalizing to new domains. Hu et al. (2017) work on the problem of controllable text generation, which is difficult in sequence-to-sequence architectures, including generative models of dialog.

Wang et al. (2016) describe the problem of the *overwhelming implicit language model* in image captioning model decoders. They state that the decoder learns a language generation model along with a policy, however, during the process of captioning certain inputs, the decoder’s implicit language model overwhelms the policy and, as such, generates a specific output regardless of the input (e.g., if it generates ‘giraffe’, it may always output ‘a giraffe standing in a field’, regardless of the image). In dialog modelling, this problem is observed in the output of dialog models fine-tuned with reinforcement learning (Lewis et al., 2017; Zhao et al., 2019). Using reinforcement learning to fine-tune a decoder, will likely place a strong emphasis on improving the decoder’s policy and un-learn the implicit language model of the decoder. To this end, Zhao et al. (2019) proposes Latent Action Reinforcement Learning which does not update the decoder during reinforcement learning.

The methods proposed in this paper aim to mitigate these issues by explicitly modelling structure. Particularly interesting is that the structured models will reduce the effect of the *overwhelming implicit language model* by explicitly modelling the

NLG (i.e., a conditioned language model). This should lessen the divergent effect of reinforcement learning (Lewis et al., 2017; Zhao et al., 2019).

2.3 Fusion Methods

This paper aims to incorporate several pre-trained *dialog modules* into a neural dialog model. A closely related branch of research is the work done on fusion methods, which attempts to integrate pre-trained language models into sequence-to-sequence networks. Integrating language models in this manner is a form of incorporating structure into neural architectures. The simplest such method, commonly referred to as **Shallow Fusion**, is to add a language modeling term, $p_{LM}(y)$, to the cost function during inference (Chorowski and Jaitly, 2016).

To improve on this, Gulcehre et al. (2015) proposed **Deep Fusion**, which combines the states of a pre-trained machine translation models decoder and a pre-trained language model by concatenating them using a gating mechanism with trained parameters. The gating mechanism allows us to decide how important the language model and decoder states are at each time step in the inference process. However, one major drawback of Deep Fusion is that the sequence-to-sequence model is trained independently from the language model, and has to learn an implicit language model from the training data.

Cold Fusion (Sriram et al., 2017) deals with this problem by training the sequence-to-sequence model along with the gating mechanism, thus making the model aware of the pre-trained language model throughout the training process. The decoder does not need to learn a language model from scratch, and can thus learn more task-specific language characteristics which are not captured by the pre-trained language model (which has been trained on a much larger, domain-agnostic corpus).

3 Methods

This section describes the methods employed in the task of dialog response generation. In addition to the baseline model proposed by Budzianowski et al. (2018), several methods of incorporating structure into end-to-end neural dialog models are explored.

3.1 Sequence-to-Sequence

The baseline model for dialog generation, depicted in Figure 2, consists of a standard encoder-decoder framework (Sutskever et al., 2014), augmented with a belief tracker (obtained from the annotations of the dialog state) and a database vector. The dialog system is tasked with producing the appropriate system response, given a dialog context, an oracle belief state representation and a vector corresponding to the database output.

The dialog context is encoded using an LSTM (Hochreiter and Schmidhuber, 1997) sequence-to-sequence network (Sutskever et al., 2014). Experiments are conducted with and without an attention mechanism (Bahdanau et al., 2015). Given the final encoder hidden state, h_t^e , the belief state vector, v_{bs} , and the database vector, v_{db} , Equation 1 describes how the initial decoder hidden state is obtained.

$$h_0^d = \tanh(W_e h_t^e + W_{bs} v_{bs} + W_{db} v_{db} + b) \quad (1)$$




Figure 2: A diagram of the baseline sequence-to-sequence architecture. The attention mechanism is not visualized, however experiments are conducted both with and without attention.

3.2 Neural Dialog Modules

As seen in Figure 1, a traditional dialog system consists of the NLU, the DM and the NLG. The NLU maps a natural language input to a belief state representation (BS). The DM uses the belief state and some database output, to produce dialog acts (DA) for the system response. The NLG uses the dialog acts to produce a natural language response.

A neural *dialog module* is constructed for each of these three components. A visualization of these architectures is shown in Figure 3. The NLU architecture uses an LSTM encoder to map the

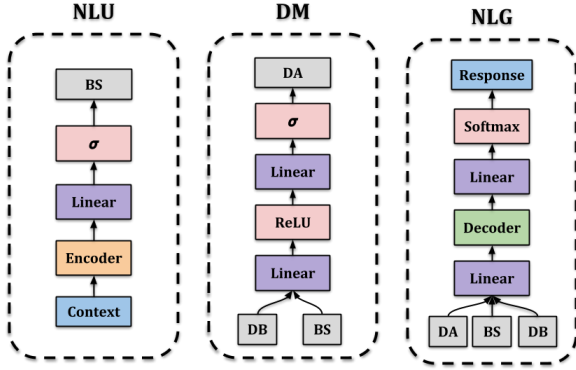


Figure 3: A visualization of the neural architectures for each of the three modules of traditional dialog systems.

natural language input to a latent representation, h_t , which is then passed through a linear layer and a sigmoid function to obtain a multi-label prediction of the belief state. The DM architecture projects the belief state and database vector into a latent space, through the use of a linear layer with a ReLU activation, which is then passed through another linear layer and a sigmoid function to predict the dialog act vector. The neural architecture corresponding to the NLG is a conditioned language model with its initial hidden state given by a linear encoding of the dialog acts, belief state and database vectors.

The following equations define the structure of the modules, where the gt subscript on an intermediate variable denotes the use of the ground-truth value:

$$bs = \text{NLU}(\text{context}) \quad (2)$$

$$da = \text{DM}(bs_{gt}, db) \quad (3)$$

$$\text{response} = \text{NLG}(bs_{gt}, db, da_{gt}) \quad (4)$$

3.3 Naïve Fusion

Naïve Fusion (NF) is a straightforward mechanism for using the neural dialog modules for end-to-end dialog response generation.

3.3.1 Zero-Shot Naïve Fusion

During training, each dialog module is trained independently, meaning that it is given the ground truth input and supervision signal. However, during inference, the intermediate values (e.g., the dialog act vector) do not necessarily exist and the outputs of other neural modules must be used instead. For example, the DM module is trained given the ground-truth belief state as input, however during inference it must rely on the belief state predicted by the NLU module. This results

in a propagation of errors, as the DM and NLG may receive imperfect input.

Zero-Shot Naïve Fusion combines the pre-trained neural modules at inference time. The construction of the response conditioned on the context, is described as follows:

$$bs = \text{NLU}(\text{context}) \quad (5)$$

$$\text{response} = \text{NLG}(bs, db, \text{DM}(bs, db)) \quad (6)$$

3.3.2 Naïve Fusion with Fine-Tuning

Since the forward propagation described in Equations 5 and 6 is continuous and there is no sampling procedure until the response is generated, Naïve Fusion can be fine-tuned for the end-to-end task of dialog generation. The pre-trained neural modules are combined as described above, and fine-tuned on the task of dialog generation using the same data and learning objective as the baseline.

3.4 Multitask Fusion

Structure can be incorporated into neural architectures through the use of multi-tasking. Multitask Fusion (MF) is a method where the end-to-end generation task is learned simultaneously with the aforementioned dialog modules. The multi-tasking setup is seen in Figure 4.




Figure 4: A depiction of Multitask Fusion, where the individual neural modules are learned simultaneously with the end-to-end task of dialog generation. The dashed boxes contain the individual components, while the red arrows depict forward propagation for the end-to-end task. The red arrows are the process used during response generation.

By sharing the weights of the end-to-end architecture and each respective module, the learned representations should become stronger and more structured in nature. For example, the encoder is

shared between the NLU module and the end-to-end task. As such, it will learn to both represent the information necessary for predicting the belief state vector and any additional information useful for generating the next utterance.

3.5 Structured Fusion Networks

The Structured Fusion Networks (SFNs) we propose, depicted in Figure 5, use the independently pre-trained neural dialog modules for the task of end-to-end dialog generation. Rather than fine-tuning or multi-tasking the independent modules, SFNs aim to learn a higher-level model on top of the neural modules to perform the task of end-to-end response generation.

The output of the NLU is concatenated at each time-step of the encoder input. The output of the DM is similarly concatenated to the input of the linear layer between the encoder and the decoder of the higher-level model. The output of the NLG, in the form of logits at a decoding time-step, is combined with the hidden state of the decoder via cold-fusion (Sriram et al., 2017). Given the NLG output as l_t^{NLG} and the higher-level decoder hidden state as s_t , the cold-fusion method is described as follows:

$$h_t^{NLG} = DNN(l_t^{NLG}) \quad (7)$$

$$g_t = \sigma(W[s_t; h_t^{NLG}] + b) \quad (8)$$

$$s_t^{CF} = [s_t; g_t \circ h_t^{NLG}] \quad (9)$$

$$y_t = softmax(DNN(s_t^{CF})) \quad (10)$$

By pre-training the modules and using their structured outputs, the higher-level model does not have to *re-learn* and *re-model* the dialog structure (i.e., representing the belief state and dialog acts). Instead, it can focus on the more abstract modelling that is necessary for the task, including recognizing and encoding complex natural language input, modelling a policy, and effectively converting a latent representation into a natural language output according to the policy.

The SFN architecture may seem complicated due to the redundancy of the inputs. For example, the context is passed to the model in two places and the database vector in three places. This redundancy is necessary for two reasons. First, each of the neural modules must function independently and thus needs sufficient inputs. Second, the higher-level model should be able to function

well independently. If any of the neural modules was to be removed, the SFN should be able to perform reasonably. This means that the higher-level module should not rely on any of the neural modules to capture information about the input and therefore allow the neural modules to focus only on representing the structure. For example, if the context was not passed into the higher-level encoder and instead only to the NLU module, then the NLU may no longer be able to sufficiently model the belief state and may instead have to more explicitly model the context (e.g., as a bag-of-words representation).

Several variations of training SFNs are considered during experimentation, enumerated as follows. (1) The pre-trained neural modules are kept frozen, as a way of ensuring that the structure is not deteriorated. (2) The pre-trained neural modules are fine-tuned for the end-to-end task of response generation. This ensures that the model is able to abandon or modify certain elements of the structure if it helps with the end-to-end task. (3) The pre-trained modules are multi-tasked with the end-to-end task of response generation. This ensures that the structure is maintained and potentially strengthened while also allowing the modules to update and improve for the end-to-end task.

4 Experiments

4.1 Dataset

The dialog systems are evaluated on the MultiWOZ dataset (Budzianowski et al., 2018), which consists of ten thousand human-human conversations covering several domains. The MultiWOZ dataset contains conversations between a tourist and a clerk at an information center which fall into one of seven domains - attraction, hospital, police, hotel, restaurant, taxi, train. Individual conversations span one to five of the domains. Dialogs were collected using the Wizard-of-Oz framework, where one participant plays the role of an automated system.

Each dialog consists of a goal and multiple user and system utterances. Each turn is annotated with two binary vectors: a belief state vector and a dialog act vector. A single turn may have multiple positive values in both the belief state and dialog act vectors. The belief state and dialog act vectors are of dimensions 94 and 593, respectively.

Several metrics are used to evaluate the models. BLEU (Papineni et al., 2002) is used to com-




Figure 5: The Structured Fusion Network. The grey dashed boxes correspond to the pre-trained neural dialog modules. A higher-level is learned on top of the pre-trained modules, as a mechanism of enforcing structure in the end-to-end model.

pute the word overlap between the generated output and the reference response. Two task-specific metrics, defined by Budzianowski et al. (2018), Inform rate and Success rate, are also used. Inform rate measures how often the system has provided the appropriate entities to the user. Success rate measures how often the system answers all the requested attributes. Similarly to Budzianowski et al. (2018), the best model is selected during validation using the combined score which is defined as $BLEU + 0.5 \times (Inform + Success)$. This combined score is also reported as an evaluation metric.

4.2 Experimental Settings

The hyperparameters match those used by Budzianowski et al. (2018): embedding dimension of 50, hidden dimension of 150, and a single-layer LSTM. All models are trained for 20 epochs using the Adam optimizer (Kingma and Ba, 2014), with a learning rate of 0.005 and batch size of 64. The norm of the gradients are clipped to 5 (Pascanu et al., 2012). Greedy decoding is used during inference.

All previous work uses the ground-truth belief state vector during training and evaluation. Therefore the experiments with the SFNs have the NLU module replaced by an "oracle NLU" which always outputs the ground-truth belief state. Table 4 in the Appendix shows experimental results which demonstrate that using only the ground-truth be-

lief state results in the best performance.

4.3 Reinforcement Learning

A motivation of explicit structure is the hypothesis that it will reduce the effects of the implicit language model, and therefore mitigate degenerate output after reinforcement learning. This hypothesis is evaluated by fine-tuning the SFNs with reinforcement learning. The setup for this experiment is similar to that of Zhao et al. (2019): (1) the model produces a response conditioned on a ground-truth dialog context, (2) the success rate is evaluated for the generated response, (3) using the success rate as the reward, the policy gradient is calculated at each word, and (4) the parameters of the model are updated. A learning rate of $1e-5$ is used with the Adam optimizer (Kingma and Ba, 2015).

Reinforcement learning is used to fine-tune the best performing model trained in a supervised learning setting. During this fine-tuning, the neural dialog modules (i.e., the NLU, DM and NLG) are frozen. Only the high-level model is updated during reinforcement learning. Freezing maintains the structure, while still updating the higher level components. Since the structure is maintained, it is unnecessary to alternate between supervised and reinforcement learning.

Model	BLEU	Inform	Success	Combined Score
Supervised Learning				
Seq2Seq (Budzianowski et al., 2018)	18.80	71.29%	60.29%	84.59
Seq2Seq w/ Attn (Budzianowski et al., 2018)	18.90	71.33%	60.96%	85.05
Seq2Seq (Ours)	20.78	61.40%	54.50%	78.73
Seq2Seq w/ Attn (ours)	20.36	66.50%	59.50%	83.36
3-layer HDSA (Chen et al., 2019)	23.60	82.90%	68.90%	99.50
Naïve Fusion (Zero-Shot)	7.55	70.30%	36.10%	60.75
Naïve Fusion (Fine-tuned Modules)	16.39	66.50%	59.50%	83.36
Multitasking	17.51	71.50%	57.30%	81.91
Structured Fusion (Frozen Modules)	17.53	65.80%	51.30%	76.08
Structured Fusion (Fine-tuned Modules)	18.51	77.30%	64.30%	89.31
Structured Fusion (Multitasked Modules)	16.70	80.40%	63.60%	88.71
Reinforcement Learning				
Seq2Seq + RL (Zhao et al., 2019)	1.40	80.50%	79.07%	81.19
LiteAttnCat + RL (Zhao et al., 2019)	12.80	82.78%	79.20%	93.79
Structured Fusion (Frozen Modules) + RL	16.34	82.70%	72.10%	93.74

Table 1: Experimental results for the various models. This table compares two classes of methods: those trained with supervised learning and those trained with reinforcement learning. All bold-face results are statistically significant ($p < 0.01$).

4.4 Results

Experimental results in Table 1 show that our Structured Fusion Networks (SFNs) obtain strong results when compared to both methods trained with and without the use of reinforcement learning. Compared to previous methods trained only with supervised learning, SFNs obtain a **+4.26** point improvement over seq2seq baselines in the combined score with strong improvement in both Success and Inform rates. SFNs are outperformed by the recently published HDSA (Chen et al., 2019) models which relies on BERT (Devlin et al., 2018) and conditioning on graph structured dialog acts. When using reinforcement learning, SFNs match the performance of LiteAttnCat (Zhao et al., 2019) on the combined score. Though the Inform rate is equivalent and the Success rate is lower (albeit still better than all supervised methods), the BLEU score of SFNs is much better with an improvement of **+3.54** BLEU over LiteAttnCat.

In the reinforcement learning setting, the improved BLEU can be attributed to the explicit structure of the model. This structure enables the model to optimize for the reward (Success rate) without resulting in degenerate output (Lewis et al., 2017).

SFNs obtain the highest combined score when the modules are fine-tuned. This is likely because, while the structured modules serve as a strong ini-

tialization for the task of dialog generation, forcing the model to maintain the exact structure (i.e., frozen modules) limits its ability to learn. In fact, the end-to-end model may choose to ignore some elements of intermediate structure (e.g., a particular dialog act) which prove useless for the task of response generation.

Despite strong overall performance, SFNs do show a **-2.27** BLEU drop when compared to the strongest seq2seq baseline and a **-5.09** BLEU drop compared to HDSA. Though it is difficult to ascertain the root cause of this drop, one potential reason could be that the dataset contains many social niceties and generic statements (e.g., "happy anniversary") which are difficult for a structured model to effectively generate (since it is not an element of the structure) while a free-form sequence-to-sequence network would not have this issue.

To a lesser degree, multi-tasking (i.e., multi-tasked modules) would also prevent the model from being able to ignore some elements of the structure. However, the SFN with multitasked modules performs best on the Inform metric with a **+9.07%** improvement over the seq2seq baselines and a **+3.10%** over other SFN-based methods. This may be because the Inform metric measures how many of the requested attributes were answered, which benefits from a structured representation of the input.

Zero-Shot Naïve Fusion performs very poorly, suggesting that the individual components have difficulty producing good results when given imperfect input. Though the NLG module performs extremely well when given the oracle dialog acts (28.97 BLEU; 106.02 combined), its performance deteriorates significantly when given the predicted dialog acts. This observation is also applicable to Structured Fusion with frozen modules.

HDSA (Chen et al., 2019) outperforms SFN possibly due to the use of a more sophisticated Transformer model (Vaswani et al., 2017) and BERT pre-training (Devlin et al., 2018). A unique advantage of SFNs is that the architecture of the *neural dialog modules* is flexible. The performance of HDSA could potentially be integrated with SFNs by using the HDSA model as the NLG module of an SFN. This is left for future work, as the HDSA model was released while this paper was already in review.

These strong performance gains reaffirm the hypothesis that adding explicit structure to neural dialog systems results in improved modelling ability particularly with respect to dialog policy as we see in the increase in Inform and in Success. The results with reinforcement learning suggest that the explicit structure allows *controlled fine-tuning* of the models, which prevents divergent behavior and degenerate output.

4.5 Human Evaluation

To supplement the results in Table 1, human evaluation was used to compare seq2seq, SFN, SFN fine-tuned with reinforcement learning, and the ground-truth human response. Workers on Amazon Mechanical Turk (AMT) were asked to read the context, and score the *appropriateness* of each response on a Likert scale (1-5). One hundred context-response pairs were labeled by three workers each. The results shown in Table 2 demonstrate that SFNs with RL outperform the other methods in terms of human judgment. These results indicate that in addition to improving on automated metrics, SFNs result in user-favored responses.

5 Analysis

5.1 Limited Data


Structured Fusion Networks (SFNs) should outperform sequence-to-sequence (seq2seq) networks in reduced data scenarios due to the explicit

Model	Avg Rating	≥ 4	≥ 5
Seq2Seq	3.00	40.21%	9.61%
SFN	3.02	44.84%	11.03%
SFN + RL	3.12	44.84%	16.01%
Human	3.76	59.79%	34.88%


Table 2: Results of human evaluation experiments. The ≥ 4 and ≥ 5 columns indicate the percentage of system outputs which obtained a greater than 4 and 5 rating, respectively.

structure. While a baseline method would require large amounts of data to learn to infer structure, SFNs do this explicitly.

The performance of seq2seq and SFNs are determined, when training on 1%, 5%, 10% and 25% of the training data (total size of $\sim 55,000$ utterances). The supervised-learning variant of SFNs with fine-tuned modules is used. The pre-training of the modules and fine-tuning of the full model is done on the same data split. The full data is used during validation and testing.



(a)



(b)

Figure 6: Variation of Inform (a) and Success (b) rate at different amounts of training data.

The results in Figure 6 show the Inform and Success rates for different amounts of training data. SFNs significantly outperform the seq2seq model in low-data scenarios. Notably, improve-

ment is considerably higher in the most extreme low-data scenario, when only 1% of the training data (~ 550 dialogs) is used. As the amount of training data increases, the gap between the two models stabilizes. The effectiveness at extreme low-data scenarios reaffirms the hypothesis that explicit structure makes SFNs less data-hungry than sequence-to-sequence networks.

5.2 Domain Generalizability

The explicit structure of SFNs should facilitate effective domain generalizability. A domain transfer experiment was constructed to evaluate the comparative ability of seq2seq and SFNs. The models were both trained on a reduced dataset that largely consists of out-of-domain examples and evaluated on in-domain examples. Specifically, 2000 out-of-domain training examples and only 50 in-domain training examples were used. The restaurant domain of MultiWOZ was selected as in-domain.

Model	BLEU	Inform	Success
Seq2Seq	10.22	35.65%	1.30%
SFN	7.44	47.17%	2.17%

Table 3: Results of the domain transfer experiment comparing sequence-to-sequence and Structured Fusion Networks. All bold-face results are statistically significant ($p < 0.01$).

The results, seen on Table 3, show that SFNs perform significantly better on both the Inform (+**11.52%**) and Success rate. Although SFNs have a slightly higher Success rate, both models perform poorly. This is expected since the models would be unable to answer all the requested attributes when they have seen little domain data – their language model would not be tuned to the in-domain task. The **-2.78** BLEU reduction roughly matches the BLEU difference observed on the main task, therefore it is not an issue specific to domain transfer.

6 Conclusions and Future Work

This paper presents several methods of incorporating explicit structure into end-to-end neural models of dialog. We created Structured Fusion Networks, comprised of pre-trained *dialog modules* and a higher-level end-to-end network, which obtain strong results on the MultiWOZ dataset both with and without the use of reinforcement learning. SFNs are further shown to be robust to divergence during reinforcement learning, effective

in low data scenarios and better than sequence-to-sequence on the task of domain transfer.

For future research, the explicit structure of SFNs has been shown to have multi-faceted benefits; another potential benefit may be interpretability. It would be interesting to investigate the use of SFNs as more interpretable models of dialog. While domain generalizability has been demonstrated, it would be useful to further explore the nature of generalizability (e.g., task transfer, language style transfer). Another potential avenue of research is whether the explicit structure of SFNs could potentially allow swapping the dialog modules without any fine-tuning. Structured Fusion Networks highlight the effectiveness of using explicit structure in end-to-end neural networks, suggesting that exploring alternate means of incorporating structure would be a promising direction for future work.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of ICLR*.
- Dan Bohus and Alexander I Rudnicky. 2009. The ravenclaw dialog management framework: Architecture and systems. *Computer Speech & Language*, 23(3):332–361.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. Multiwoz-a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. *arXiv preprint arXiv:1810.00278*.
- Wenhu Chen, Jianshu Chen, Pengda Qin, Xifeng Yan, and William Yang Wang. 2019. Semantically conditioned dialog response generation via hierarchical disentangled self-attention. *arXiv preprint arXiv:1905.12866*.
- Jan Chorowski and Navdeep Jaitly. 2016. Towards better decoding and language model integration in sequence to sequence models. *arXiv preprint arXiv:1612.02695*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Emily Dinan, Varvara Logacheva, Valentin Malykh, Alexander Miller, Kurt Shuster, Jack Urbanek, Douwe Kiela, Arthur Szlam, Iulian Serban, Ryan Lowe, et al. 2019. The second conversational

- intelligence challenge (convai2). *arXiv preprint arXiv:1902.00098*.
- Çaglar Gulcehre, Orhan Firat, Kelvin Xu, Kyunghyun Cho, Loïc Barrault, Huihui Lin, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2015. On using monolingual corpora in neural machine translation. *arXiv preprint arXiv:1503.03535*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P Xing. 2017. Toward controlled generation of text. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1587–1596. JMLR. org.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR*.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of ICLR*.
- Mike Lewis, Denis Yarats, Yann N Dauphin, Devi Parikh, and Dhruv Batra. 2017. Deal or no deal? end-to-end learning for negotiation dialogues. *arXiv preprint arXiv:1706.05125*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2015. A diversity-promoting objective function for neural conversation models. *arXiv preprint arXiv:1510.03055*.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios P Spithourakis, Jianfeng Gao, and Bill Dolan. 2016a. A persona-based neural conversation model. *arXiv preprint arXiv:1603.06155*.
- Jiwei Li, Will Monroe, Alan Ritter, Michel Galley, Jianfeng Gao, and Dan Jurafsky. 2016b. Deep reinforcement learning for dialogue generation. *arXiv preprint arXiv:1606.01541*.
- Bing Liu and Ian Lane. 2016. Joint online spoken language understanding and language modeling with recurrent neural networks. *arXiv preprint arXiv:1609.01462*.
- Shikib Mehri, Evgeniia Razumovskaia, Tiancheng Zhao, and Maxine Eskenazi. 2019. Pretraining methods for dialog context representation learning. *arXiv preprint arXiv:1906.00414*.
- Tom M Mitchell. 1980. *The need for biases in learning generalizations*. Department of Computer Science, Laboratory for Computer Science Research .
- Kaixiang Mo, Yu Zhang, Qiang Yang, and Pascale Fung. 2018. Cross-domain dialogue policy transfer via simultaneous speech-act and slot alignment. *arXiv preprint arXiv:1804.07691*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318. Association for Computational Linguistics.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2012. Understanding the exploding gradient problem. *CoRR*.
- Iulian Vlad Serban, Alessandro Sordani, Yoshua Bengio, Aaron C Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *AAAI*, volume 16, pages 3776–3784.
- Anuroop Sriram, Heewoo Jun, Sanjeev Satheesh, and Adam Coates. 2017. Cold fusion: Training seq2seq models together with language models. *arXiv preprint arXiv:1708.06426*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Jörg Tiedemann. 2009. News from opus-a collection of multilingual parallel corpora with tools and interfaces. In *Recent advances in natural language processing*, volume 5, pages 237–248.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 6000–6010.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Zhuhao Wang, Fei Wu, Weiming Lu, Jun Xiao, Xi Li, Zitong Zhang, and Yueting Zhuang. 2016. Diverse image captioning via grouptalk. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, pages 2957–2964. AAAI Press.
- Chien-Sheng Wu, Andrea Madotto, Genta Winata, and Pascale Fung. 2017. End-to-end recurrent entity network for entity-value independent goal-oriented dialog learning. In *Dialog System Technology Challenges Workshop, DSTC6*.
- Tiancheng Zhao and Maxine Eskenazi. 2018. Zero-shot dialog generation with cross-domain latent actions. *arXiv preprint arXiv:1805.04803*.
- Tiancheng Zhao, Allen Lu, Kyusong Lee, and Maxine Eskenazi. 2017. Generative encoder-decoder models for task-oriented spoken dialog systems with chatting capability. *arXiv preprint arXiv:1706.08476*.

Tiancheng Zhao, Kaige Xie, and Maxine Eskenazi.
2019. Rethinking action spaces for reinforcement learning in end-to-end dialog agents with latent variable models. *arXiv preprint arXiv:1902.08858*.

A Belief State Ablation Study

All previous research working on dialog generation for the MultiWOZ dataset uses the ground-truth belief state vector during training and evaluation. Therefore for fair comparability, the SFN experiments in our paper had the NLU module replaced by an "oracle NLU" which always outputs the ground-truth belief state.

An ablation experiment was performed to ascertain whether providing *only* the ground-truth belief state was the optimal solution. Several methods of combining the ground-truth belief state with the pre-trained NLU module were explored. These methods are enumerated as follows:

- (1) **Ground-Truth Only:** The setting used in the primary experiments, shown in Table 1 of the main paper. Only the ground-truth belief state vector is used.
- (2) **Predicted Only:** Only the belief state predicted by the pre-trained NLU module is used.
- (3) **Sum:** The predicted and ground-truth belief states are summed, before being used by all upper layers.
- (4) **Linear:** The predicted and ground-truth belief states are concatenated and passed through a linear layer.

These experiments are performed using the best model, Structured Fusion Networks with fine-tuned modules. The results are shown in Table 4.

Model	BLEU	Inform	Success	Comb.
GT	18.51	77.30%	64.30%	89.31
Pred	16.88	73.80%	58.60%	83.04
Sum	15.93	72.90%	60.80%	82.78
Linear	15.42	66.80%	54.80%	76.22

Table 4: Results of the domain transfer experiment comparing sequence-to-sequence and Structured Fusion Networks. All bold-face results are statistically significant ($p < 0.01$).

It is observed that adding the pre-trained NLU does not provide any additional performance benefit, when the ground-truth belief state is already provided. As such, combinations of the ground-truth and predicted belief state actually perform worse than either of the methods independently because of (1) additional parameters to be learned,

especially in the case of the *Linear* method, and (2) a conflicting trade-off between fine-tuning a learned NLU module and using the ground-truth belief state.

B Qualitative Examples

Table 5 shows several examples of dialogs from the test set of MultiWOZ, along with the produced response from three different models: sequence-to-sequence networks, Structured Fusion Networks, and Structured Fusion Networks fine-tuned with reinforcement learning. These examples serve to provide insight into the respective strengths and weaknesses of the different models. A few noteworthy observations from the four examples are enumerated below:

- (1) SFN fine-tuned with RL **consistently provides more attribute information.** It provides at least one attribute in every example response, for a total of 14 total attributes across the four examples. This, along with the high Success score of this model, is a consequence of the reinforcement learning fine-tuning which directly optimizes the Success score and rewards the model for providing the correct attributes.
- (2) Seq2Seq **produces more generic responses.** In the second and fourth examples, the outputs produced by Seq2Seq are generic and are unrelated to the context. This suggests that the Seq2Seq model has an overwhelming implicit language model, which produces generic and dull responses. On the other hand, the explicit structure in SFN mitigates the effect of the implicit language model by relying on an *explicit* language model (in the form of the NLG).
- (3) Seq2Seq **requests attributes which the user has already provided.** In the first and third example, the Seq2Seq output is requesting parameters that the user explicitly provided. This highlights the lack of structured representations in the Seq2Seq model. SFN which explicitly models the structure of the input effectively understands and captures the information provided in the context.

Example 1

Dialog Context:

USER: hello . i am looking for a [value_pricerange] place to eat in the [value_area] .
can you help me ?

Sequence-to-Sequence Response:

i have [value_count] options for you . do you have a preference on price range ?

Structured Fusion Network Response:

[restaurant_name] is a [value_pricerange] -ly priced [value_food] restaurant in the [value_area] .
would you like to make a reservation ?

Structured Fusion Network w/ Reinforcement Learning Response:

[restaurant_name] is a [value_pricerange] [value_food] restaurant in the [value_area] . it is
located at [restaurant_address] . the postcode is [restaurant_postcode] . the phone number
is [restaurant_phone]

Example 2

Dialog Context:

...

USER: i want to book this hotel for the same group of people and [value_count] nights starting
from [value_day] . i'll need a reference number please .

SYSTEM: your reservation is confirmed and the reference number is ##

USER: i am sorry ? were you able to book the express by holiday inn for [value_count] people ?

Sequence-to-Sequence Response:

you are welcome . have a great day !

Structured Fusion Network Response:

you are all set . your reference number is [hotel_reference] . can i help you with anything else ?

Structured Fusion Network w/ Reinforcement Learning Response:

your booking was successful . the reference number is [hotel_reference] .

Example 3

Dialog Context:

USER: can you see if there are any entertainment attractions in the [value_area] of town ?

Sequence-to-Sequence Response:

there are [value_count] museums in the [value_area] . do you have a preference for area ?

Structured Fusion Network Response:

there are [value_count] museums in the [value_area] . [attraction_name] is a museum and
the [attraction_name] is a [value_count] .

Structured Fusion Network w/ Reinforcement Learning Response:

[attraction_name] is a museum in the [value_area] and the [attraction_name] is located
at [attraction_address] . the postcode is [attraction_postcode] .

Example 4

Dialog Context:

USER: i need a taxi . i'll be departing from [restaurant_name] .

SYSTEM: i can help you with that . when do you need to leave ?

USER: i would like to leave after [value_time] please .

SYSTEM: where will you be going ?

USER: i'll be going to [restaurant_name] .

Sequence-to-Sequence Response:

you are welcome . have a good time !

Structured Fusion Network Response:

you are welcome . have a great day !

Structured Fusion Network w/ Reinforcement Learning Response:

i have booked for you a [taxi_type] , the contact number is [taxi_phone] . is there anything else i can
help you with ?

Table 5: Four examples of dialog contexts from the dataset, and the responses generated by three different models: Seq2Seq, SFN, and SFN with RL.

Flexibly-Structured Model for Task-Oriented Dialogues

Lei Shu¹*, Piero Molino², Mahdi Namazifar², Hu Xu¹, Bing Liu¹, Huaixiu Zheng², Gokhan Tur²

¹Department of Computer Science, University of Illinois at Chicago

²Uber AI,

¹{lshu3, hxu48, liub}@uic.edu,

²{piero, mahdin, huaixiu.zheng, gokhan}@uber.com

Abstract

This paper proposes a novel end-to-end architecture for task-oriented dialogue systems. It is based on a simple and practical yet very effective sequence-to-sequence approach, where language understanding and state tracking tasks are modeled jointly with a structured copy-augmented sequential decoder and a multi-label decoder for each slot. The policy engine and language generation tasks are modeled jointly following that. The copy-augmented sequential decoder deals with new or unknown values in the conversation, while the multi-label decoder combined with the sequential decoder ensures the explicit assignment of values to slots. On the generation part, slot binary classifiers are used to improve performance. This architecture is scalable to real-world scenarios and is shown through an empirical evaluation to achieve state-of-the-art performance on both the Cambridge Restaurant dataset and the Stanford in-car assistant dataset¹.

1 Introduction

A traditional task-oriented dialogue system is often composed of a few modules, such as natural language understanding, dialogue state tracking, knowledge base (KB) query, dialogue policy engine and response generation. Language understanding aims to convert the input to some predefined semantic frame. State tracking is a critical component that models explicitly the input semantic frame and the dialogue history for producing KB queries. The semantic frame and the corresponding belief state are defined in terms of informable slots values and requestable slots. Informable slot values capture information provided by the user

so far, e.g., $\{price=cheap, food=italian\}$ indicating the user wants a cheap Italian restaurant at this stage. Requestable slots capture the information requested by the user, e.g., $\{address, phone\}$ means the user wants to know the address and phone number of a restaurant. Dialogue policy model decides on the system action which is then realized by a language generation component.

To mitigate the problems with such a classic modularized dialogue system, such as the error propagation between modules, the cascade effect that the updates of the modules have and the expensiveness of annotation, end-to-end training of dialogue systems was recently proposed (Liu and Lane, 2018; Williams et al., 2017; Lowe et al., 2017; Li et al., 2018; Liu et al., 2018; Budzianowski et al., 2018; Bordes et al., 2017; Wen et al., 2017b; Serban et al., 2016, among others). These systems train one whole model to read the current user’s utterance, the past state (that may contain all previous interactions) and generate the current state and response.

There are two main approaches for modeling the belief state in end-to-end task-oriented dialogue systems in the literature: the *fully structured* approach based on classification (Wen et al., 2017b,a), and the *free-form* approach based on text generation (Lei et al., 2018). The fully structured approaches (Ramadan et al., 2018; Ren et al., 2018) use the full structure of the KB, both its schema and the values available in it, and assumes that the sets of informable slot values and requestable slots are fixed. In real-world scenarios, this assumption is too restrictive as the content of the KB may change and users’ utterances may contain information outside the pre-defined sets. An ideal end-to-end architecture for state tracking should be able to identify the values of the informable slots and the requestable slots, easily adapt to new domains, to the changes in the content of the KB, and to the

Work mostly performed as an intern at Uber AI Labs

¹The code is available at <https://github.com/uber-research/FSDM>

occurrence of words in users’ utterances that are not present in the KB at training time, while at the same time providing the right amount of inductive bias to allow generalization.

Recently, a free-form approach called TSCP (Two Stage Copy Net) (Lei et al., 2018) was proposed. This approach does not integrate any information about the KB in the model architecture. It has the advantage of being readily adaptable to new domains and changes in the content of the KB as well as solving the out-of-vocabulary word problem by generating or copying the relevant piece of text from the user’s utterances in its response generation. However, TSCP can produce invalid states (see Section 4). Furthermore, by putting all slots together into a sequence, it introduces an unwanted (artificial) order between different slots since they are encoded and decoded sequentially. It could be even worse if two slots have overlapping values, like departure and arrival airport in a travel booking system. As such, the unnecessary order of the slots makes getting rid of the invalid states a great challenge for the sequential decoder. As a summary, both approaches to state tracking have their weaknesses when applied to real-world applications.

This paper proposes the Flexibly-Structured Dialogue Model (FSDM) as a new end-to-end task-oriented dialogue system. The state tracking component of FSDM has the advantages of both fully structured and free-form approaches while addressing their shortcomings. On one hand, it is still structured, as it incorporates information about slots in KB schema; on the other hand, it is flexible, as it does not use information about the values contained in the KB records. This makes it easily adaptable to new values. These desirable properties are achieved by a separate decoder for each informable slot and a multi-label classifier for the requestable slots. Those components explicitly assign values to slots like the fully structured approach, while also preserving the capability of dealing with out-of-vocabulary words like the free-form approach. By using these two types of decoders, FSDM produces only valid belief states, overcoming the limitations of the free-form approach. Further, FSDM has a new module called response slot binary classifier that adds extra supervision to generate the slots that will be present in the response more precisely before generating the final textual agent response (see Section 3 for details).

The main contributions of this work are

1. FSDM, a task-oriented dialogue system with a new belief state tracking architecture that overcomes the limits of existing approaches and scales to real-world settings;
2. a new module, namely the response slot binary classifier, that helps to improve the performance of agent response generation;
3. FSDM achieves state-of-the-art results on both the Cambridge Restaurant dataset (Wen et al., 2017b) and the Stanford in-car assistant dataset (Eric et al., 2017) without the need for fine-tuning through reinforcement learning

2 Related Work

Our work is related to end-to-end task-oriented dialogue systems in general (Liu and Lane, 2018; Williams et al., 2017; Lowe et al., 2017; Li et al., 2018; Liu et al., 2018; Budzianowski et al., 2018; Bordes et al., 2017; Hori et al., 2016; Wen et al., 2017b; Serban et al., 2016, among others) and those that extend the Seq2Seq (Sutskever et al., 2014) architecture in particular (Eric et al., 2017; Fung et al., 2018; Wen et al., 2018). Belief tracking, which is necessary to form KB queries, is not explicitly performed in the latter works. To compensate, Eric et al. (2017); Xu and Hu (2018a); Wen et al. (2018) adopt a copy mechanism that allows copying information retrieved from the KB to the generated response. Fung et al. (2018) adopt Memory Networks (Sukhbaatar et al., 2015) to memorize the retrieved KB entities and words appearing in the dialogue history. These models scale linearly with the size of the KB and need to be retrained at each update of the KB. Both issues make these approaches less practical in real-world applications.

Our work is also akin to modularly connected end-to-end trainable networks (Wen et al., 2017b,a; Liu and Lane, 2018; Liu et al., 2018; Li et al., 2018; Zhong et al., 2018). Wen et al. (2017b) includes belief state tracking and has two phases in training: the first phase uses belief state supervision, and then the second phase uses response generation supervision. Wen et al. (2017a) improves Wen et al. (2017b) by adding a policy network using latent representations so that the dialogue system can be continuously improved through reinforcement learning. These methods utilize classification as a way to decode the belief state.

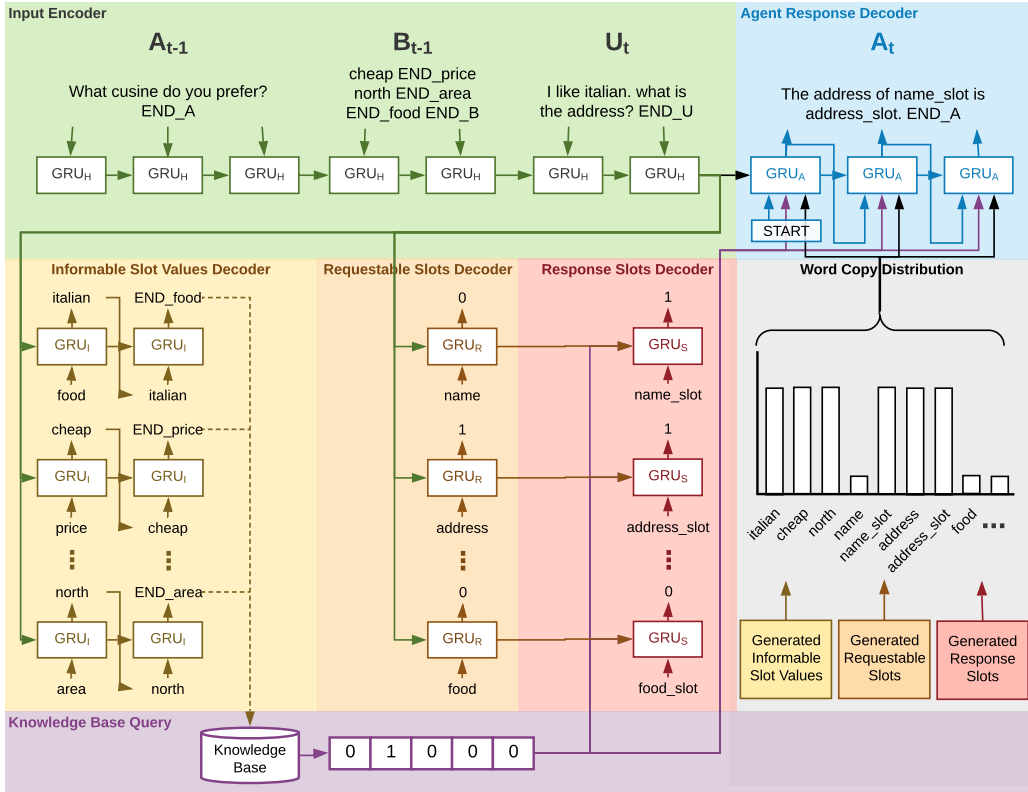


Figure 1: FSDM architecture illustrated by a dialogue turn from the Cambridge Restaurant dataset with the following components: an input encoder (green), a belief state tracker (yellow for the informable slot values, orange for the requestable slots), a KB query component (purple), a response slot classifier (red), a component that calculates word copy probability (grey) and a response decoder (blue). Attention connections are not drawn for brevity.

Lei et al. (2018) decode the belief state as well as the response in a free-form fashion, but it tracks the informable slot values without an explicit assignment to an informable slot. Moreover, the arbitrary order in which informable slot values and requestable slots are encoded and decoded suggests that the sequential inductive bias the architecture provides may not be the right one.

Other works (Jang et al., 2016; Henderson et al., 2014; Bapna et al., 2017; Kobayashi et al., 2018; Xu and Hu, 2018b) focus on the scalability of DST to large or changing vocabularies. Rastogi et al. (2017) score a dynamically defined set of candidates as informable slot values. Dernoncourt et al. (2016) addresses the problem of large vocabularies with a mix of rules and machine-learned classifiers.

3 Methodology

We propose a fully-fledged task-oriented dialogue system called Flexibly-Structured Dialogue Model (FSDM), which operates at the turn level. Its overall architecture is shown in Figure 1, which illustrates one dialogue turn. Without loss of generality, let us assume that we are on the t -th turn of a dia-

logue. FSDM has three (3) inputs: agent response and belief state of the $t - 1$ -th turn, and user utterance of the t -th turn. It has two (2) outputs: the belief state for the t -th turn that is used to query the KB, and the agent response of the t -th turn based on the query result. As we can see, belief tracking is the key component that turns unstructured user utterance and the dialogue history into a KB-friendly belief state. The success of retrieving the correct KB result and further generating the correct response to complete a task relies on the quality of the produced belief state.

FSDM contains five (5) components that work together in an end-to-end manner as follows: (1) The input is encoded and the last hidden state of the encoder serves as the initial hidden state of the belief state tracker and the response decoder; (2) Then, the belief state tracker generates a belief state $B_t = \{I_t, R_t\}$, where I_t is the set of constraints used for the KB query generated by the informable slots value decoder and R_t is the user requested slots identified by the requestable slots multi-label classifier; (3) Given I_t , the KB query component queries the KB and encodes the number of records

returned in a one-hot vector d_t ; (4) The response slot binary classifier predicts which slots should appear in the agent response S_t ; (5) Finally, the agent response decoder takes in the KB output d_t , a word copy probability vector \mathcal{P}^c computed from I_t , R_t , S_t together with an attention on hidden states of the input encoder and the belief decoders, and generates a response A_t .

3.1 Input Encoder

The input contains three parts: (1) the agent response A_{t-1} , (2) the belief state B_{t-1} from the $(t-1)$ -th turn and (3) the current user utterance U_t . These parts are all text-based and concatenated, and then consumed by the input encoder. Specifically, the belief state B_{t-1} is represented as a sequence of informable slot names with their respective values and requestable slot names. As an example, the sequence $\langle \text{cheap, end_price, italian, end_food, address, phone, end_belief} \rangle$ indicates a state where the user informed cheap and Italian as KB query constraints and requested the address and phone number.

The input encoder consists of an embedding layer followed by a recurrent layer with Gated Recurrent Units (GRU) (Cho et al., 2014). It maps the input $A_{t-1} \circ B_{t-1} \circ U_t$ (where \circ denotes concatenation) to a sequence of hidden vectors $\{h_i^E | i = 1, \dots, |A_{t-1} \circ B_{t-1} \circ U_t|\}$ so that $h_i^E = \text{GRU}_H(e^{A_{t-1} \circ B_{t-1} \circ U_t})$ where e is the embedding function that maps from words to vectors. The output of the input encoder is its last hidden state h_l^E , which is served as the initial state for the belief state and response decoders as discussed next.

3.2 Informable Slot Value Decoder

The belief state is composed of informable slot values I_t and the requestable slots R_t . We describe the generation of the former in this subsection and the latter in the next subsection.

The informable slot values track information provided by the user and are used to query the KB. We allow each informable slot to have its own decoder to resolve the unwanted artificial dependencies among slot values introduced by TSCP (Lei et al., 2018). As an example of artificial dependency, ‘italian; expensive’ appears a lot in the training data. During testing, even when the gold informable value is ‘italian; moderate’, the decoder may still generate ‘italian; expensive’. Modeling

one decoder for each slot exactly associates the values with the corresponding informable slot.

The informable slot value decoder consists of GRU recurrent layers with a copy mechanism as shown in the yellow section of Figure 1. It is composed of weight-tied GRU generators that take the same initial hidden state h_l^E , but have different start-of-sentence symbols for each unique informable slot. This way, each informable slot value decoder is dependent on the encoder’s output, but it is also independent of the values generated for the other slots. Let $\{k^I\}$ denote the set of informable slots. The probability of the j^{th} word $P(y_j^{k^I})$ being generated for the slot k^I is calculated as follows: (1) calculate the attention with respect to the input encoded vectors to obtain the context vector $c_j^{k^I}$, (2) calculate the generation score $\phi_g(y_j^{k^I})$ and the copy score $\phi_c(y_j^{k^I})$ based on the current step’s hidden state $h_j^{k^I}$, (3) calculate the probability using the copy mechanism:

$$\begin{aligned} c_j^{k^I} &= \text{Attn}(h_{j-1}^{k^I}, \{h_i^E\}), \\ h_j^{k^I} &= \text{GRU}_I((c_j^{k^I} \circ e^{y_j^{k^I}}), h_{j-1}^{k^I}), \\ \phi_g(y_j^{k^I}) &= W_g^{K^I} \cdot h_j^{k^I}, \\ \phi_c(y_j^{k^I}) &= \tanh(W_c^{K^I} \cdot h_j^{k^I}) \cdot h_j^{k^I}, \\ y_j^{k^I} &\in A_{t-1} \cup B_{t-1} \cup U_t, \\ P(y_j^{k^I} | y_{j-1}^{k^I}, h_{j-1}^{k^I}) &= \text{Copy}(\phi_c(y_j^{k^I}), \phi_g(y_j^{k^I})), \end{aligned} \quad (1)$$

where for each informable slot k^I , $y_0^{k^I} = k^I$ and $h_0^{k^I} = h_l^E$, $e^{y_j^{k^I}}$ is the embedding of the current input word (the one generated at the previous step), and $W_g^{K^I}$ and $W_c^{K^I}$ are learned weight matrices. We follow (Gu et al., 2016) and (Bahdanau et al., 2015) for the copy $\text{Copy}(\cdot, \cdot)$ and attention $\text{Attn}(\cdot, \cdot)$ mechanisms implementation respectively.

The loss for the informable slot values decoder is calculated as follows:

$$\begin{aligned} \mathcal{L}^I &= - \frac{1}{|\{k^I\}|} \frac{1}{|Y^{k^I}|} \sum_{k^I} \sum_j \\ &\log P(y_j^{k^I} = z_j^{k^I} | y_{j-1}^{k^I}, h_{j-1}^{k^I}), \end{aligned} \quad (2)$$

where Y^{K^I} is the sequence of informable slot value decoder predictions and z is the ground truth label.

3.3 Requestable Slot Binary Classifier

As the other part of a belief state, requestable slots are the attributes of KB entries that are explicitly requested by the user. We introduce a separate

multi-label requestable slots classifier to perform binary classification for each slot. This greatly resolves the issues of TSCP that uses a single decoder with each step having unconstrained vocabulary-size choices, which may potentially lead to generating non-slot words. Similar to the informable slots decoders, such a separate classifier also eliminates the undesired dependencies among slots.

Let $\{k^R\}$ denote the set of requestable slots. A single GRU cell is used to perform the classification. The initial state h_t^E is used to pay attention to the input encoder hidden vectors to compute a context vector c^{k^R} . The concatenation of c^{k^R} and e^{k^R} , the embedding vector of one requestable slot k^R , is passed as input and h_t^E as the initial state to the GRU. Finally, a sigmoid non-linearity is applied to the product of a weight vector W_y^R and the output of the GRU h^{k^R} to obtain y^{k^R} , which is the probability of the slot being requested by the user.

$$\begin{aligned} c^{k^R} &= \text{Attn}(h_t^E, \{h_i^E\}), \\ h^{k^R} &= \text{GRU}_R((c^{k^R} \circ e^{k^R}), h_t^E), \\ y^{k^R} &= \sigma(W_y^R \cdot h^{k^R}). \end{aligned} \quad (3)$$

The loss function for all requestable slot binary classifiers is:

$$\mathcal{L}^R = - \frac{1}{|\{k^R\}|} \sum_{k^R} z^{k^R} \log(y^{k^R}) + (1 - z^{k^R}) \log(1 - y^{k^R}). \quad (4)$$

3.4 Knowledge Base Query

The generated informable slot values $I_t = \{Y^{k^I}\}$ are used as constraints of the KB query. The KB is composed of one or more relational tables and each entity is a record in one table. The query is performed to select a subset of the entities that satisfy those constraints. For instance, if the informable slots are $\{price=cheap, area=north\}$, all the restaurants that have attributes of those fields equal to those values will be returned. The output of this component, the one-hot vector d_t , indicates the number of records satisfying the constraints. d_t is a five-dimensional one-hot vector, where the first four dimensions represent integers from 0 to 3 and the last dimension represents 4 or more matched records. It is later used to inform the response slot binary classifier and the agent response decoder.

3.5 Response Slot Binary Classifier

In order to incorporate all the relevant information about the retrieved entities into the response,

FSDM introduces a new response slot binary classifier. Its inputs are requestable slots and KB queried result d_t and the outputs are the response slots to appear in the agent response. Response slots are the slot names that are expected to appear in a de-lexicalized response (discussed in the next subsection). For instance, assume the requestable slot in the belief state is ‘‘address’’ and the KB query returned one candidate record. The response slot binary classifier may predict name_slot, address_slot and area_slot, which are expected to appear in an agent response as ‘‘name_slot is located in address_slot in the area_slot part of town’’².

The response slots $\{k^S\}$ map one-to-one to the requestable slots $\{k^R\}$. The initial state of each response slot decoder is the last hidden state of the corresponding requestable slot decoder. In this case, the context vector c^{k^S} is obtained by paying attention to all hidden vectors in the informable slot value decoders and requestable slots classifiers. Then, the concatenation of the context vector c^{k^S} , the embedding vector of the response slot e^{k^S} and the KB query vector d_t are used as input to a single GRU cell. Finally, a sigmoid non-linearity is applied to the product of a weight vector W_y^S and the output of the GRU h^{k^S} to obtain a probability y^{k^S} for each slot that is going to appear in the answer.

$$\begin{aligned} c^{k^S} &= \text{Attn}(h^{k^R}, \\ &\{h_i^{k^I} | k^I \in K^I, i \leq |Y^{k^I}|\} \cup \{h^{k^R} | k^R \in K^R\}), \\ h^{k^S} &= \text{GRU}_S((c^{k^S} \circ e^{k^S} \circ d_t), h^{k^R}), \\ y^{k^S} &= \sigma(W_y^S \cdot h^{k^S}). \end{aligned} \quad (5)$$

The loss function for all response slot binary classifiers is:

$$\mathcal{L}^S = - \frac{1}{|\{k^S\}|} \sum_{k^S} z^{k^S} \log(y^{k^S}) + (1 - z^{k^S}) \log(1 - y^{k^S}). \quad (6)$$

3.6 Word Copy Probability and Agent Response Decoder

Lastly, we introduce the agent response decoder. It takes in the generated informable slot values, requestable slots, response slots, and KB query result and generates a (de-lexicalized) response. We adopt a copy-augmented decoder (Gu et al., 2016) as architecture. The canonical copy mechanism only takes a sequence of word indexes as inputs but

² Before the agent response is presented to the user, those slot names are replaced by the actual values of the KB entries.

does not accept the multiple Bernoulli distributions we obtain from sigmoid functions. For this reason, we introduce a vector of independent word copy probabilities \mathcal{P}^C , which is constructed as follows:

$$\mathcal{P}^C(w) = \begin{cases} y^{k^R}, & \text{if } w = k^R, \\ y^{k^S}, & \text{if } w = k^S, \\ 1, & \text{if } w \in I_t, \\ 0, & \text{otherwise,} \end{cases} \quad (7)$$

where if a word w is a requestable slot or a response slot, the probability is equal to their binary classifier output; if a word appears in the generated informable slot values, its probability is equal to 1; for the other words in the vocabulary the probability is equal to 0. This vector is used in conjunction with the agent response decoder prediction probability to generate the response.

The agent response decoder is responsible for generating a de-lexicalized agent response. The response slots are substituted with the values of the results obtained by querying the KB before the response is returned to the user.

Like the informable slot value decoder, the agent response decoder also uses a copy mechanism, so it has a copy probability and generation probability. Consider the generation of the j^{th} word. Its generation score ϕ_g is calculated as:

$$\begin{aligned} c_j^{A^E} &= \text{Attn}(h_{j-1}^A, \{h_i^E\}), \\ c_j^{A^B} &= \text{Attn}(h_{j-1}^A, \{h_i^{k^I} | k^I \in K^I, i \leq |Y^{k^I}|\} \\ &\cup \{h^{k^R} | k^R \in K^R\} \cup \{h^{k^S} | k^S \in K^S\}), \\ h_j^A &= \text{GRU}_A((c_j^{A^E} \circ c_j^{A^B} \circ e_j^A \circ d_t), h_{j-1}^A), \\ \phi_g(y_j^A) &= W_g^A \cdot h_j^A, \end{aligned} \quad (8)$$

where $c_j^{A^E}$ is a context vector obtained by attending to the hidden vectors of the input encoder, $c_j^{A^B}$ is a context vector obtained by attending to all hidden vectors of the informable slot value decoder, requestable slot classifier and response slot classifier, and W_g^A is a learned weight matrix. The concatenation of the two context vectors $c_j^{A^E}$ and $c_j^{A^B}$, the embedding vector e_j^A of the previously generated word and the KB query output vector d_t is used as input to a GRU. Note that the initial hidden state is $h_0^A = h_l^E$. The copy score ϕ_c is calculated as:

$$\phi_c(y_j^A) = \begin{cases} \mathcal{P}^C(y_j^A) \cdot \tanh(W_c^A \cdot h_{y_j^A}^A) \cdot h_j^A, \\ \text{if } y_j^A \in I_t \cup K^R \cup K^S, \\ \mathcal{P}^C(y_j^A), \text{ otherwise,} \end{cases} \quad (9)$$

where W_c^A is a learned weight matrix. The final

CamRest: restaurant reservation			
dialogue split	train: 408	dev: 136	test: 136
# of keys	informable: 3	requestable: 7	response: 7
database record	99		
KVRET: navigation, weather, calendar scheduling			
dialogue split	train: 2425	dev: 302	test: 302
# of keys	informable: 10	requestable: 12	response: 12
database record	284		

Table 1: Dataset

probability is:

$$P(y_j^A | y_{j-1}^A, h_{j-1}^A) = \text{Copy}(\phi_g(y_j^A), \phi_c(y_j^A)). \quad (10)$$

Let z denote the ground truth de-lexicalized agent response. The loss for the agent response decoder is calculated as follows where Y^A is the sequence of agent response decoder prediction:

$$\mathcal{L}^A = -\frac{1}{|Y^A|} \sum_j \log P(y_j^A = z_j^A | y_{j-1}^A, h_{j-1}^A). \quad (11)$$

3.7 Loss Function

The loss function of the whole network is the sum of the four losses described so far for the informable slot values \mathcal{L}^I , requestable slot \mathcal{L}^R , response slot \mathcal{L}^S and agent response decoders \mathcal{L}^A , weighted by α hyperparameters:

$$\mathcal{L} = \alpha^I \mathcal{L}^I + \alpha^R \mathcal{L}^R + \alpha^S \mathcal{L}^S + \alpha^A \mathcal{L}^A. \quad (12)$$

The loss is optimized in an end-to-end fashion, with all modules trained simultaneously with loss gradients back-propagated to their weights. In order to do so, ground truth results from database queries are also provided to the model to compute the d_t , while at prediction time results obtained by using the generated informable slot values I_t are used.

4 Experiments

We tested the FSDM on the Cambridge Restaurant dataset (CamRest) (Wen et al., 2017b) and the Stanford in-car assistant dataset (KVRET) (Eric et al., 2017) described in Table 1.

4.1 Preprocessing and Hyper-parameters

We use NLTK (Bird et al., 2009) to tokenize each sentence. The user utterances are precisely the original texts, while all agent responses are de-lexicalized as described in (Lei et al., 2018). We obtain the labels for the response slot decoder from the de-lexicalized response texts. We use 300-dimensional GloVe embeddings (Pennington et al., 2014) trained on 840B words. Tokens not present

Dataset	CamRest						KVRET					
	Inf P	Inf R	Inf F ₁	Req P	Req R	Req F ₁	Inf P	Inf R	Inf F ₁	Req P	Req R	Req F ₁
TSCP/RL [†]	0.970	0.971	0.971	0.983	0.935	0.959	0.936	0.874	0.904	0.725	0.485	0.581
TSCP [†]	0.970	0.971	0.971	0.983	0.938	0.960	0.934	0.890	0.912	0.701	0.435	0.526
FSDM/Res	0.979	0.984	0.978	0.994	0.947	0.967	0.918	0.930	0.925	0.812	0.993	0.893
FSDM	0.983*	0.986*	0.984*	0.997*	0.952	0.974*	0.92	0.935*	0.927*	0.819*	1.000*	0.900*

Table 2: Turn-level performance results. **Inf**: Informable, **Req**: Requestable, **P**: Precision, **R**: Recall. Results marked with [†] are computed using available code, and all the other ones are reported from the original papers. * indicates the improvement is statistically significant with $p = 0.05$.

Dataset	CamRest			KVRET		
	BLEU	EMR	SuccF ₁	BLEU	EMR	SuccF ₁
NDM	0.212	0.904	0.832	0.186	0.724	0.741
LIDM	0.246	0.912	0.840	0.173	0.721	0.762
KVRN	0.134	-	-	0.184	0.459	0.540
TSCP	0.253	0.927	0.854	0.219	0.845	0.811
TSCP/RL [†]	0.237	0.915	0.826	0.195	0.809	0.814
TSCP [†]	0.237	0.913	0.841	0.189	0.833	0.81
FSDM/St	0.245	-	0.847	0.204	-	0.809
FSDM/Res	0.251	0.924	0.855	0.209	0.834	0.815
FSDM	0.258*	0.935*	0.862*	0.215	0.848*	0.821*

Table 3: Dialogue level performance results. **SuccF₁**: Success F₁ score, **EMR**: Entity Match Rate. Results marked with [†] are computed using available code, and all the other ones are reported from the original papers. * indicates the improvement is statistically significant with $p = 0.05$.

user msg	what is the date and time of my next meeting and who will be attending it ?
	belief state
GOLD	informable slot (event=meeting), requestable slot (date, time, party)
TSCP	'meeting' '(EOS_Z1)' 'date' ',' 'party'
FSDM	event=meeting date=True time=True party = True
	agent response
GOLD	your next meeting is with party_SLOT on the date_SLOT at time_SLOT.
TSCP	your next meeting is at time_SLOT on date_SLOT at time_SLOT .
FSDM	you have a meeting on date_SLOT at time_SLOT with party_SLOT

Table 4: Example of generated belief state and response for calendar scheduling domain

in GloVe are initialized to be the average of all other embeddings plus a small amount of random noise to make them different from each other.

We optimize both training and model hyperparameters by running Bayesian optimization over the product of validation set BLEU, EMR, and SuccF₁ using skopt³. The model that performed the best on the validation set uses Adam optimizer (Kingma and Ba, 2015) with a learning rate of 0.00025 for minimizing the loss in Equation 12 for both datasets. We apply dropout with a rate of 0.5 after

³<https://scikit-optimize.github.io/>

the embedding layer, the GRU layer and any linear layer for CamRest and 0.2 for KVRET. The dimension of all hidden states is 128 for CamRest and 256 for KVRET. Loss weights α^I , α^R , α^S , α^A are 1.5, 9, 8, 0.5 respectively for CamRest and 1, 3, 2, 0.5 for KVRET.

4.2 Evaluation Metrics

We evaluate the performance concerning belief state tracking, response language quality, and task completion. For belief state tracking, we report precision, recall, and F₁ score of informable slot values and requestable slots. BLEU (Papineni et al., 2002) is applied to the generated agent responses for evaluating language quality. Although it is a poor choice for evaluating dialogue systems (Liu et al., 2016), we still report it in order to compare with previous work that has adopted it. For task completion evaluation, the Entity Match Rate (EMR) (Wen et al., 2017b) and Success F₁ score (SuccF₁) (Lei et al., 2018) are reported. EMR evaluates whether a system can correctly retrieve the user’s indicated entity (record) from the KB based on the generated constraints so it can have only a score of 0 or 1 for each dialogue. The SuccF₁ score evaluates how a system responds to the user’s requests at dialogue level: it is F₁ score of the response slots in the agent responses.

4.3 Benchmarks

We compare FSDM with four baseline methods and two ablations.

NDM (Wen et al., 2017b) proposes a modular end-to-end trainable network. It applies delexicalization on user utterances and responses.

LIDM (Wen et al., 2017a) improves over NDM by employing a discrete latent variable to learn underlying dialogue acts. This allows the system to be refined by reinforcement learning.

KVRN (Eric et al., 2017) adopts a copy-augmented Seq2Seq model for agent response generation and uses an attention mechanism on the KB.

It does not perform belief state tracking.

TSCP/RL (Lei et al., 2018) is a two-stage Copy-Net which consists of one encoder and two copy-mechanism-augmented decoders for belief state and response generation. **TSCP** includes further parameter tuning with reinforcement learning to increase the appearance of response slots in the generated response. We were unable to replicate the reported results using the provided code⁴, hyperparameters, and random seed, so we report both the results from the paper and the average of 5 runs on the code with different random seeds (marked with †).

FSDM is the proposed method and we report two ablations: in **FSDM/St** the whole state tracking is removed (informable, requestable and response slots) and the answer is generated from the encoding of the input, while in **FSDM/Res**, only the response slot decoder is removed.

4.4 Result Analysis

At the turn level, FSDM and FSDM/Res perform better than TSCP and TSCP/RL on belief state tracking, especially on requestable slots, as shown in Table 2. FSDM and FSDM/Res use independent binary classifiers for the requestable slots and are capable of predicting the correct slots in all those cases. FSDM/Res and TSCP/RL do not have any additional mechanism for generating response slot, so FSDM/Res performing better than TSCP/RL shows the effectiveness of flexible-structured belief state tracker. Moreover, FSDM performs better than FSDM/Res, but TSCP performs worse than TSCP/RL. This suggests that using RL to increase the appearance of response slots in the response decoder does not help belief state tracking, but our response slot decoder does.

FSDM performs better than all benchmarks on the dialogue level measures too, as shown in Table 3, with the exception of BLEU score on KVRET, where it is still competitive. Comparing TSCP/RL and FSDM/Res, the flexibly-structured belief state tracker achieves better task completion than the free-form belief state tracker. Furthermore, FSDM performing better than FSDM/Res shows the effectiveness of the response slot decoder for task completion. The most significant performance improvement is obtained on CamRest by FSDM, confirming that the additional inductive bias helps to generalize from smaller datasets. More impor-

tantly, the experiment confirms that, although making weaker assumptions that are reasonable for real-world applications, FSDM is capable of performing at least as well as models that make stronger limiting assumptions which make them unusable in real-world applications.

4.5 Error Analysis

We investigated the errors that both TSCP and FSDM make and discovered that the sequential nature of the TSCP state tracker leads to the memorization of common patterns that FSDM is not subject to. As an example (Table 4), TSCP often generates “date; party” as requestable slots even if only “party” and “time” are requested like in “what time is my next activity and who will be attending?” or if “party”, “time” and “date” are requested like in “what is the date and time of my next meeting and who will be attending it?”. FSDM produces correct belief states in these examples.

FSDM misses some requestable slots in some conditions. For example, consider the user’s utterance: “I would like their address and what part of town they are located in”. The ground-truth requestable slots are ‘address’ and ‘area’. FSDM only predicts ‘address’ and misses ‘area’, which suggests that the model did not recognize ‘what part of town’ as being a phrasing for requesting ‘area’. Another example is when the agent proposes “the name_SLOT is moderately priced and in the area_SLOT part of town . would you like their location ?” and the user replies “i would like the location and the phone number, please”. FSDM predicts ‘phone’ as a requestable slot, but misses ‘address’, suggesting it doesn’t recognize the connection between ‘location’ and ‘address’. The missing requestable slot issue may propagate to the agent response decoder. These issues may arise due to the use of fixed pre-trained embeddings and the single encoder. Using separate encoders for user utterance, agent response and dialogue history or fine-tuning the embeddings may solve the issue.

5 Conclusion

We propose the flexibly-structured dialogue model, a novel end-to-end architecture for task-oriented dialogue. It uses the structure in the schema of the KB to make architectural choices that introduce inductive bias and address the limitations of fully structured and free-form methods. The experiment suggests that this architecture is competitive

⁴<https://github.com/WING-NUS/sequicity>

with state-of-the-art models, while at the same time providing a more practical solution for real-world applications.

Acknowledgments

We would like to thank Alexandros Papangelis, Janice Lam, Stefan Douglas Webb and SIGDIAL reviewers for their valuable comments.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations, San Diego, California, USA*.
- Ankur Bapna, Gökhan Tür, Dilek Z. Hakkani-Tür, and Larry P. Heck. 2017. Towards zero-shot frame semantic parsing for domain scaling. In *INTER-SPEECH*.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural Language Processing with Python*. O'Reilly Media, Inc.
- Antoine Bordes, Y-Lan Boureau, and Jason Weston. 2017. Learning end-to-end goal-oriented dialog. In *International Conference on Learning Representations, Toulon, France*.
- Paweł Budzianowski, Iigo Casanueva, Bo-Hsiang Tseng, and Milica Gai. 2018. Towards end-to-end multi-domain dialogue modelling. *Technical Report CUED/F-INFENG/TR.706, Cambridge University*.
- Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *EMNLP*, pages 1724–1734. ACL.
- Franck Dernoncourt, Ji Young Lee, Trung H. Bui, and Hung H. Bui. 2016. Robust dialog state tracking for large ontologies. In *IWSDS*, volume 427 of *Lecture Notes in Electrical Engineering*, pages 475–485. Springer.
- Mihail Eric, Lakshmi Krishnan, Francois Charette, and Christopher D. Manning. 2017. Key-value retrieval networks for task-oriented dialogue. In *SIGDIAL Conference*, pages 37–49. Association for Computational Linguistics.
- Pascale Fung, Chien-Sheng Wu, and Andrea Madotto. 2018. Mem2seq: Effectively incorporating knowledge bases into end-to-end task-oriented dialog systems. In *ACL (1)*, pages 1468–1478. Association for Computational Linguistics.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O. K. Li. 2016. Incorporating copying mechanism in sequence-to-sequence learning. In *ACL (1)*. The Association for Computer Linguistics.
- Matthew Henderson, Blaise Thomson, and Steve J. Young. 2014. Robust dialog state tracking using delexicalised recurrent neural networks and unsupervised adaptation. *2014 IEEE Spoken Language Technology Workshop (SLT)*, pages 360–365.
- Takaaki Hori, Hai Wang, Chiori Hori, Shinji Watanabe, Bret Harsham, Jonathan Le Roux, John R. Hershey, Yusuke Koji, Yi Jing, Zhaocheng Zhu, and Takeyuki Aikawa. 2016. Dialog state tracking with attention-based sequence-to-sequence learning. In *SLT*, pages 552–558. IEEE.
- Youngsoo Jang, Jiyeon Ham, Byung-Jun Lee, Youngjae Chang, and Kee-Eung Kim. 2016. Neural dialog state tracker for large ontologies by attention mechanism. *2016 IEEE Spoken Language Technology Workshop (SLT)*, pages 531–537.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations, San Diego, California, USA*.
- Yuka Kobayashi, Takami Yoshida, Kenji Iwata, Hiroshi Fujimura, and Masami Akamine. 2018. Out-of-domain slot value detection for spoken dialogue systems with context information. *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 854–861.
- Wenqiang Lei, Xisen Jin, Min-Yen Kan, Zhaochun Ren, Xiangnan He, and Dawei Yin. 2018. Sequicity: Simplifying task-oriented dialogue systems with single sequence-to-sequence architectures. In *ACL*.
- Xiujun Li, Sarah Panda, Jingjing Liu, and Jianfeng Gao. 2018. Microsoft dialogue challenge: Building end-to-end task-completion dialogue systems. volume abs/1807.11125.
- Bing Liu and Ian Lane. 2018. End-to-end learning of task-oriented dialogs. In *Proceedings of the NAACL-HLT*.
- Bing Liu, Gokhan Tur, Dilek Hakkani-Tur, Pararth Shah, and Larry Heck. 2018. Dialogue learning with human teaching and feedback in end-to-end trainable task-oriented dialogue systems. In *NAACL*.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Michael Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *EMNLP*, pages 2122–2132. The Association for Computational Linguistics.
- Ryan Thomas Lowe, Nissan Pow, Iulian Vlad Serban, Laurent Charlin, Chia-Wei Liu, and Joelle Pineau. 2017. Training end-to-end dialogue systems with

- the ubuntu dialogue corpus. *Dialogue and Discourse*, 8(1):31–65.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*, pages 311–318. ACL.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543. ACL.
- Osman Ramadan, Paweł Budzianowski, and Milica Gasic. 2018. Large-scale multi-domain belief tracking with knowledge sharing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 432–437.
- Abhinav Rastogi, Dilek Hakkani-Tur, and Larry Heck. 2017. Scalable multi-domain dialogue state tracking. In *Proceedings of IEEE ASRU*.
- Liliang Ren, Kaige Xie, Lu Chen, and Kai Yu. 2018. Towards universal dialogue state tracking. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2780–2786.
- Iulian Vlad Serban, Alessandro Sordani, Yoshua Bengio, Aaron C. Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *AAAI*, pages 3776–3784. AAAI Press.
- Sainbayar Sukhbaatar, Arthur Szlam, Jason Weston, and Rob Fergus. 2015. End-to-end memory networks. In *NIPS*, pages 2440–2448.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*, pages 3104–3112.
- Haoyang Wen, Yijia Liu, Wanxiang Che, Libo Qin, and Ting Liu. 2018. Sequence-to-sequence learning for task-oriented dialogue with dialogue state representation. pages 3781–3792.
- Tsung-Hsien Wen, Yishu Miao, Phil Blunsom, and Steve J. Young. 2017a. Latent intention dialogue models. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 3732–3741. PMLR.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gasic, Lina M Rojas Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017b. A network-based end-to-end trainable task-oriented dialogue system. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, volume 1, pages 438–449. ACL.
- Jason D. Williams, Kavosh Asadi, and Geoffrey Zweig. 2017. Hybrid code networks: practical and efficient end-to-end dialog control with supervised and reinforcement learning. In *ACL (1)*, pages 665–677. Association for Computational Linguistics.
- Puyang Xu and Qi Hu. 2018a. An end-to-end approach for handling unknown slot values in dialogue state tracking. In *ACL (1)*, pages 1448–1457. Association for Computational Linguistics.
- Puyang Xu and Qi Hu. 2018b. An end-to-end approach for handling unknown slot values in dialogue state tracking. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1448–1457. Association for Computational Linguistics.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2018. Global-locally self-attentive encoder for dialogue state tracking. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1458–1467. Association for Computational Linguistics.

FriendsQA: Open-Domain Question Answering on TV Show Transcripts

Zhengzhe Yang

Computer Science
Emory University
Atlanta, GA, USA

zhengzhe.yang@emory.edu

Jinho D. Choi

Computer Science
Emory University
Atlanta, GA, USA

jinho.choi@emory.edu

Abstract

This paper presents FriendsQA, a challenging question answering dataset that contains 1,222 dialogues and 10,610 open-domain questions, to tackle machine comprehension on everyday conversations. Each dialogue, involving multiple speakers, is annotated with several types of questions regarding the dialogue contexts, and the answers are annotated with certain spans in the dialogue. A series of crowdsourcing tasks are conducted to ensure good annotation quality, resulting a high inter-annotator agreement of 81.82%. A comprehensive annotation analytics is provided for a deeper understanding in this dataset. Three state-of-the-art QA systems are experimented, R-Net, QANet, and BERT, and evaluated on this dataset. BERT in particular depicts promising results, an accuracy of 74.2% for answer utterance selection and an F1-score of 64.2% for answer span selection, suggesting that the FriendsQA task is hard yet has a great potential of elevating QA research on multiparty dialogue to another level.

1 Introduction

Question answering (QA) has received lots of hype over the recent years as deep learning models have progressively pushed the limit of machine comprehension to the level of human intelligence. Several systems have demonstrated their superiority over human for answering quizbowl questions (Ferrucci, 2011; Yamada et al., 2017). Strong evidences have been found that advance neural network models will likely surpass human performance for answering open-domain questions in a foreseeable future (Devlin et al., 2018; Liu et al., 2019). Nonetheless, no system has reached such high intelligence for understanding contexts in dialogue, although it is the most natural means of human communication. Moreover, the amount of data in this form has increased at a faster rate than any other type of textual data (Newport, 2014; Gonçalves, 2017).

Many datasets have been presented for various QA tasks (Section 2.1). While numerous models have shown remarkable results with these datasets (Section 2.2), the evidence passages, where the contexts of questions are derived from, mostly reside within wiki articles, newswire, (non-)fictional stories, or children’s books, but not from multiparty dialogue. Contextual understanding in dialogue is challenging because it needs to interpret contents composed by multiple speakers, and anticipate colloquial language filled with sarcasms, metaphors, humors, etc. This inspires us to create a new dataset, FriendsQA, that aims to enhance machine comprehension on this domain. Dialogues in this dataset are excerpted from transcripts of the TV show *Friends*, that is the world-wide and also go-to show for English learners to get familiarized with everyday conversations.

Section 3 describes the FriendsQA dataset with annotation details. Section 4 describes the architectures of QA systems experimented on this dataset. Finally, Section 5 shows the experimental results with an in-depth error analysis. To the best of our knowledge, FriendsQA is the first dataset that is publicly available and challenges span-based QA on multiparty dialogue with everyday topics. The contributions of this work include:

- An open-domain question answering dataset on multiparty dialogue comprising 1,222 dialogues, 10,610 questions, and 21,262 answer spans.
- A comprehensive corpus analytics to ensure its validity as a deep learning resource and explain the diverse nature of this dataset for QA.
- Model comparisons between three state-of-the-art QA systems trained on this dataset to project its practicality in real applications.
- A thorough error analysis to illustrate major challenges found in this task and make suggestions to future research on the dialogue domain.

2 Related Work

2.1 QA Datasets

The NLP community has been dedicated to produce three types of question answering (QA) datasets. The first is for reading comprehension QA, where the model picks answers for multiple choice questions regarding the evidence passages. MCTest is an open-domain dataset comprising short fictional stories (Richardson et al., 2013). RACE is a large dataset compiled from English assessments for 12-18 years old students (Lai et al., 2017). TQA gives passages from middle school science lessons and textbooks (Kembhavi et al., 2017). SciQ gives passages from science exams collected via crowdsourcing (Welbl et al., 2017). DREAM gives multi-party dialogue passages from English-as-a-foreign-language exams (Sun et al., 2019).

The second is for cloze-style QA, for which the model fills in the blanks that obliterate certain contents in sentences describing the evidence passages. CNN/Daily Mail targets on entities in bullet points summarizing articles from *CNN* and *Daily News* (Hermann et al., 2015). Children’s Book Test focuses on named entities, nouns, verbs, and prepositions in passages from children’s books (Hill et al., 2016). Who-did-What gives description sentences and evidence passages extracted from news articles in English Gigaword (Onishi et al., 2016). Book-Test is similar to Children’s Book Test but 60 times larger (Bajgar et al., 2016).

The third is for span-based QA, where the model finds the answer contents as spans in the evidence passages. bAbI aims to reinforce learning on event types and infer a sequence of event descriptions (Weston et al., 2016). WikiQA (Yang et al., 2015) and SQuAD (Rajpurkar et al., 2016) use Wikipedia, whereas NewsQA (Trischler et al., 2017) use CNN articles as evidence passages. MS MARCO gives questions involving zero to multiple answer contents from web documents (Nguyen et al., 2016). TriviaQA is compiled by trivia enthusiasts to challenge machine comprehension (Joshi et al., 2017). CoQA focuses on conversational flows between a questioner and an answerer (Reddy et al., 2018).

2.2 QA Systems for the Past Two Years

Wang et al. (2017) presented R-Net that used gated attention-based recurrent networks and refined QA representation with self-matching attention. Shen et al. (2017) presented ReasonNet that took multiple turns to reason over the relationships between

query, documents, and answers. Cui et al. (2017) presented the Attention Over Attention Reader to better capture similarities between questions and answer contents. Hu et al. (2017) presented the Reinforced Mnemonic Reader to combine the memorized attention with new attention. Vaswani et al. (2017) applied self-attention to QA, which became known as the Transformer.

Huang et al. (2018) presented FusionNet that kept the history of word representations and used multi-level attention. Salant and Berant (2018) presented a standard neural architecture with rich contextualized word representations. Liu et al. (2018) presented Stochastic Answer Network (SAN) with a stochastic prediction dropout layer as the final layer. Yu et al. (2018) presented QANet with CNN and self-attention to combine local and global interactions. Peters et al. (2018) presented the Embeddings from Language Models (ELMo) that used bi-directional LSTM and Devlin et al. (2018) presented the Bidirectional Encoder Representations (BERT) that used deep-layered transformers to generate contextualized word embeddings.

2.3 Character Mining

The Character Mining dataset provides transcripts of the TV show *Friends* as well as annotation for several tasks. Chen and Choi (2016) annotated the first two seasons for character identification, that is an entity linking task identifying personal mentions with character names. Chen et al. (2017) extended this annotation to the next two seasons and added annotation of ambiguous mentions. Zhou and Choi (2018) added annotation of plural mentions to those four seasons for character identification. Zahiri and Choi (2018) annotated the first four seasons for fine-grained emotion detection. Finally, Ma et al. (2018) annotated selected dialogues from all ten seasons for a cloze-style reading comprehension task.

2.4 FriendsQA vs. Other Dialogue QA

Three datasets have been presented for QA on dialogue. CoQA (Reddy et al., 2018) aims to answer questions that are part of one-to-one conversations, whereas FriendsQA focuses on questions asked by third-parties listening to multiparty dialogues. Ma et al. (2018) also provides a dataset based on transcripts of *Friends*; however, their work aims to cloze-style QA restricted by PERSON entities, while we broadly focus on span-based QA with open-domain questions. Similarly, DREAM (Sun et al., 2019), although their passages are based on

(a) Challenges with entity resolution. In this example (season 4, episode 12), $\{you_1, boys_2, us_3\}$ refer to the boys and $\{you_4, we_8\}$ refer to the girls. Many pronouns are used to refer different people, which makes it difficult to find the answer span for a question like “*who forced Rachel to raise the stakes*” by simply matching strings.

Rachel	Y’know what, you ₁ are mean boys ₂ , who are just being mean!
Joey	Hey, don’t get mad at us ₃ ! No one forced you ₄ to raise the stakes!
Rachel	That is not true. She ₅ did! She ₆ forced me ₇ !
Monica	Hey, we ₈ would still be living here if you ₉ hadn’t gotten the question wrong!

(b) Challenges with metaphors. In this example (season 1, episode 4), Joey mishears ‘*omnipotent*’ as ‘*I’m impotent*’ so that he metaphorically refers it to as “*Little Joey’s dead*”, which makes it difficult to answer a question like “*why would Joey want to kill himself for being omnipotent*”.

Monica	Hey, Joey, what would you do if you were omnipotent ?
Joey	Probably kill myself!
Monica	Excuse me?
Joey	Hey, if Little Joey’s dead , then I got no reason to live!

(c) Challenges with sarcasm. In this example (season 3, episode 1), Chandler is being sarcastic about him making pancakes, which makes it difficult to answer a question like “*did Chandler make pancakes*”.

Chandler	Morning.
Joey	Morning, hey, you made pancakes?
Chandler	Yeah, like there’s any way I could ever do that.

Table 1: Challenges with entity resolution, metaphors, and sarcasm in understanding dialogue contexts for QA.

dialogue, tackles multiple-choice questions, which suit well for evaluating reading comprehension, but not necessarily for practical QA applications.

3 FriendsQA Dataset

For the generation of the FriendsQA dataset, 1,222 scenes from the first four seasons of the Character Mining dataset are selected (Section 2.3). Scenes with fewer than five utterances are discarded (83 of them), and each scene is considered an independent dialogue. FriendQA can be viewed as answer span selection, where questions are asked for some contexts in a dialogue and the model is expected to find certain spans in the dialogue containing answer contents. The dialogue aspects of this dataset, however, make it more challenging than other datasets comprising passages in formal languages (Section 2.1). Three challenging aspects that are commonly found in dialogue QA are illustrated in Table 1.

3.1 Crowdsourcing

All annotation tasks are conducted on the Amazon Mechanical Turk. TALEN, a web-based tool for named entity annotation (Mayhew and Roth, 2018), is extended for our QA annotation such that it displays a dialogue segmented into a sequence of utterances with speaker names, and asks crowd workers to first generate questions then select spans or utterance IDs in the dialogue containing the answer contents (Section 3.2). Prior to the annotation, crowd workers are required to pass a quiz regarding the dialogue context, to verify if they have a good un-

derstanding in this context. Upon the submission, it validates the annotation by running several quality assurance tests (Section 3.3).

3.2 Phase 1: Question-Answer Generation

For each dialogue, the crowd workers are required to generate at least 4 out of six types of questions, $\{who, what, when, where, why, how\}$, regarding the dialogue contexts. Every question must be answerable; in other words, there needs to be at least one contiguous answer span in the dialogue. The crowd workers are allowed to select more than one answer span per question if appropriate. If multiple mentions of the same entity are to be considered, annotators are instructed to select ones that fit the best for the question. For Q2 in Table 2, although multiple mentions of *Casey* are found in this dialogue, only the first three are selected as the answer because the other mentions are not relevant to this particular question (e.g., *Casey* in U08). This type of selective answer spans adds another level of difficulty to the task of FriendsQA.

Annotators are also allowed to select the speaker names as the answer spans. This is useful for *who* questions asking about certain speakers yet no mentions of them are found in the dialogue (e.g., *Chandler* has no explicit mention in Table 2). Moreover, when an entire utterance is considered the answer, which happens often with *why* and *how* questions, annotators are asked to select the corresponding utterance ID instead of the whole utterance to reduce span-related errors (e.g., U13 for Q5 in Table 2).

(a) A dialogue excerpted from *Friends* (season 4, episode 7).

U01	[Scene: <u>Central Perk</u> , Joey is getting a phone number from a woman (<u>Casey</u>) as Chandler watches from the doorway.]
U02	<u>Casey</u> : Here you go.
U03	Joey: Great! All right, so I'll call you later.
U04	<u>Casey</u> : Great!
U05	Chandler: Hey-Hey-Hey! Who was that?
U06	Joey: That would be Casey. <u>We're going out tonight.</u>
U07	Chandler: Goin' out, huh? Wow! Wow! So things didn't work out with Kathy, huh? Bummer.
U08	Joey: No, <u>things are fine with Kathy.</u> I'm having a late dinner with her <u>tonight</u> , right after <u>my early dinner with Casey.</u>
U09	Chandler: What?
U10	Joey: Yeah-yeah. And the craziest thing is that I just ate a whole pizza by myself!
U11	Chandler: Wait! You're going out with Kathy!
U12	Joey: Yeah. Why are you getting so upset?
U13	Chandler: Well, I'm upset for you. I mean, dating an endless line of beautiful women must be very unfulfilling for you.

(b) Six types of questions: {*who, what, when, where, why, how*}.

Q1	<u>What</u> is Joey going to do with Casey tonight?	Q4	<u>Where</u> are Joey and Chandler?
Q2	<u>Who</u> is Joey getting a phone number from?	Q5	<u>Why</u> is Chandler upset?
Q3	<u>When</u> will Joey have dinner with Kathy?	Q6	<u>How</u> are things between Joey and Kathy?

Table 2: A sample dialogue from the FriendsQA dataset comprising six types of questions, where the answer spans are annotated on the dialogue contents. Each utterance has the utterance ID, the speaker name, and the text. The answer spans for Q[1-6] are indicated by solid underlines, wavy underlines, double underlines, dashed underlines, **bold font**, and dotted underlines, respectively.

3.3 Quality Assurance

Each MTurk annotation job gives up to 6 questions and their answer spans, which are validated by the following tests before the submission:

1. Are there at least 4 types of questions annotated?
2. Does each question have at least one answer span associated with it?
3. Does any question have too much string overlaps with the original text in the dialogue?

The first test ensures that there are sufficiently large and diverse enough questions generated for developing practical QA models. The second test checks if there are any inappropriate associations between questions and answer spans. Finally, the third test prevents from creating mundane questions by copying and pasting the original text from the dialogue. No annotation job is accepted unless it passes all of these assurance tests.

3.4 Phase 2: Verification and Paraphrasing

All dialogues with the questions and answer spans annotated by the first phase (Section 3.2) are again put to the second phase. During the second phase, annotators are asked to first verify whether or not the answer spans are appropriate for the questions, and fix ones that are not or add more if necessary. Annotators are then asked to revise questions that

are either unanswerable or too ambiguous. Finally, they are asked to paraphrase the questions, resulting two sets of questions for every dialogue where one is a paraphrase of the other. The same quality assurance tests (Section 3.3) with an additional test of checking string overlaps between the questions from phases 1 and 2 are run to preserve the challenging level of this dataset.

3.5 Four Rounds of Annotation

The same F1-score metric used for the evaluation of span-based QA systems (Rajpurkar et al., 2016) is used to measure the inter-annotation agreement (ITA) between the answer spans annotated by the phases 1 and 2 (Sections 3.2 and 3.4, respectively). Four rounds of crowdsourcing tasks are conducted to stabilize the quality of our annotation, where two randomly selected episodes from Seasons 1-4 are used for annotation, respectively. After each round, ITA is measured and a sample set of annotation is manually checked. Then, the annotation guidelines are updated based on this assessment. The column A from the rows R1 ~ R4 in Table 3 illustrates the progressive ITA improvements over these four rounds. The followings show summaries of actions performed after each round (R[1-4]: round 1-4):

R1 We observe that the questions are often too ambiguous for humans to answer; thus, we update the guidelines and request annotators to make the questions as explicit as possible.

	S	Q	Q _p	Q _r	A	A _p	F1	F1 _p	EM	EM _p
R1	24	122	98	62	264	216	66.59	83.42	48.15	61.17
R2	26	242	185	57	484	368	72.86	83.99	50.00	57.69
R3	30	264	213	66	528	422	75.34	83.12	48.92	53.97
R4	37	370	296	75	740	593	76.01	88.17	52.25	60.78
S1	288	2,908	2,560	627	5,824	5,123	69.93	79.78	42.78	49.01
S2	259	2,682	2,314	587	5,372	4,633	69.21	80.86	44.01	51.73
S3	291	2,908	2,546	610	5,826	5,099	72.12	81.92	47.22	53.88
S4	267	2,768	2,398	594	5,553	4,808	72.26	83.27	49.52	57.41
Total	1,222	12,264	10,610	2,678	2,4591	21,262	71.17	81.82	46.35	53.55

Table 3: Statistics of the FriendsQA dataset. The R[1-4] rows show the statistics for the rounds 1-4, and the S[1-4] rows show the statistics for Seasons 1-4, respectively. S: # of dialogues, Q: # of questions, Q_p: Q after pruning, Q_r: # of revised questions during phase 2, A: # of answer spans, A_p: A after pruning, F1: F1-score to measure ITA, F1_p: F1 after pruning, EM: exact matching score to measure ITA, EM_p: EM after pruning.

R2 We observe the 6.27% improvement on ITA from the first round; thus, we add more examples of questions and answer spans to the guidelines without updating other contents.

R3 We observe another 2.48% improvement on ITA from the second round; no update is made to the guidelines.

R4 We observe a marginal ITA improvement of 0.67% from the third round, which implies that our annotation guidelines are stabilized. Thus, all of the rest episodes are pushed for annotation.

3.6 Question/Answer Pruning

Once all annotation is collected, each question from phase 1 is represented by the bag-of-words model using TF-IDF scores and compared against its revised counterpart from phase 2 if available. About 21.8% of the questions from phase 1 are revised during phase 2. If the cosine similarity between the two questions is below 0.8, they are not considered similar so that the question and its answer spans from phase 1 are discarded because that question requires a major revision to be answerable. Even when the questions are considered similar, if the F1 score between their answer spans is below 20, they are still discarded because annotators do not seem to agree on the answer. As a result, 13.5% of the questions and answer spans from phase 1 are pruned out from our final dataset.

3.7 Inter-annotator Agreement

Table 3 show the overall statistics of the FriendsQA dataset. There is a total of 1,222 dialogues, 10,610 questions, and 21,262 answer spans in this dataset after pruning (Section 3.6). Note that annotators were not asked to paraphrase questions during the second phase of the first round (R1 in Table 3), so

the number of questions in R1 is about twice less than ones from the other rounds. The final inter-annotator agreement scores are 81.82% and 53.55% for the F1 and exact matching scores respectively, indicating high-quality annotation in our dataset.

3.8 Question Types vs. Answer Categories

Table 4 shows the statistics between the question types and answer categories, where answers to each question type are further analyzed into categories. Questions show balanced distributions across different types, implying good diversity of the dataset. The analysis of answer categorization is performed manually among 250 randomly sampled questions.

Type	Count	Answer Categories (%)		
What	2,058	Factual:	100.00	
Where	1,896	Factual:	77.78	Abstract: 22.22
Who	1,847	Speaker:	30.56	Content: 69.44
Why	1,688	Explicit:	73.53	Implicit: 26.47
How	1,628	Explicit:	77.42	Implicit: 22.58
When	1,493	Absolute:	62.07	Relative: 37.93

Table 4: Statistics of the question types as well as the answer categories.

What No distinct categorization is found for answers to *what* questions, which are entirely factual. This is because annotators are mostly driven by factoid contents for the generation of *what* questions.

Where Answers to *where* questions can be categorized into factual and abstract, meaning that they are either concrete facts (e.g., named entities) or abstract concepts (e.g., *the wild, out there*), where the majority is driven by factoid contents (77.78%).

Who Answers to *who* questions can be annotated on either speaker names or utterance contents. The majority of *who* questions (69.44%) finds their answers in the utterance contents.

Why and How Answers to *why* and *how* questions are categorized into explicit and implicit such that they are either directly answering the questions (e.g., why doesn't Joey want to throw the chair out? → *Joey: I built this thing with my own hand*), or indirectly implying the answers (e.g., How are Joey and Chandler going to get to Monica's place? → *Joey: we're not gonna have to walk there, right?*). Explicit answers are more common for both *why* (73.53%) and *how* (77.42%) questions.

When Answers to *when* questions can be categorized into absolute and relative such that they can be either exact timing (e.g., clock time, specific date, holiday) or timing of action relative to another event (e.g., I called her *while I was watching TV*). About two third of the answers are considered explicit for *when* questions.

4 State-of-the-Art QA Systems

Three of state-of-the-art QA systems, R-Net based on recurrent neural networks (RNN) (Section 4.1), QANet based on convolutional neural networks (CNN) with self-attention (Section 4.2), and BERT based on deep feed-forward neural networks with transformers (Section 4.3), are used to validate our dataset as a practical resource for building advanced deep learning models. All models will output two positions which will be combined to form answer spans. These systems are chosen because they give a good survey among different types of neural networks in combination with attention mechanisms that are dominant in the research of contemporary question answering.

4.1 R-Net

R-Net held the 1st place on the SQuAD leaderboard at the time of its publication (Wang et al., 2017). It builds representations for questions and evidence passages using RNN and presents a self-matching mechanism to aggregate key information from the evidence passages, in order to compensate the limitedly memorized information from RNN. The same configuration described in the original paper is used to train models for our experiments.

4.2 QANet

QANet is another state-of-the-art open-domain QA system utilizing CNN and self-attention (Yu et al., 2018). Dramatic is the speed-up gained by QANet, which enables to perform data augmentation. Their original configuration cannot be fit in a 12GB GPU

machine using our dataset; thus, the configuration is compromised for our experiments as follows:

- The number of filters: 96 instead of 128,
- The number of attention heads: 1 instead of 8.

Given this configuration, its performance may not be optimal but at least can be directly compared to other models trained on the FriendsQA dataset.

4.3 BERT

The Bidirectional Encoder Representations from Transformers (BERT) pushed all current state-of-the-art scores to another level (Devlin et al., 2018). Trained with the masked language model on next sentence prediction tasks, BERT shows extremely promising results on several tasks in NLP. The pre-trained decapitalized BERT model with 12-layers is fine-tuned on our dataset. The larger BERT model with 24-layers again cannot be fit in a 12GB GPU machine; thus, it is not used for our experiments.

5 Experiments

For our experiments, all dialogues from Table 3 are randomly shuffled and redistributed as the training (80%), development (10%), and test (10%) as shown in Table 5.

Set	Dialogues	Questions	Answers
Training	977	8,535	17,074
Development	122	1,010	2,057
Test	123	1,065	2,131

Table 5: Data split for our experiments.

5.1 Model Development

Each instance consists of an evidence dialogue, a question and an answer span. Utterance IDs, annotated to indicate the whole utterances being answer spans (Section 3.2), are preprocessed and replaced by the actual spans on the dialogue contents. Since each question can have multiple answers, the following strategies are experimented to acquire one gold answer span for each training instance:

Shortest The shortest answer span is chosen and all the other spans are discarded from training.

Longest The longest answer span is chosen and all the other spans are discarded from training.

Multiple The question is paired with every answer to create multiple instances. For example, a question q with two answer spans, a_1 and a_2 , generate two instances, (q, a_1) and (q, a_2) , and trained independently.

Model	Shortest-Answer Strategy			Longest-Answer Strategy			Multiple-Answer Strategy		
	UM	SM	EM	UM	SM	EM	UM	SM	EM
R-Net	45.41 (±1.16)	35.69 (±1.28)	25.55 (±1.60)	49.50 (±0.54)	37.26 (±0.72)	23.77 (±0.42)	43.77 (±0.56)	33.97 (±0.75)	23.02 (±1.30)
QANet	42.12 (±3.21)	34.04 (±0.03)	22.89 (±0.42)	46.21 (±4.51)	34.55 (±1.87)	21.15 (±1.21)	47.10 (±1.30)	35.38 (±1.33)	23.16 (±1.15)
BERT	72.61 (±0.20)	63.64 (±0.42)	48.33 (±1.41)	72.16 (±1.93)	60.36 (±1.53)	43.23 (±1.83)	74.18 (±0.21)	64.15 (±0.29)	48.96 (±0.42)

Table 6: Results from the three state-of-the-art QA systems. All models are experimented three times and their average scores with standard deviations are reported. UM: Utterance Match, SM: Span Match, EM: Exact Match.

5.2 Evaluation Metrics

Two tasks are experimented, answer utterance selection and answer span selection, with the FriendsQA dataset. The utterance match (UM) is used to evaluate answer utterance selection, which checks if the predicted answer span a_i^p resides within the same utterance u_i^g as the gold answer span a_i^g , and is measured as follows: (n : # of questions):

$$\text{UM} = \frac{1}{n} \sum_{i=1}^n (1 \text{ if } a_i^p \in u_i^g; \text{ otherwise, } 0)$$

Following Rajpurkar et al. (2016), the span match (SM) is adapted to evaluate answer span selection, where each a_i^p is treated as a bag-of-tokens (ϕ) and compared to the bag-of-tokens of a_i^g ; the macro-average F1 score across all questions is measured for the final evaluation (P : precision, R : recall):

$$\text{SM} = \frac{1}{n} \sum_{i=1}^n \frac{2 \cdot P(\phi(a_i^p), \phi(a_i^g))R(\phi(a_i^p), \phi(a_i^g))}{P(\phi(a_i^p), \phi(a_i^g)) + R(\phi(a_i^p), \phi(a_i^g))}$$

Additionally, the exact match (EM) is used to evaluate answer span selection that checks the exact span match between the gold and predicted answers.

5.3 Results

Table 6 shows results from 9 models trained by the three state-of-the-art systems in Section 5.2 using the three answer selection strategies in Section 5.1. All experiments are run three times and their average scores with standard deviations are reported. BERT and QANet perform better with the multiple-answer strategy, that gives more training instances per question, whereas R-Net performs better with the other strategies. This could be due to R-Net’s self-matching mechanism that gets confused when multiple answers are provided for training the same question. BERT models significantly outperform ones from the other two systems in all evaluations. Since our hyper-parameters are tuned around grids provided by the original papers, it is possible that

these results are still suboptimal, which points out another important property of BERT that it is not as sensitive to different QA datasets.

Type	Dist.	UM	SM	EM
What	19.70%	77.43	69.39	55.04
Where	18.28%	84.35	78.86	65.93
Who	17.17%	74.12	64.34	55.29
Why	15.76%	60.47	50.03	27.14
How	14.65%	65.52	52.04	32.64
When	14.44%	80.65	65.81	51.98

Table 7: Results with respect to question types using BERT and the multiple-answer strategy.

Table 7 shows results from BERT’s multiple answer models by question types. Answers to *where* and *when* questions are mostly factoid, which show the highest performance. On the other hand, answers to *why* and *how* usually span out to longer sequences, leading to worse performance. Answers to *who* and *what* questions give a good mixture of proper and common nouns and show moderate performance.




Figure 1: Increasing score with top-20 answer candidates. From top to bottom: Utterance Match, Span Match and Exact Match.

Figure 1 shows improvement of BERT’s multiple-answer models by accepting the top- k answer predictions; the scores are measured by picking the best matching answer within the top- k predictions. UM surpasses 90% and SM approaches to 90%

when $k = 14$ and 20 , respectively. More importantly, the gap between UM and SM gets smaller as k increases, which implies that FriendsQA is not only learnable by deep learning but also can be enhanced by re-ranking the answer predictions.

5.4 Error Analysis

An extensive error analysis is manually performed on 100 randomly sampled, exact unmatched predictions ($F1 = 0$) to provide insights for future research. Figure 2 shows six types of errors that become evident through this analysis.

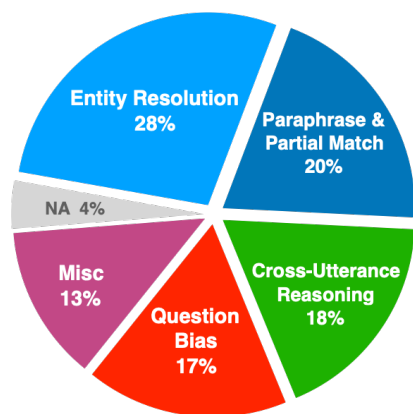


Figure 2: The distribution of six error types analyzed in 100 sampled predictions. NA: Noise in annotation.

Entity Resolution This type is the most frequent and often occurs when many of the same entities are mentioned in multiple utterances. The recurring use of coreference and anaphora can be confusing. This error also occurs when the QA system is asked about a specific person, but predicts wrong people. For example, the question asks for Chandler’s opinion about marriage, but the model matches comments from Joey instead due to the lack of referent resolution made in those comments.

Paraphrase and Partial Match This error type may be even challenging for humans without inside knowledge. Answers can be expressed in numerous ways through paraphrasing, abstraction, nicknames in dialogue, signifying the difficulty in FriendsQA. Moreover, answers might also be partially correct, especially for *why* and *how* questions, which could be acceptable in practice.

Cross-Utterance Reasoning This type reveals an universal challenge in understanding human-to-human conversation. To correctly predict an answer span in dialogue, the system should be equipped with the ability to reason across multiple utterances

back and forth, especially if a story or an event unfolds gradually, scatters in different places, and is told by different speakers.

Question Bias This type occurs when the answer predictions overly rely on the question types. For *why* questions, the model tends to blindly selects spans following certain keywords such as *because* even though they are placed in wrong utterances since the model is learned to be biased to the term *because*, neglecting other important factors that might otherwise lead to the correct answers.

Noise in Annotation (NA) Our dataset, although it gives high inter-annotator agreement (Sec. 3.7), it still includes noise caused by wrong spans, ambiguous or unanswerable questions, or typos.

Miscellaneous Errors in this category have no apparent cause to understand why the model predicts these answers, which often seem irrelevant to the questions so that they need more investigation.

Given this analysis, we hope many challenges can be overcome by future studies. For instance, coreferent mentions, especially plural mentions, should be more intelligently processed (Zhou and Choi, 2018). Moreover, the speaker information, which are currently treated as the first tokens in utterances, can be better encoded to give more insights.

6 Conclusion

This paper presents an open-domain question answering dataset called FriendsQA, compiled from transcripts of the TV show *Friends*. An extensive and comprehensive analysis is performed on this dataset to show its validity, difficulty and diversity. Three state-of-the-art models are run and compared, and show the full potential of FriendsQA as a rich QA research resource. Finally, erroneous answer predictions are sampled out for a further analysis to offer insightful retrospective. All our resources are publicly available.¹

For future work, the question-type (Table 7) and error analyses (Section 5.4) can serve as guidelines to further enhance the QA model performance. Top- k answer analysis also brings up another challenging but tangible task to re-rank the answer predictions. More tasks such as answer existence prediction and an utterance-based model to select among utterance candidates can also be issued.

¹FriendsQA: <https://github.com/emorynlp/question-answering>

References

- Ondrej Bajgar, Rudolf Kadlec, and Jan Kleindienst. 2016. Embracing data abundance: Book-Test Dataset for Reading Comprehension. *arXiv*, 1610.00956.
- Henry Yu-Hsin Chen and Jinho D. Choi. 2016. Character Identification on Multiparty Conversation: Identifying Mentions of Characters in TV Shows. In *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, SIG-DIAL'16, pages 90–100.
- Henry Yu-Hsin Chen, Ethan Zhou, and Jinho D. Choi. 2017. Robust Coreference Resolution and Entity Linking on Dialogues: Character Identification on TV Show Transcripts. In *Proceedings of the 21st Conference on Computational Natural Language Learning*, CoNLL'17.
- Yiming Cui, Zhipeng Chen, Si Wei, Shijin Wang, Ting Liu, and Guoping Hu. 2017. Attention-over-Attention Neural Networks for Reading Comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, ACL'17, pages 593–602.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv*, 1810.04805.
- David A. Ferrucci. 2011. IBM's Watson/DeepQA. In *Proceedings of the 38th Annual International Symposium on Computer Architecture*, ISCA'11.
- Pedro Gonçalves. 2017. 10 graphs that show why your business should be available through messaging apps.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching Machines to Read and Comprehend. In *Annual Conference on Neural Information Processing Systems*, NIPS'15, pages 1693–1701.
- Felix Hill, Antoine Bordes, Sumit Chopra, and Jason Weston. 2016. The Goldilocks Principle: Reading Children's Books with Explicit Memory Representations. In *Proceedings of the 6th International Conference on Learning Representations*, ICLR'16.
- Minghao Hu, Yuxing Peng, Zhen Huang, Xipeng Qiu, Furu Wei, and Ming Zhou. 2017. Reinforced Mnemonic Reader for Machine Reading Comprehension. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, IJCAI'18, pages 4099–4106.
- Hsin-Yuan Huang, Chenguang Zhu, Yelong Shen, and Weizhu Chen. 2018. FusionNet: Fusing via Fully-aware Attention with Application to Machine Comprehension. In *Proceedings of the International Conference on Learning Representations*, page ICLR'18.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. TriviaQA: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, ACL'17, pages 1601–1611.
- Aniruddha Kembhavi, Minjoon Seo, Dustin Schwenk, Jonghyun Choi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Are You Smarter Than a Sixth Grader? Textbook Question Answering for Multimodal Machine Comprehension. In *The IEEE Conference on Computer Vision and Pattern Recognition*, CVPR'17.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. RACE: Large-scale ReAding Comprehension Dataset From Examinations. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP'17, pages 785–794.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019. Multi-Task Deep Neural Networks for Natural Language Understanding. *arXiv*, 1901.11504.
- Xiaodong Liu, Yelong Shen, Kevin Duh, and Jianfeng Gao. 2018. Stochastic Answer Networks for Machine Reading Comprehension. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, ACL'18, pages 1694–1704.
- Kaixin Ma, Tomasz Jurczyk, and Jinho D. Choi. 2018. Challenging Reading Comprehension on Daily Conversation: Passage Completion on Multiparty Dialog. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2039–2048, New Orleans, Louisiana. Association for Computational Linguistics.
- Stephen Mayhew and Dan Roth. 2018. TALEN: Tool for Annotation of Low-resource ENTities. In *Proceedings of the ACL System Demonstrations*, ACL:DEMO'18, pages 80–86.
- Frank Newport. 2014. The New Era of Communication Among Americans.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. MS MARCO: A Human Generated Machine Reading Comprehension Dataset. In *Proceedings of the NIPS Workshop on Cognitive Computation: Integrating Neural and Symbolic Approaches*.
- Takeshi Onishi, Hai Wang, Mohit Bansal, Kevin Gimpel, and David McAllester. 2016. Who did What: A Large-Scale Person-Centered Cloze Dataset. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, EMNLP'16, pages 2230–2235.

- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep Contextualized Word Representations](#). In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, NAACL'18, pages 2227–2237.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ Questions for Machine Comprehension of Text](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP'16, pages 2383–2392.
- Siva Reddy, Danqi Chen, and Christopher D. Manning. 2018. [CoQA: A Conversational Question Answering Challenge](#). *arXiv*, 1808.07042.
- Matthew Richardson, Christopher J.C. Burges, and Erin Renshaw. 2013. [MCTest: A Challenge Dataset for the Open-Domain Machine Comprehension of Text](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, EMNLP'13, pages 193–203.
- Shimi Salant and Jonathan Berant. 2018. [Contextualized Word Representations for Reading Comprehension](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, ACL'18, pages 554–559.
- Yelong Shen, Po-Sen Huang, Jianfeng Gao, and Weizhu Chen. 2017. [ReasonNet: Learning to Stop Reading in Machine Comprehension](#). In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD'17, pages 1047–1055.
- Kai Sun, Dian Yu, Jianshu Chen, Dong Yu, Yejin Choi, and Claire Cardie. 2019. [DREAM: A Challenge Dataset and Models for Dialogue-Based Reading Comprehension](#). *Transactions of the Association for Computational Linguistics*, 7.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2017. [NewsQA: A Machine Comprehension Dataset](#). In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 191–200.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention Is All You Need](#). In *Proceedings of the 31st Conference on Neural Information Processing Systems*, NIPS'17, pages 5998–6008.
- Wenhui Wang, Nan Yang, Furu Wei, Baobao Chang, and Ming Zhou. 2017. [Gated Self-Matching Networks for Reading Comprehension and Question Answering](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, ACL'17, pages 189–198.
- Johannes Welbl, Nelson F. Liu, and Matt Gardner. 2017. [Crowdsourcing Multiple Choice Science Questions](#). In *Proceedings of the 3rd Workshop on Noisy User-generated Text*, pages 94–106.
- Jason Weston, Antoine Bordes, Sumit Chopra, Alexander M. Rush, Bart van Merriënboer, Armand Joulin, and Tomas Mikolov. 2016. [Towards AI-Complete Question Answering: A Set of Prerequisite Toy Tasks](#). In *Proceedings of the 5th International Conference on Learning Representations*, ICLR'16.
- Ikuya Yamada, Ryuji Tamaki, Hiroyuki Shindo, and Yoshiyasu Takefuji. 2017. [Studio Ousia's Quiz Bowl Question Answering System at NIPS HCQA 2017](#). In *Human-Computer Question Answering Competition at the 31 Annual Conference on Neural Information Processing Systems*, NIPS-HCQA'17.
- Yi Yang, Wen-tau Yih, and Christopher Meek. 2015. [WIKIQA: A Challenge Dataset for Open-Domain Question Answering](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP'15, pages 2013–2018.
- Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V. Le. 2018. [QANet: Combining Local Convolution with Global Self-Attention for Reading Comprehension](#). In *Proceedings of the 6th International Conference on Learning Representations*, ICLR'18.
- Sayed Zahiri and Jinho D. Choi. 2018. [Emotion Detection on TV Show Transcripts with Sequence-based Convolutional Neural Networks](#). In *Proceedings of the AAAI Workshop on Affective Content Analysis*, AFFCON'18, New Orleans, LA.
- Ethan Zhou and Jinho D. Choi. 2018. [They Exist! Introducing Plural Mentions to Coreference Resolution and Entity Linking](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 24–34.

Foundations of Collaborative Task-Oriented Dialogue: What’s in a Slot?¹

Philip R. Cohen
Laboratory for Dialogue Research
Faculty of Information Technology
Monash University

Abstract

In this paper, we examine the foundations of task-oriented dialogues, in which systems are requested to perform tasks for humans. We argue that the way this dialogue task has been framed has limited its applicability to processing simple requests with atomic “slot-fillers”. However, such dialogues can contain more complex utterances. Furthermore, situations for which it would be desirable to build task-oriented dialogue systems, e.g., to engage in collaborative or multiparty dialogues, will require a more general approach. In order to provide such an approach, we give a logical analysis of the “intent+slot” dialogue setting that overcomes these limitations.

1 Introduction

An important problem that forms the core for many current spoken dialogue systems is that of “slot-filling” — the system’s ability to acquire required and optional attribute-values of the user’s requested action, for example, finding the date, time, and number of people for booking a restaurant reservation, or the departure date, departure time, destination, airline, arrival date, arrival time, etc. for booking a flight (Bobrow et al., 1977, Zue et al., 1991). If a required argument is missing, the system asks the user to supply it. Although this may sound simple, building such systems is more complex than one might suppose. For example, real task-related dialogues may be constraint-based rather than slot-filling, and are usually collaborative, such that dialogue participants may together fill slots,

and people go beyond what was literally requested to address higher-level goals.

In this paper, we discuss the limitations of the general slot-filling approach, and provide a formal theory that can be used not only to build slot-filling task-oriented dialogue systems, but also other types of dialogues, especially multiparty and collaborative ones. We argue first that without being explicit about the mental states and the logical forms that serve as their contents, systems are too tightly bound to the specific and limited conversational task of a single user’s getting a system to perform an action.

1.1 Intent+Slots (I+S)

The spoken language community has been working diligently to enable users to ask systems to perform actions. This requires the system to recover the user’s “intent” from the spoken language, meaning the action the system is being requested to perform, and the arguments needed to perform it, termed “slots”.² The most explicit definition of “slot” we can find is from (Henderson, 2015) in describing the Dialog State Tracking Challenge (DSTC2/3):

The slots and possible slot values of a slot-based dialog system specify its domain, i.e. the scope of what it can talk about and the tasks that it can help the user complete. The slots inform the set of possible actions the system can take, the possible semantics of the user utterances, and the possible dialog states... For each slot $s \in S$, the set of possible values for the slot is denoted V_s .

Henderson goes on to describe a system’s dialog state and two potentially overlapping slot

¹ Inspired by Woods (1975), “What’s in a Link: Foundations for Semantic Networks”

² See <https://developer.amazon.com/docs/custom-skills/create-intents-utterances-and-slots.html> for an example of the commercial interest in “intent + slots”.

types, so-called “informable” and “requestable” slots, denoted by sets S_{inf} and S_{req} , respectively.

The term *dialog state* loosely denotes a full representation of what the user wants at any point from the dialog system. The dialog state comprises all that is used when the system makes its decision about what to say next. ... the dialog state at a given turn consists of:

- The **goal constraint** for every informable slot $s \in S_{inf}$. This is an assignment of a value $v \in V_s$ that the user is specifying as a constraint, or a special value *Dontcare*, which means the user has no preference, or *None*, which means the user is yet to specify a valid goal for the slot.
- A set of **requested slots**, the current list of slots that the user has asked the system to inform. This is a subset of S_{req} .^{3,4} (Henderson, 2015) ...

Most papers in the field at best have informal definitions of “intent” and “slot”. In order to clarify these concepts, we frame their definitions in a logic with a precise semantics. We find the following topics require further explication.

2 Limitations of Slot-Filling

2.1 Representation of Actions

The DSTC proposes a knowledge representation of actions with a fixed set of slots, and atomic values with which to fill them, such as `reserve(restaurant=Mykonos, cuisine=Greek, Location = North)` to represent the user’s desire that the system reserve Mykonos, a Greek restaurant in the north of town, or `reserve(restaurant=none, cuisine=Greek, Location = dontcare)`, which apparently says that the user wants the system to reserve a Greek restaurant anywhere. However, missing from this representation is the agent of the action. At a minimum, we need to be able to distinguish between the user’s performing and the system’s performing an action. Thus, such a representation cannot directly accommodate the user’s saying “I want to eat at Guillaume” because the user is not explicitly requesting the system to perform an action.⁵ Also missing are variables used as values, especially *shared* variables. This severely limits the kinds of utterances people can provide. For example, it would prevent the

system from representing the meaning of “I want you to reserve that Greek restaurant in the north of Cambridge *that John ate at last week.*”

2.2 Restrictions on Logical Forms (LFs)

Next, the slot-filling approach limits the set of logical forms the dialogue system can consider by requiring the user to supply an atomic value (including *Dontcare* and *None*) to fill a slot. For example, slot-filling systems can be trained to expect simple atomic responses like “7pm” to such questions as “*what time do you want me to reserve a table?*” However, I+S systems typically will not accept such reasonable responses as “*not before 7pm,*” “*between 7 and 8 pm,*” or “*the earliest time available.*” What’s missing from these systems are true logical forms that employ a variety of relations and operators, such as **and**, **or**, **not**, **all**, **if-then-else**, **some**, **every**, **before**, **after**, **count**, superlatives, comparatives, as well as proper variables. Critically, adequate meaning representations are compositional often employing relative clauses, such as the LF underlying “*What are the three best Chinese or Japanese restaurants that are within walking distance of Century Link Field?*” Compositional utterances often require scoped representations, as in “*What is the closest parking to the Japanese restaurant nearest to the Space Needle?*” which has two superlative expressions, one embedded within the other. These phenomena are also problematic for requests, as in: *Book a table at the closest good Italian restaurant to the Orpheum Theater on Monday for 4 people.* Although current I+S systems cannot parse or represent such utterances (Ultes et al. 2018), complex logical forms such as those underlying the above can now be produced robustly from competent semantic parsers (e.g., (Duong et al., 2017; Wang et al., 2015)). What we claim is necessary is to move from an I+S representation language of actions with attributes and atomic values to a true logical form language with which to represent the meaning of users’ utterances.

2.3 Explicit Attitudes

However, this is still not sufficient. The I+S approach, as incorporated into the DSTC 2 (Henderson, 2015), says that the dialogue state

³ This appears to be the reverse of the definition in (Gašić et al., 2016, p. 557)

⁴ At least implicitly, the DSTC must allow a distinguished symbol (e.g., ‘?’) to indicate what slot values are being requested. Alternatively, we have seen `request(<attribute>`

with an unstated value, meaning the user is asking for the value of the attribute.

⁵ In order to handle this as an indirect request, a system would need to reason about users’ plans and how the system can help the user achieve them.

“loosely denotes a full representation of what the user wants at any point from the dialog system”, but treats as implicit the desire attitude associated with the intent content. Thus, when a user says “I want you to reserve for Monday” the notion of “want” is taken to be just syntactic sugar and is generally thrown away, resulting in a representation that looks like this: `inform(reserve(day = monday))`. But this is too simplistic for a real system as there are many types of utterances about actions that a user might provide that cannot be so expressed. For example, the user might want to personalize the system by telling it never to book a particular restaurant, i.e., the user wants the system *not* to perform an action. Moreover, a virtual assistant positioned in a living room may be expected to help multiple people, either as individuals or as a group. A system needs to keep separate the actions and parameters characterizing one person’s desires from another’s, or else it will be unable to follow a discussion between two parties about an action. For example, John says *he wants* the system to reserve Vittorio’s for he and Sue on Monday, and Sue says *she wants* the reservation on Tuesday. In addition to specifying agents for actions, we need to specify the agent of the inform, so that we can separate what John and Sue each said, as in: `inform(agent=john, reserve(patron=[john,sue],day=monday))`, and `inform(agent=sue, reserve(patron=[john,sue], day=tuesday))`. But, since I+S slots encode the *speaker’s* desire, how can *John’s* saying “Sue wants you to reserve Monday” be represented? Does this utterance fill slots in Sue’s desired reservation action, both of theirs, or neither? And what if Sue replies “no, I don’t”? What then is in the `day` slot for Sue? Dontcare? She didn’t say she doesn’t care what day a table is reserved. In fact, she *does* care — she does *not want* a reservation on Monday. By merely having an implicit attitude, we cannot represent this.⁶

All these representational weaknesses compound. Imagine John’s being asked by the system “when do you want me to reserve Vittorio’s?” and he replies “whenever Sue wants.” Again, whose slot and attitude is associated with the utterance— John’s or Sue’s?

⁶Some researchers have advocated a “negate($a=x$)” action with an informal semantics that the user does not want the slot a to be filled with the value x (Young et al., 2010). In the multiparty case, one would need to be more explicit about whose slot and desire this is.

Without a shared variable, agents for actions, and explicit desires, we cannot represent this either.

2.4 Mixed initiative and collaboration

Finally, in the dialogue below, apart from the fact that I+S cannot represent utterance (1), question (2) is answered with a subdialogue starting at question (3) that shifts the dialogue initiative (Bohus and Rudnicky, 2002; Horvitz, 2007; Litman and Allen, 1987; Morbini et al., 2012). In utterances (4) and (6), the system is proposing a value and in (5) and (7), the user is rejecting or accepting the proposal. Thus, *both* system and user are *collaboratively* filling the slot (Clark and Wilkes-Gibbs, 1986), not just one or the other. I+S systems cannot do this.

- (1) U: Please book a reservation at the closest good restaurant to the Orpheum Theater on Monday for 4 people.
- (2) S: OK, I recommend Guillaume. What time would you like to eat?
- (3) U: what’s the earliest time available?
- (4) S: 6 pm
- (5) U: too early
- (6) S: how about 7 pm?
- (7) U: OK

2.5 Dialogue state and belief

The DSTC approach to I+S represents dialogue state in terms of the user’s desires. We claim that task-oriented dialogue systems, especially those that could engage in multiparty conversations, will also need to explicitly represent other mental states, including but not limited to people’s beliefs.⁷ The naive approach to representing beliefs is as an embedded database (Cohen, 1978; Moore, 1977). Such an approach could perhaps work until one attempts to deal with vague beliefs. For example, you know Joe is sitting by a window and able to look outside. You can reasonably ask Joe “Is it raining?” because you believe that *either* Joe believes it is raining, *or* Joe believes it is not raining, i.e., Joe knows whether it is raining or not. This is different than believing that Joe believes that $\text{Rain} \vee \sim\text{Rain}$, which is a tautology. But to use the database approach, what should the system put into Joe’s database? It can’t put in Rain , and it can’t put in $\sim\text{Rain}$, or else it would not need to ask. It needs to represent something

⁷ This is a different notion of “belief” than “belief state” as used in POMDP dialogue modeling (Williams & Young, 2007).

more vague – *that* Joe knows if it is raining, a concept that was described as $\text{KNOWIF} \stackrel{\text{def}}{=} (\text{BEL } x \text{ P}) \vee (\text{BEL } x \sim \text{P})$ (Allen 1979; Cohen and Levesque, 1990b; Cohen and Perrault, 1979; Miller et al., 2017; Perrault and Allen, 1980; Sadek et al., 1997, Steedman and Petrick, 2015). In the case of a multiparty dialogue system, the system should direct the yes/no question of whether it is raining to the person whom it believes knows the answer without having to know what they think it is.

2.6 Knowledge acquisition

Any task-oriented dialogue system will need to acquire information, usually by asking wh-questions, which we have argued will require it to deal somehow with variables. Again, for a multiparty context, in order to ask a wh-question, the system should be asking someone whom it thinks knows the answer. We need to be able to represent such facts as “John knows Mary’s mobile phone number”, which is different from saying “John knows Mary has a mobile phone number”. In the former case, I could ask John the question “what is Mary’s phone number?”, while in the latter case, it would be uncertain whether he could reply. This ability to represent an agent’s knowing the referent of a description, was called KNOWREF (Allen 1979; Cohen and Levesque, 1990b; Cohen and Perrault, 1979; Perrault and Allen, 1980), Bref (Sadek et al., 1997), or KNOWS_VAL (Young et al., 2010), and is intimately related to the concept of quantifying-into a modal operator (Barcan, 1946; Kaplan, 1968; Kripke, 1967; Quine, 1956), about which a huge amount of philosophical ink has been spilled. For a database approach to representing belief, the problem here revolves around what to put in the database to represent Mary’s phone number. One cannot put in a constant, or one is asserting that to be her phone number. And one cannot put in an ordinary variable, since that provides no more information than the existentially quantified proposition that she has a phone number, not that John knows what it is! Over the years, various researchers have attempted to incorporate special types of constants (Cohen, 1978; Konolige, 1987), but to no avail because the logic of these constants requires that they encode all the modal operators in whose scope they are quantified. Rather, one needs to represent and reason with quantified beliefs like

$$\exists X (\text{BEL john phone_number(mary,X)})$$

To preview our logic below, we define some syntactic sugar using roles and Prolog syntax (and a higher-order schematic variable ranging over predicates Pred):

$$(\text{KNOWREF agent:X variable:Var predicate:Pred}) \stackrel{\text{def}}{=} \exists \text{Var} (\text{BEL } x \text{ Pred}), \text{ with Var bound in Pred}$$

In other words, the agent X knows the referent of the description ‘ Var such that Pred ’. For example, we can represent “John knows Mary’s phone number” as

$$(\text{KNOWREF agent:john,variable:Ph, predicate:phone_number(mary,Ph)})$$

In summary, a system’s beliefs about other agents cannot simply be a database. Rather, the system needs to be able to represent such beliefs without having precise information about what those beliefs are.⁸ If it can do so, it can separate what it takes to be one agent’s beliefs from another’s, which would be needed for a multiparty dialogue system. Dialogue state for task-oriented dialogue systems is thus considerably more complex than envisioned by I+S approaches.

3 Logic of Task-Oriented Conversation

Let us now cast the I+S dialogue setting into a logical framework. We will examine intent vs. intention, semantics of slots, and dialogue state.

3.1 What is an Intent?

How does the action description in such utterances as those above relate to an “intent”? First, let us assume “intent” bears some relation to “intention”. What appears to be the use within the spoken language community is that an “intent” is the action content of a user request that (somehow) encodes the user’s intention. To be precise here, we need to review some earlier work that can form the basis for a logic of task-oriented conversation.

3.2 The Language \mathcal{L}

We will use Cohen and Levesque’s (1990) formal language and model theory for expressing the relations among belief, goal, and intention (see Appendix for precise description of \mathcal{L}). Other formal languages that handle belief and intention (e.g., (Rao and Georgeff, 1995)) may do just as

⁸ Note that this has nothing to do with uncertainty in the probabilistic sense. I can be certain that John knows Mary’s phone number, but still not know what it is.

well, but this will provide the expressivity we need. The language \mathcal{L} is a first-order multi-modal logical language with basic predicates, arguments, constants, functions, objects, quantifiers, variables, roles, values (atomic or variables), actions, lists, temporal operators (Eventually (\diamond , LATER), DOES and DONE), and two mental states, BEL and GOAL. The logic does not consider agents' preferences, assuming the agent has chosen those it finds superior (according to some metric such as expected utility). These are called GOALS in the logic. Unlike preferences, at any given time, goals are consistent, but they can change in the next instant. As is common, we refer to this as a BDI logic. See the Appendix for examples of well-formed formulas.

3.3 Possible worlds semantics

Again from (Cohen and Levesque, 1990), the propositional attitudes BEL and GOAL are given a relatively standard possible worlds semantics, with two accessibility relations B and G . However, for modelling slot-filling, we are critically interested in the semantics of "quantifying-in" (Barcan, 1946; Kaplan, 1968; Kripke, 1967; Quine, 1956). Briefly, a variable valuation function v in the semantics assigns some value chosen from the domain of the world and time at which the formula is being satisfied. When "quantifying-into" a BEL or GOAL formula, that value is chosen and then the BEL or GOAL formula is satisfied. As is standard in modal logic after (Kripke, 1967), the semantics of these modal operators is given in terms of a universal quantifier ranging over B - and G -related possible worlds. Thus, the semantics of satisfying $\exists y(\text{BEL } x \text{ } p(y))$ in world W is that there is a single value that is assigned by the variable assignment function v to y , such that for all worlds W' that are B -related to W , $p(y)$ is true in W' . In other words, the value assigned to y is the same for all the related worlds W' . If the quantifier is within the scope of the modal operator as in $(\text{BEL } x \exists y p(y))$, then a different value could be assigned to the variable in each B -related world. Likewise, one can quantify into GOAL, and even iterated modalities or modalities of different agents. This gives rise to the theorems below, and analogous ones for GOAL.

$$\models \exists y (\text{BEL } x \text{ } p(y)) \supset \not\vdash (\text{BEL } x \exists y p(y)), \text{ and}$$

$$\models \text{BEL } x \text{ } p(c) \supset \not\vdash \exists y (\text{BEL } x \text{ } p(y)) \text{ for constant } c.$$

This paper shows why quantifying into BEL and GOAL is key for slot-filling systems.

3.4 Persistent goals and intentions

Cohen and Levesque (1990) defined a concept of an internal commitment, namely an agent's adopting a relativized persistent goal (PGOAL x P Q), to be an achievement goal P that x believes to false but desires to be true in the future, and agent x will not give up P as an achievement goal at least until it believes P to be satisfied, impossible, or irrelevant (i.e., x believes $\sim Q$). If the agent believes $\sim Q$, it can drop the PGOAL. More formally, they have:

$$\begin{aligned} (\text{PGOAL } x \text{ } P \text{ } Q) =_{\text{def}} & (\text{GOAL } x (\text{LATER } P)) \wedge (\text{BEL } x \sim P) \wedge \\ & (\text{BEFORE } ((\text{BEL } x \text{ } P) \vee (\text{BEL } x \square \sim P) \vee (\text{BEL } x \sim Q)) \\ & \quad \sim (\text{GOAL } x (\text{LATER } P)) \end{aligned}$$

They also defined an *intention* to be a persistent goal to perform an action. More formally:

$$(\text{INTEND } x \text{ } A \text{ } Q) =_{\text{def}} (\text{PGOAL } x (\text{DONE } x \text{ } A) \text{ } Q).$$

In other words, an agent x intending to do an action A is internally committed (i.e., has a PGOAL) to having performed the action A in the future. So, an intention is a future-directed commitment towards an action.

3.5 What is a slot?

Given this language, how would one represent a DSTC slot, which incorporates the user's desire? We propose to separate the attitude, action, and role-value list, then reassemble them. First, we consider the role:value argument in an action expression, using upper case variables (as in Prolog), such as `reserve(patron:P, restaurant:R, day:D, time:T, num_eaters:N)`. Here, `restaurant:R` is the role:value expression. Next, we need to add the desire attitude (as a PGOAL) in order to express such phrases "the day Joe *wants* me to reserve Vittorio's Ristorante for him." Here is how we would express it as part of the system's belief:

$$\begin{aligned} (1) \exists \text{Day} & \\ & (\text{PGOAL } \text{joe } \exists [T, N] \\ & \quad (\text{DONE } \text{sys } \text{reserve}([\text{patron}:\text{joe}, \\ & \quad \quad \quad \text{restaurant:vittorios}, \\ & \quad \quad \quad \text{day:Day}, \text{time:T}, \\ & \quad \quad \quad \text{num_eaters:N]))) \text{ } Q) \end{aligned}$$

In other words, there is a Day on which Joe is committed to there being a Time, and number of eaters N such that the system reserves Vittorio's

on that **Day** at that **Time** and with **N** eaters. The system has represented Joe as being picky about what day he wants the system to reserve Vittorio's (e.g., as a creature of habit, he always wants to eat there on Monday), but the system does not know what day that is. Here, we have quantified **Day** into the PGOAL, but the rest of the variables are existentially quantified within the PGOAL. That means that Joe has made no choice about the **Time** or **Number** of people. But because the system has this representation, it can reasonably ask Joe "What day do you want me to reserve Vittorio's?". We can now also represent the day Joe does *not* want the system to reserve, can distinguish between the day Joe wants the system to reserve and the day Sue wants, and we can even equate the two, saying that Joe wants the system to reserve on whatever day Sue wants (See section 2.7). So the DSTC "slot" **day** turns out to have a variable in an action expression all right, but one that is now quantified into an intention or PGOAL operator. This explicit representation enables the system to discuss the action with or without anyone's wanting to perform it, and to differentiate between agents' attitudes, which is essential for multiparty dialogues.

3.6 Where do the slot-filling goals and intentions come from?

In order to know what action to perform, an agent needs to know the values of the required arguments of an action. (Allen and Perrault, 1980; Appelt, 1985; Cohen and Perrault, 1979; Moore, 1977)⁹. In the case of the task-oriented dialogue setting, in which the agents are intended to be cooperative, we will have all agents obey the following rule. (We suppress roles below and hereafter.)

For any agents X and Y (who could be the same):

If: (BEL Y (PGOAL X (DONE Y A) Q)),

Then for the set of required but unfilled obligatory arguments **Args**, assert

(2) (PGOAL Y
 (KNOWREF Y Args (PGOAL X (DONE Y A)),
 (PGOAL X (DONE Y A) Q)),

⁹ Required arguments will be stipulated as part of a meta-data template in the system's knowledge base. Knowing the values for arguments of actions is not the only case in which having to know an argument is required. For

In other words, assuming Y is the system and X is the user, this rule says that if the system believes the user is committed to the system's doing an action A (as would be the result of a request), then the system is committed to knowing the referents of all required arguments of the action A that the user wants the system to perform.¹⁰ That is, the system is committed to knowing the *user's* desired "slot" values in the action that the user wants the system to perform. For example, if the system believes the user wants the system to do the action of reserving Vittorio's Ristorante for the user, then the system adopts a persistent goal to know the **Time**, **Day**, and **Num**, for which the user wants the system to reserve Vittorio's.¹¹

Notice that this holds no matter how the system comes to infer that the user wants it to do an action. For example, the system could make an indirect offer and the user could accept (Smith and Cohen, 1996), as in *System*: "Would you like me to reserve vittorio's for you?" *User*: "Sure". Here, the offer is stated as a question about what the user wants the system to do, and the positive reply provides the system with the rule antecedent above.

3.7 Application of the logic to I+S: Expressing problematic user responses

Let us now apply the logic to handle some of the expressions we claimed were problematic for an I+S approach. Assume the system has asked the user: "What time do you want me to reserve Vittorio's Ristorante?" We start with the base case, i.e. with the user's supplying an atomic value, and assume the representation of the question has only the **Time** variable quantified-in.

User: "7 pm".

Essentially, we unify the variable quantified into the PGOAL with the atom 7pm, resulting in:

(PGOAL usr \exists [Day,N]
 (DONE sys reserve([usr, vittorios,Day,7pm, N]))
 Q)

This is classic slot-filling.

User: "I don't know". The system would need to assert into its database a formula like the following (assume the action variable A

example, for the system to determine the number of available seats at a restaurant, it needs to know the date.

¹⁰ When X and Y are the same agent, (PGOAL X (DONE X A)) is exactly the definition of an intention.

¹¹ Formula (1) is a consequence of this.

represents the act of reserving Vittorio’s for the user, and that it has a free variable Time):

\sim (KNOWREF usr Time
(PGOAL usr (DONE usr, A) Q))

In doing so, the system should retract its previous KNOWREF belief that enabled it to ask the original question. How a system responds to this statement of ignorance is a different matter. For example, it might then ask someone else if it came to believe that person knows the answer. Thus, if the user then said “but Mom knows” and the system believes the user, the system could then ask Mom the question.

User: “I don’t care”. There are only two approaches we have seen to handling this in the I+S literature. One is to put the Dontcare atom into the value of a slot (Henderson, 2015). However, it is not clear what this means. It does not mean the same thing as “I don’t know.” It might be the equivalent of a variable, as it matches anything as a slot value, but that begs the question of variables in slots. To express “I don’t care” in the logic, we can define CAREREF, a similar concept to KNOWREF:

(CAREREF x Var Pred) =_{def} \exists Var (GOAL x Pred), where Var is free in Pred. Then for “I don’t care”, one could say: \sim (CAREREF x Var Pred) with the formal semantics that there is no specific value v for Var towards which x has a goal that Pred be true of it.

Rather than have a distinguished “don’t care” value in a slot, Bapna et al. (2017) create a “don’t_care(slot)” intent, with the informal meaning that the user does not care about what value fills that slot.¹² Here, it is not clear if this applies on a slot-by-slot basis, or on an intent+slot basis. For example, if it is on a slot-by-slot basis, then if the user says “I don’t care” to the question “Do you want me to reserve Monday at 7pm or Tuesday at 6pm?” it would lead to four don’t_care(slot) intent expressions. Would these be disjunctions? How would the relation between Monday and 7pm be expressed?

By contrast, we can define a comparable concept to KNOWIF,

(CAREIF x P) =_{def} (GOAL x P) \vee (GOAL x \sim P)

such that one can say “x doesn’t care whether P”, as \sim (CAREIF x P), with the obvious logical interpretation. With CAREIF, one could express

the reply “I don’t care” to the above disjunctive question as:

\sim (CAREIF usr
(LATER
(DONE sys reserve([usr, mond, 7pm])) \vee
(DONE sys reserve([usr, tues, 6pm]))))

User: “before 8 pm.” Because all that the I+S approach can do is to put atomic values in slots or leave them unfilled, the only approach possible here is to put some atom like before_8_pm into the slot. If one tried to give a semantics for this, it might be a function call or λ -expression that would somehow be interpreted as a comparative relation with whatever value eventually fills the slot. But, one would need a different comparison relation for every time value, not to mention for other more complex expressions such as not_before_7_pm_or_after_9_pm, or between_7_pm_and_9_pm. How would the system infer that these are the same condition? Instead, one might think we only need a method to append new constraints to the quantified persistent goal “slot” expression, as in

\exists Time (PGOAL usr
 \exists [Day,Num]
(DONE sys
reserve([usr,vittorios,Day,Time,Num]))
 \wedge (BEFORE Time 8:15_pm))

However, as a representation of the reply, the above is not quite what we want. Here, the user has implicated (Grice, 1975) that she does not have a goal for a particular time such that she wants a reservation at that time. Rather, she wants whatever time she eats to be before 8:15 pm. So, in fact, we want this constraint to be embedded within the scope of the existential quantifier:

(PGOAL usr \exists [Day,Time,Num]
((DONE sys reserve([usr,vittorios,
Day,Time, Num]))
 \wedge (BEFORE Time 8:15_pm)))

The reason we need an inference like a Gricean implicature is that the system would need to reason that in response to the question, if the user knew the answer, she would have told me, and she didn’t, so she (probably) doesn’t know the answer. Thus, the system needs to assert a weaker PGOAL.

¹² Notice that “intent” for Bapna et al. does not indicate an action being requested, so their notion of intent is different

from that of (Henderson, 2015) or that used by Amazon Alexa.

User: “*whenever Mary wants.*” To represent the content of this utterance, one can equate the quantified-in variables T_1, T_2 (and ignoring Q):

$$\begin{aligned} & \exists[T_1, T_2] (\text{equals } T_1, T_2) \wedge \\ & ((\text{PGOAL } \text{usr } \exists[\text{Day}, \text{Num}] \\ & (\text{DONE } \text{sys } \text{reserve}([\text{usr}, \text{vittorios}, \text{Day}, T_1, \text{Num}]))) \wedge \\ & (\text{PGOAL } \text{mary } \exists[\text{Day}, \text{Num}] \\ & (\text{DONE } \text{sys } \text{reserve}([\text{mary}, \text{vittorios}, \text{Day}, T_2, \text{Num}]))) \end{aligned}$$

If the system learns that Mary wants the reservation to be at 7 pm, it can infer that the User wants it then too.

The above examples show that the logic can represent users’ utterances in response to slot-filling questions that supply constraints on slot values, but not the values themselves.

4 Towards Best Practices

This paper has provided a logical definition of the DSTC 2/3 slot (and I+S slots more generally) as a quantified-in formula stating the value that the agent wants an action’s role to have. In addition, the logic presented here captures a more general concept than what I+S supports, in that it can express multiple agents’ desires as well as non-atomic constraints on attribute-value in logical forms.

Still, our purpose here is not merely clarity and good hygiene, but ultimately to build systems that can engage in explainable, collaborative, multiparty dialogues. Below we sketch how to build systems that can handle the above issues, some of which we have implemented in a prototype system that uses the logic in this paper to engage in collaborative knowledge-based dialogues, including slot-filling. A report on this system and approach will be provided in a subsequent paper.

4.1 Enabling an operational semantics

Systems based on a BDI logic will often have a belief-desire-intention architecture that serves as an operational semantics for the logic (Rao and Georgeff, 1995). By “operational semantics”, we mean that the system’s operation behaves (or at least approximates) the requirements of the logic. For example, the adoption of a persistent goal to achieve a state of affairs results in finding a plan to achieve it, which then results in the agent’s intending to perform the planned action. If the system finds a persistent goal/intention to be achieved, impossible or irrelevant, it drops that mental state, which causes an unraveling of other mental states as well. Our system in fact reasons

with the formulas shown here, engaging in slot-filling and related question-answering dialogues. However, other systems may be able to make such distinctions without explicit logical reasoning.

4.2 A plan-based approach to dialogue

We advocate a plan-based model of dialogue (Allen, 1979, Allen and Perrault, 1980; Allen et al., 1995; Appelt, 1985; Cohen 1978; Cohen and Perrault, 1979; Cohen and Levesque, 1990b; Galescu et al., 2017; Litman and Allen 1987; Perrault and Allen, 1980; Sadek et al., 1997; Steedman and Petrick, 2007; Stone, 2004; Traum and Hinkelman, 1992) such that the same planning and plan recognition algorithms can apply to both physical, digital, and communicative acts. When applied to communicative acts, the system plans to alter its own and the users’ beliefs, goals, and intentions. For example, goal (2) as applied to the slot expression in (1) will cause it to plan the wh-question “*what day would you like me to reserve Vittorio’s?*” to alter the speaker’s KNOWREF in goal (2) (see Appendix for definition of whq). Conversely, as a collaborator, on identifying a user’s speech act, the system asserts the user’s goal was to achieve the effect of the speech act. Based on that effect, the system attempts to recognize the user’s larger plan, to debug that plan, and to plan to overcome obstacles to it so that the user may achieve his/her higher level goals (Allen, 1979; Cohen, 1978; Cohen et al., 1982). In this way, a system can engage in collaborative non-I+S dialogues such as User: “Where is Dunkirk playing?” System: “It’s playing at the Roxy theater at 7:30pm, however it is sold out. But you can watch it on Netflix.” Finally, the system is in principle explainable because everything it says has a plan behind it.

4.3 A hybrid approach to handling task-oriented dialogue variability.

In order to incorporate such an approach into a useful dialogue system, we advocate building a semantic parser using the crowd-sourced “overnight” approach (Duong et al., 2018; Wang et al., 2015), which maps crowd-paraphrased utterances onto LFs derived from a backend API or data/knowledge base. This methodology involves: 1) Creating a grammar of LFs whose predicates are chosen from the backend application/data base, 2) using that grammar to generate a large number of LFs, 3) generating a “clunky” paraphrase of an LF, and 4) collecting

enough crowd-sourced natural paraphrases of those clunky paraphrases/LFs¹³. A neural network semantic parser trained over such a corpus can handle considerable utterance variability, including the creation of logical forms both for I+S utterances, and for complex utterances not supportable by I+S approaches. In the past, we have used this method to generate a corpus of utterances and logical forms that supported the semantic parsing/understanding of the complex utterances in Section 2.2 (Duong et al., 2017; Duong et al., 2018).

Whereas much utterance variability and uncertainty can be captured via the above approach, we believe there is less variability at the level of the goal/intention lifecycle, which includes goal adoption, commitment, planning, achievement, failure, abandonment, reformulation, etc. (Galescu et al., 2018; Johnson et al., 2018). This goal lifecycle would be directly supported by the BDI architecture and therefore would be available for every domain. Rather than train a dialogue system end-to-end where we would need many examples of each of these goal relationships, we believe a domain independent dialogue manager can be written once, parameterized by the contents of the knowledge representation (Allen et al., 2019; Galescu et al., 2018). Beyond learning to map utterances to logical forms, the system needs to learn how to map utterances in context to goal relationships. For example, what does “too early” in Utterance (5) of Section 2.4 mean? Is that a rejection of a contextually-specified proposal? The system also needs to learn how actions in the domain may lead to goals for which the user may want the system’s assistance. In order to be helpful to the user, the system must recognize the user’s goals and plan that led to his/her utterance(s) (Allen and Perrault, 1980; Sukthankar et al., 2014; Vered et al., 2016). One approach is to collect the action data needed to support plan recognition via crowdsourcing and text mining (Branavan et al., 2012; Fast et al., 2016; Jiang and Riloff, 2018). The upshot will be a collaborative dialogue manager that can be used directly in a dialogue system, or can become a next generation user simulator with which to train a dialogue manager (Schatzman et al., 2007; Shah et al., 2018).

Acknowledgments

This paper benefitted from insightful comments by the reviewers, Drs. Mark Johnson, Lizhen Qu, and Mor Vered.

5 References

- Allen, J. F. A plan-based approach to speech act recognition, PhD Thesis, Dept. of Computer Science, University of Toronto, 1979.
- Allen, J. F. and Perrault, C. R., Analyzing intention in utterances, *Artificial intelligence* 15 (3), 143-178.
- Allen, J. F., Schubert, L. K., Ferguson, G., Heeman, P., Hwang, C. H., Kato, T., Light, M., Martin, N., Miller, B., Poesio, M., Traum, D. R., The TRAINS project: A case study in building a conversational planning agent *Journal of Experimental and Theoretical Artificial Intelligence*, 1995
- Allen, J. F., André, E., Cohen, P. R., Hakkani-Tür, D., Kaplan, R., Lemon, O., Traum, D., Challenge discussion: Advancing multimodal dialogue, Chapter 5 in *Handbook of Multimodal-Multisensor Interfaces*, Oviatt, S. L., Schuller, B., Cohen, P. R., Sonntag, D., Potamianos, G., and Krüger, A., ACM Press/Morgan and Claypool Publishers, 2019.
- Appelt, D. *Planning English Sentences*, Cambridge University Press, Cambridge, UK, 1985
- Barcan, R. C., A Functional Calculus of First Order Based on Strict Implication, *Journal of Symbolic Logic*, 11, 1946.
- Bapna, A., Tür, G., Hakkani-Tür, D., and Heck, L., Sequential dialogue context modelling for spoken language understanding, *Proc. of SIGDIAL*, 2017, 103-114.
- Bobrow, D. G., Kaplan, R. M., Kay, M., Norman, D. A., Thompson, H., and Winograd, T. GUS, a frame-driven dialog system. *Artificial Intelligence*, 8(2), 1977, 155-173.
- Bohus, D. and Rudnicky, A. I., The RavenClaw dialogue management framework, *Computer Speech and Language*, 23, 2009, 332-361.
- Branavan, R.K., Kushman, N., Lei, T., Barzilay, R. Learning High-Level Planning from Text, Proc. ACL-12, 2012, 126-135.
- Clark, H. H., and Wilkes-Gibbs, D., Referring as a collaborative process, *Cognition*(22), 1986, 1-39

¹³ This might take longer than overnight (vs. Wang et al. 2015).

- Cohen, P. R. On knowing what to say: Planning speech acts. PhD Thesis, Dept. of Computer Science, University of Toronto, 1978.
- Cohen, P. R. and Levesque, H. J., Intention is choice with commitment, *Artificial Intelligence*, 42 (2-3), 1990, 213-261.
- Cohen, P. R., and Levesque, H. J. , Rational Interaction as the Basis for Communication *Intentions in Communication*, Cohen, P. R., Morgan, J. and Pollack, M.E., MIT Press, 1990a.
- Cohen, P. R. and Perrault, C. R., Elements of a plan-based theory of speech acts, *Cognitive Science*, 3(3), 1979.
- Cohen, P. R., Perrault, C. R., and Allen, J. F., Beyond question-answering, in Strategies for Natural Language Processing, Lehnert, W. and Ringle, M. (eds), Lawrence Erlbaum Associates, 1982.
- Duong, L., Afshar, H., Estival, D., Pink, G., Cohen, P. R., and Johnson M. Multilingual Semantic Parsing and Code-switching, *Proc. of the 21st Conf. on Computational Natural Language Learning (CoNLL 2017)*, 2017, pp. 379-389.
- Duong, L., Afshar, H. Estival, D., Pink, G., Cohen, P., Johnson M., Active learning for deep semantic parsing. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, 2018, 43-48.
- Fast, E., McGrath, W., Rajpurkar, P. and Bernstein, M., Augur: Mining Human Behaviors from Fiction to Power Interactive Systems. *Proc. of the 2016 CHI Conference on Human Factors in Computing Systems*, ACM Press, 2016.
- Galescu, L., Teng, C. M., Allen J. F., and Pereira, I. Cogent: A Generic Dialogue System Shell Based on a Collaborative Problem Solving Model, *Proceedings of SigDial*, 2018, 400-409.
- Gašić, M., Mrkšić, N., Rojas-Barahona, L. M., Su, P-H., Ultes, S., Vandyke, D., Wen, T-H., and Young, S., Dialogue manager domain adaptation using Gaussian process reinforcement learning, *Computer Speech and Language* 45, 2016, 552-569.
- Grice, H.P. Logic and Conversation, *Syntax and Semantics*, vol.3 P. Cole and J. Morgan (eds.), Academic Press, 1975.
- Henderson, M., Machine learning for dialog state tracking: A review, *Proceedings of The First International Workshop on Machine Learning in Spoken Language Processing*, 2015.
- Horvitz, E., Reflections on challenges and promises of mixed-initiative interaction, *AI Magazine*, 28(2), 2007, 19-22.
- Johnson B., Floyd M.W., Coman A., Wilson M.A., Aha D.W. Goal reasoning and trusted autonomy. In: Abbass H., Scholz J., Reid D. (eds), *Foundations of Trusted Autonomy. Studies in Systems, Decision and Control*, vol 117. Springer, 47-66, 2018.
- Kaplan, D. Quantifying in, *Synthese* 19(1/2), 1968, 178-214.
- Konolige, K., On the relation between autoepistemic logic and circumscription: Preliminary Report, *Proc. of IJCAI*, 1989, 1213-1218.
- Kripke, S. A. Semantical Considerations on Modal Logic *Acta Philosophica Fennica* 16 1963, 83-94.
- Jiang, T., and Riloff, E., Learning prototypical goal activities for locations, *Proc. of Assoc. for Comp. Ling.*, 2018, 1297-1307.
- Larsson, S. and Traum, D. R., Information state and dialogue management in the TRINDI Dialogue Move Engine Toolkit, *Natural Language Engineering* 6(3-4), 2000, 323-340.
- Litman, D. J. and Allen, J. F., A Plan Recognition Model for Subdialogues in Conversations, *Cognitive Science*, 11, 1987, 163-200.
- Miller, T., Felli, P., Muise, C., Pearce, A. R., and Sonenberg, L. 'Knowing whether' in Proper Epistemic Knowledge Bases, *Proc. of AAAI*, 2017.
- Morbini, F., DeVault, D., Sagae, K., Gerten, J., Nazarian, A., and Traum D., FLoReS: A Forward Looking, Reward Seeking, Dialogue Manager, *Proceedings of the 4th International Workshop on Spoken Dialog Systems* November, 2012, 151-162.
- Moore, Robert C, Reasoning about knowledge and action, *Proc. of IJCAI*, 1977.
- Perrault, C. R. and Allen, J. F., A plan-based analysis of indirect speech acts, *Computational Linguistics*, 6(3-4), 1980, 167-182.
- Rao, A. and Georgeff, M. BDI-agents: From Theory to Practice". *Proceedings of the First International Conference on Multiagent Systems*, 1995.
- Rao, A. and Georgeff, M. Decision procedures for BDI logics, *Journal of Logic and Computation* 8(3), 1998.
- Quine, W. V. O. Quantifiers and propositional attitudes, *Journal of Philosophy* 53(5), 1956, 177-187.

- Sadek, D., Bretier, P., and Panaget, F., ARTIMIS: Natural dialogue meets rational agency, *Proc. IJCAI-15*, 1997, pp. 1030-1035.
- Schatzmann, J., Thomson, B., Weilhammer, K., Ye, H., and Young, S., Agenda-Based User Simulation for Bootstrapping a POMDP Dialogue System, *Proc. of NAACL-HLT*, 2007.
- Shah, P., Hakkani-Tür, D., Tür, G., Rastogi, A., Bapna, A., Nayak, N., Heck, L., Building a conversational agent overnight with dialogue self-play, *arXiv: 1801.04871v1*, Jan., 2018.
- Smith, I. A., and Cohen, P. R. Toward a semantics for an agent communications language *Proc. AAAI-96*, 24-31.
- Steedman, M. and Petrick, R. Planning dialogue actions, *Proc. of SigDial*, 2007.
- Stone, M. Intention, interpretation and the computational structure of language, *Cognitive Science* 28, 2004, 781–809.
- Sukthankar, G., Geib, C., Bui, H., Pynadath, D., and Goldman, R., *Plan, Activity, and Intent Recognition: Theory and Practice*, San Francisco: Morgan Kauffman Publishers, 2014.
- Traum, D. R. and Hinkelman, E. A., Conversation acts in task-oriented spoken dialogue, *Computational Intelligence*, 8(3), 575-599.
- Ultes, S. Budzianowski, P., Casanueva, I., Rojas-Barahona, L., Tseng B-H., Wu, Y-C., Young, S., and Gašić, M. Addressing Objects and Their Relations: The Conversational Entity Dialogue Model, *Proc. of SigDial* 2018.
- Vered, M., Kaminka, G. A. and Biham S. Online Goal Recognition through Mirroring: Humans and Agents. In *Proceedings of the Annual Conference on Advances in Cognitive Systems (ACS)*, 2016.
- Wang, Y., Berant, J., and Liang, P., Building a semantic parser overnight, *Proc. of Assoc. for Comp. Ling.*, 2015, 1332–1342.
- Williams, J. D., and Young, S. Partially observable Markov decision processes for spoken dialog systems, *Computer Speech and Language* 21 (2007), 393-422.
- Woods, W. A. What's in a Link: Foundations for Semantic Networks. In D. Bobrow and A. Collins (eds.), *Representation and Understanding: Studies in Cognitive Science*, New York: Academic Press, 1975.
- Young, S., Gašić, M., Keizer, S., Mairesse, F., Schatzmann, J., Thomson, B., and Yu, K. The Hidden Information State model: A practical framework for POMDP-based spoken dialogue management, *Computer Speech and Language* 24, 2010, pp. 150-174.
- Zue, V. W., Glass, J., Goodine, D., Hirschman, L. Leung, H. C., Phillips, M., Polifroni, J., Seneff, S. "The MIT ATIS system: Preliminary development, spontaneous speech data collection, and performance evaluation", *Proc of EUROSPEECH*, 1991, 537-540.

Appendix

The Language \mathcal{L}

Variables and constants

$\langle \text{Action-var} \rangle ::= a, b, a_1, a_2 \dots b_1, b_2 \dots e, e_1, e_2 \dots$
 $\langle \text{Agent-var} \rangle ::= x, y, x_1, x_2 \dots y_1, y_2 \dots$
 $\langle \text{Regular-var} \rangle ::= i, j, i_1, i_2 \dots j_1, j_2 \dots$
 $\langle \text{Variable} \rangle ::= \langle \text{Agent-var} \rangle | \langle \text{Action-var} \rangle | \langle \text{Regular-var} \rangle |$
 $[\langle \text{Variable}_1 \rangle \dots \langle \text{Variable}_n \rangle]$, i.e.,
 (a list of variables)

Predicates and Formulas

$\langle \text{Role} \rangle ::=$ distinguished Role symbols for a given action
 $\langle \text{Role-list} \rangle ::= [\langle \text{Role} \rangle_1 : \langle \text{Variable} \rangle_1, \dots, \langle \text{Role} \rangle_n : \langle \text{Variable} \rangle_n]$
 $\langle \text{Pred-symbol} \rangle ::=$ an element of a distinguished set of predicate symbols
 $\langle \text{Pred} \rangle ::= (\langle \text{Pred-symbol} \rangle)$

Well-formed formulas (WFFS)

$\langle \text{Wff} \rangle ::= \langle \text{Pred} \rangle | \sim \langle \text{Wff} \rangle | \langle \text{Wff} \rangle \vee \langle \text{Wff} \rangle | \langle \text{Wff} \rangle \wedge \langle \text{Wff} \rangle | \exists \langle \text{Variable} \rangle \langle \text{Wff} \rangle |$
 $\diamond \langle \text{Wff} \rangle$ — $\langle \text{Wff} \rangle$ is true eventually
 $\square \langle \text{Wff} \rangle$ — $\langle \text{Wff} \rangle$ is always true (note that $\square \langle \text{Pred} \rangle =_{\text{def}} \sim \diamond \sim \langle \text{Pred} \rangle$)
 $\langle \text{Variable} \rangle = \langle \text{Variable} \rangle$
 $(\text{DOES } \langle \text{Action-expr} \rangle)$ — $\langle \text{Action-expr} \rangle$ happens next,
 $(\text{DONE } \langle \text{Action-expr} \rangle)$ — $\langle \text{Action-expr} \rangle$ has *just* happened,
 $(\text{Agt } \langle \text{Agent-var} \rangle \langle \text{Action-var} \rangle)$: $\langle \text{Agent-var} \rangle$ is the only agent of $\langle \text{Action-var} \rangle$,
 $(\text{BEL } \langle \text{Agent-var} \rangle \langle \text{Wff} \rangle)$ — meaning $\langle \text{Wff} \rangle$ follows from $\langle \text{Agent-var} \rangle$'s beliefs,
 $(\text{GOAL } \langle \text{Agent-var} \rangle \langle \text{Wff} \rangle)$ — meaning $\langle \text{Wff} \rangle$ follows from $\langle \text{Agent-var} \rangle$'s goals,
 $\langle \text{Time-proposition} \rangle ::= \langle \text{Numeral} \rangle$
 $(\text{LATER } \langle \text{Wff} \rangle)$::= $\sim \langle \text{Wff} \rangle \wedge \diamond \langle \text{Wff} \rangle$
 $\langle \text{Wff} \rangle$ is false now but eventually true
 $(\text{BEFORE } \langle \text{Wff} \rangle_1 \langle \text{Wff} \rangle_2)$
 before $\langle \text{Wff} \rangle_1$ becomes true,

Action expressions:

$\langle \text{Action-name} \rangle ::=$ an element of a designated set of action names
 $\langle \text{Action-expr} \rangle ::=$
 $\langle \text{Action-var} \rangle$ or one of the following:
 $\langle \text{Action} \rangle ::= \langle \text{Action-name} (\text{Role-list}) \rangle$
 $\langle \text{Action-expr} \rangle ; \langle \text{Action-expr} \rangle$ — sequential action,
 $\langle \text{Action-expr} \rangle | \langle \text{Action-expr} \rangle$ — nondeterministic choice action,
 $\langle \text{Wff} \rangle ?$ — test action
 $\langle \text{Action-expr} \rangle || \langle \text{Action-expr} \rangle$ — concurrent action
 $\langle \text{Action-expression} \rangle^*$: iterative action.

Examples of Well-formed Formulas:

$(\text{DONE } \text{john eat}(\text{john}, \text{vittorios}, \text{mond}, \text{7pm}))$
 Joe has just eaten at Vittorio's on Monday at 7pm.
 $(\text{GOAL } \text{john } \diamond (\text{DONE } \text{john eat}(\text{john}, \text{vittorios}, \text{mond}, \text{7pm})))$
 Joe's goal is to eventually have eaten at Vittorio's at Monday at 7pm.
 $(\text{BEL } \text{john } (\text{PGOAL } \text{mary } (\text{KNOWREF } \text{john } \text{variable:Time } (\text{PGOAL } \text{mary } \text{eat}(\text{mary}, \text{vittorios}, \text{mond}, \text{Time}))))$
 John believes Mary has a persistent goal for him to know the time that Mary wants to eat at Vittorio's on Monday.

Speech Act definitions

$\text{whq}([\text{Speaker}, \text{Listener}, \text{Var}, \text{Pred}])$	
Precondition:	$(\text{KNOWREF } \text{Listener}, \text{Var}, \text{Pred})$
Effect:	$(\text{KNOWREF } \text{Speaker}, \text{Var}, \text{Pred})$
Constraint:	$\text{Speaker} \neq \text{Listener}$
$\text{informref}([\text{Speaker}, \text{Listener}, \text{Var}, \text{Pred}])$	
Precondition:	$(\text{KNOWREF } \text{Speaker}, \text{Var}, \text{Pred})$
Effect:	$(\text{KNOWREF } \text{Listener}, \text{Var}, \text{Pred})$
Constraint:	$\text{Speaker} \neq \text{Listener}$
$\text{ynq}([\text{Speaker}, \text{Listener}, \text{Pred}]),$	
Precondition:	$(\text{KNOWIF } \text{Listener } \text{Pred})$
Effect:	$(\text{KNOWIF } \text{Speaker } \text{Pred})$
Constraint:	$\text{Speaker} \neq \text{Listener}$
$\text{inform}([\text{Speaker}, \text{Listener}, \text{Pred}])$	
Precondition:	$(\text{BEL } \text{Speaker } \text{Pred})$
Effect:	$(\text{BEL } \text{Listener } \text{Pred})$
Constraint:	$\text{Speaker} \neq \text{Listener}$

Speaker-adapted neural-network-based fusion for multimodal reference resolution

Diana Kleingarn

Ruhr University Bochum
diana.kleingarn@rub.de

Nima Nabizadeh

Ruhr University Bochum
nima.nabizadeh@rub.de

Martin Heckmann

Honda Research Institute Europe GmbH
martin.heckmann@honda-ri.de

Dorothea Kolossa

Ruhr University Bochum
dorothea.kolossa@rub.de

Abstract

Humans use a variety of approaches to reference objects in the external world, including verbal descriptions, hand and head gestures, eye gaze or any combination of them. The amount of useful information from each modality, however, may vary depending on the specific person and on several other factors. For this reason, it is important to learn the correct combination of inputs for inferring the best-fitting reference. In this paper, we investigate speaker-dependent and independent fusion strategies in a multimodal reference resolution task. We show that without any change in the modality models, only through an optimized fusion technique, it is possible to reduce the error rate of the system on a reference resolution task by more than 50%.

1 Introduction

Reference resolution is of vital importance when human-machine interaction is expected to become natural and be integrated into everyday life. Humans have at their disposal a broad range of modalities to refer to objects in their environment, including verbal and material signals (Clark, 2005). Equipping machines with the capability to correctly interpret such reference resolutions raises the question of how to fuse the information derived from the different modalities.

Popular fusion methods in this domain can be categorized along two dimensions. The first is at which level of processing the fusion happens and the second how the fusion is performed (see Atrey et al. (2010); Ramachandram and Taylor (2017) for a comprehensive overview). In so-called early fusion or feature level fusion the features derived from the different modalities are combined, whereas in late fusion or decision level fusion classification results, e.g. in the form of probabilities, are combined. Regarding the second

dimension, the methods are mainly grouped into classification-based and estimation-based methods.

As for the classification-based techniques, the modalities are usually combined at the feature level, i.e. early fusion, and the decision is obtained using a classifier. Iida et al. (2011) approached a reference resolution task in which two humans collaboratively solve a Tangram puzzle. Their method computed linguistic, gaze and task-specific features for each object of the board game and the objects were ranked using an SVM classifier. In a similar puzzle task, Funakoshi et al. (2012) proposed a model that could resolve verbal descriptions as well as gestures utilizing a Bayesian network. The Bayesian network design was later employed by Whitney et al. (2016) for interpreting referring expressions with speech and pointing gestures in a real-world cooking task.

Regarding the rule-based fusion, linear weighted fusion is one of the simplest and most widely used rule-based methods. This method combines the information from the different modalities linearly and it is assumed that the share of each modality in decision making does not change. It has been successfully utilized in multiple studies on reference resolution (Matuszek et al., 2014; Prasov and Yue Chai, 2010; Kennington et al., 2015; Kennington and Schlangen, 2017). A constraint-based rule system was used by (Holzapfel et al., 2004) where the constraints considered the time correlation of events and their semantic content for the fusion.

In this paper, we concentrate on one rule-based method used in Kennington et al. (2015). For this purpose, first, we explain the task and dataset in Sec. 2. Then, we discuss different approaches for the fusion of data in Sec. 3, including linear weighted fusion (Sec. 3.1) and our proposed neural-network-based fusion (Sec. 3.2), which

also provides the possibility of learning speaker-dependent weights (Sec. 3.3). We summarize the results in Sec. 4 and give a short conclusion and outlook on future work in Sec. 5.

2 Previous work

2.1 The TAKE Dataset




Figure 1: Example PENTO board on the TAKE dataset (Kennington and Schlangen, 2017)

The TAKE dataset was first introduced in Kousidis et al. (2013). It is a Wizard-of-Oz study, in which the participants were placed in front of a screen showing 15 pieces of a PENTO board game in random colors and shapes. The pieces were grouped into the four corners of the screen. For every episode, the shown objects and their positions on the screen was set randomly.

The participants were asked to instruct the system to select one specific PENTO piece on the board per episode. There was no instruction telling the participants how to refer to the item. According to the setup, it was possible to specify the object using spoken words, pointing gestures or eye gaze. Next, one piece was marked and the participant confirmed whether this selection was correct.

The example episode below, corresponding to Fig. 1, shows the English translation of the speech input and the true referent identifier:

- then we take now the se- so the second t that is on the top right ... out of this group there I would like to have the yellow t ... yes
- REFERENT o3

For this work, the confirmation utterance, e.g. the word “yes” in the above example, was removed, since it is not available at the time the decision is made. After this cleanup, the dataset includes 1034 episodes distributed over 7 users as shown in Table 1. The participants were native speakers, except for one, who spoke proficient but not native German.

User	Episodes	With pointing	With gaze
1	90	87	71
2	66	29	64
3	133	35	126
4	230	209	212
5	146	13	130
6	176	78	157
7	193	162	164
Total	1034	613	924

Table 1: Number of episodes, per user and cumulatively, in the TAKE dataset.

The speech, an average of 6.8 words per utterance, was transcribed using Google Web Speech as an automatic speech recognition (ASR), with a vocabulary size of 1049. Additionally, the speech was transcribed by hand, which can provide a reasonable upper bound for the results. A Microsoft Kinect above the screen captured the arm movements and an eye tracker (*Seeingmachines FaceLab*) was used to determine the eye gaze.

Since the scenes in this dataset are virtual, we can directly annotate the objects with the properties and then query the scene representation. For this simplified task, the properties are the color, the shape and the spatial relations of the pieces. Using image processing techniques described in Kennington et al. (2015), several features for each object are extracted, including the number of edges, RGB (red, green, blue) values, HSV (hue, saturation, value), its centroid, horizontal and vertical skewness, and the orientation value denoting the direction of the principal axis. These features are used for the natural language grounding described in the next section.

2.2 Model for Natural Language Understanding

The idea is to treat each word in the vocabulary as a classifier which can relate the word to the perceptual information of the objects. For this purpose, a logistic regression classifier is trained to map the visual features \mathbf{x} of each particular candidate object to a probability p_w of these features, given the word w .

$$p_w(\mathbf{x}) = \sigma(\mathbf{w}^\top \mathbf{x} + b) \quad (1)$$

Here, \mathbf{w} is the learned weight vector and σ is the logistic function. What is needed for further steps, however, is one distribution over all candidate objects per episode. To accomplish that, we can average the distribution of all time steps $n = 1 \dots N$

and normalize the prediction score of each object ($i \in I$) over all the $|I| = 15$ object candidates via

$$p_{\text{speech}}(i) = \frac{\sum_{n=1}^N p_{w_n}(\mathbf{x}_i)}{\sum_{k=1}^{|I|} \sum_{n=1}^N p_{w_n}(\mathbf{x}_k)}. \quad (2)$$

2.3 Model for Pointing Gestures and Gaze

For gaze and pointing gestures, we need a model that takes the coordinates of gaze and pointing as its input and returns a probability distribution over the object candidates, given the location of objects. This model is the same for gaze and pointing gestures.

For this purpose, we compute the average of the gaze or pointing coordinates for each episode, producing a reference point (R) for the modality. The reference point is compared to the centroid of each object (x_i, y_i) using a Gaussian distribution,

$$p_d(i) \propto \exp \left[-\frac{(x_R - x_i)^2}{2 \cdot \sigma_x^2} - \frac{(y_R - y_i)^2}{2 \cdot \sigma_y^2} \right]. \quad (3)$$

The result is then normalized over all objects to obtain p_{point} and p_{gaze} , so that the objects closer to the reference point will have a higher probability.

3 Fusion Models

3.1 Linear Fusion

For optimum performance, all three modalities need to be combined. A simple approach is to perform a rule-based late fusion by estimating a fixed weight for each modality and then summing the weighted prediction distributions, as in [Kennington \(2016\)](#):

$$p(i) = p_{\text{speech}}(i) \cdot \alpha_1 + p_{\text{point}}(i) \cdot \alpha_2 + p_{\text{gaze}}(i) \cdot (1 - \alpha_1 - \alpha_2). \quad (4)$$

The system then makes a maximum-likelihood decision according to

$$\hat{i} = \arg \max_{i \in I} p(i). \quad (5)$$

3.2 Neural-network-based Fusion

In [Sec. 3.1](#), a baseline approach to late fusion is shown. To decrease the error rate, we now propose a more flexible method, which can model non-linear relations between the modalities. For this purpose we chose a fully connected neural network with one hidden layer, 512 neurons and a rectified linear unit as the activation function.

Its inputs \mathbf{o} are the three concatenated modality vectors from (2) and (3),

$$\begin{aligned} \mathbf{o} &= [\mathbf{p}_{\text{speech}}, \mathbf{p}_{\text{point}}, \mathbf{p}_{\text{gaze}}] \quad (6) \\ \text{with } \mathbf{p} &= [p_1, \dots, p_{|I|}]. \end{aligned}$$

The output layer uses the softmax function so that the output can be interpreted as a probability distribution and used in [Eq. \(5\)](#) to obtain the estimated referent. To optimize the network parameters, we carried out preliminary tests with differently sized hidden layers and with additional reliability information, e.g., the variance of gaze or pointing information. For hand-annotated data, including the variance of all deixis coordinates of the current episode, \mathbf{V} , in the observation vector gave the best results. With this update, the network input becomes

$$\mathbf{o} = [\mathbf{p}_{\text{speech}}, \mathbf{p}_{\text{point}}, \mathbf{p}_{\text{gaze}}, \mathbf{V}]. \quad (7)$$

3.3 Speaker adaptation

Humans have different preferences in the way they refer to objects. This is also reflected in the dataset, in which many episodes from one participant are quite alike, whereas significant differences can often be observed across participants. Hence, depending on the participant, different modalities are very likely to contribute a variable amount of useful information. A model that adapts to a specific user should therefore outperform a general model.

However, judging from the small number of samples per user in [Tab. 1](#), it is evidently not promising to train a neural network using only the data of one participant. Inspired by [Saon et al. \(2013\)](#), we addressed this problem by training on the full training set and reducing to a smaller training set, containing just one user, for the last 5 % of the epochs.

4 Evaluation

We evaluate all fusion methods on the same data as [Kennington et al. \(2015\)](#) under the same four conditions: speech only, speech with gaze, speech with deixis, and speech with gaze and deixis. For this purpose, we compare the error rate $E = 100 \cdot \frac{M-C}{M}$ under all conditions, with C as the number of correctly estimated referents, among M estimates made for the test set.

However, for the linear fusion with fixed weights (fw) presented in [Sec. 3.1](#), we did not use

the weights suggested in [Kennington et al. \(2015\)](#). Instead, a grid search was run on the training data to determine optimal weights for the dataset (ow). This yielded an average improvement of 5.9% absolute for hand-annotated data and also improved all individual cases for ASR-annotated data except for the fusion of all modalities. Here, the results slightly deteriorated from 60.3% to 60.0%.

We used 10-fold cross validation to obtain an estimate of the error rate together with its standard deviation. These results are depicted in Fig. 2. As can be seen, there is a large difference in




Figure 2: Error rate (%) and standard deviation for optimized (ow) or fixed weights (fw, adapted from ([Kennington et al., 2015](#))) in (4).

performance between the results using the hand-annotated speech data vs. the ASR system, indicating a likely high number of transcription errors for the informative keywords. It can also be seen that adding more modalities consistently improves the performance.




Figure 3: Error rate (%) of the proposed neural-network-based fusion

The neural network-based fusion (Sec. 3.2) increased performance compared to the linear fusion (fw) notably and for all conditions. These results are shown in Fig. 3. We obtain the best results with an error rate of 8.9% for the fusion of all modalities using the hand-annotated data. In comparison to the fixed-weight baseline, with an error rate of 30% (see Fig. 2), the error rate is hence decreased

by 70%.

User	NN ASR	NN hand
User 1	17.5 (± 8.5) 35.8 (± 10.9)	3.4 (± 5.8) 8.5 (± 5.6)
User 2	11.7 (± 13.9) 16.2 (± 8.9)	11.0 (± 11.9) 12.1 (± 11.6)
User 3	10.8 (± 12.8) 11.9 (± 11.8)	3.5 (± 6.5) 4.5 (± 8.7)
User 4	12.3 (± 10.2) 10.5 (± 9.1)	5.7 (± 6.2) 5.2 (± 6.9)
User 5	22.6 (± 7.9) 28.3 (± 11.8)	6.5 (± 7.7) 11.3 (± 8.9)
User 6	19.1 (± 8.8) 23.4 (± 10.6)	12.5 (± 8.9) 14.8 (± 12.5)
User 7	31.0 (± 13.5) 24.0 (± 10.7)	6.4 (± 9.0) 4.7 (± 7.0)
average	18.6 (± 10.7) 22.2 (± 10.3)	7.0 (± 7.8) 8.9 (± 8.2)

Table 2: Results of the user-dependent (black) and the user-independent (gray) model in terms of error rate (%) and standard deviation σ .

Table 2 compares the results of the speaker-dependent and -independent models for each user. Here, we only report the results for the fusion of all modalities. When using the hand annotation, the speaker-adapted fusion reduces the error rate further, from 8.9% to 7.0%. But it can also be seen that the results vary largely from user to user. In particular, for user 1 (ASR data), the speaker-adapted version outperforms the other version easily, but for user 7, the original, speaker-independent version is more accurate. For hand-annotated data, the difference between the two versions is smaller, but the users for which the speaker-adapted version outperforms the other remain the same. Interestingly the speaker-adapted version performs least well for the two users with the most episodes that mostly contain gaze and pointing information, as can be seen in Table 1.

5 Conclusions

We have compared different fusion strategies for multi-modal information integration in a reference resolution task. Our results show that a fully connected neural network can reduce the error rate significantly, compared to a weighted averaging of single-modality posterior probabilities. Adapting the fusion to each specific user is also helpful to some extent, although the improvements are less clear and consistent.

In this work, we applied fairly simple models for speech, gaze and pointing, which simply use the average values of all features for the current episode. Since some words carry more semantic content than others for finding the referent, and since the coordinate sequences of gaze and pointing contain some redundancy, as well as segments of more and of less information content, future work will focus on the creation of a time-dependent model for improving multi-modal fusion.

6 Acknowledgments

We want to thank Casey Kennington for fruitful discussions and his support in setting up the model for linear fusion. Furthermore, we are grateful to Casey Kennington and Sina Zarrie for providing the data. In addition, many thanks to Julia Ringeis for supporting us in preparing the manuscript.

References

- Pradeep K Atrey, M Anwar Hossain, Abdulmotaleb El Saddik, and Mohan S Kankanhalli. 2010. [Multimodal fusion for multimedia analysis: a survey](#). *Multimedia systems*, 16(6):345–379.
- Herbert H Clark. 2005. [Coordinating with each other in a material world](#). *Discourse studies*, 7(4-5):507–525.
- Kotaro Funakoshi, Mikio Nakano, Takenobu Tokunaga, and Ryu Iida. 2012. [A unified probabilistic approach to referring expressions](#). In *Proceedings of the 13th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL 2012)*, pages 237–246.
- Hartwig Holzapfel, Kai Nickel, and Rainer Stiefelhagen. 2004. [Implementation and evaluation of a constraint-based multimodal fusion system for speech and 3d pointing gestures](#). In *Proceedings of the 6th international conference on Multimodal interfaces*, pages 175–182.
- Ryu Iida, Masaaki Yasuhara, and Takenobu Tokunaga. 2011. [Multi-modal reference resolution in situated dialogue by integrating linguistic and extralinguistic clues](#). In *Proceedings of 5th International Joint Conference on Natural Language Processing*, pages 84–92.
- Casey Kennington, Livia Dia, and David Schlangen. 2015. [A Discriminative Model for Perceptually-Grounded Incremental Reference Resolution](#). In *Proceedings of the 11th International Conference on Computational Semantics (IWCS) 2015*, pages 195–205.
- Casey Kennington and David Schlangen. 2017. [A Simple Generative Model of Incremental Reference Resolution for Situated Dialogue](#). *Comput. Speech Lang.*, 41(C):43–67.
- Casey Redd Kennington. 2016. *Incrementally Resolving References in Order to Identify Visually Present Objects in a Situated Dialogue Setting*. Ph.D. thesis, Bielefeld University.
- Spyros Kousidis, Casey Kennington, and David Schlangen. 2013. [Investigating speaker gaze and pointing behaviour in human-computer interaction with the mint. tools collection](#). In *Proceedings of the 14th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL 2013)*, pages 319–323.
- Cynthia Matuszek, Liefeng Bo, Luke Zettlemoyer, and Dieter Fox. 2014. [Learning from unscripted deictic gesture and language for human-robot interactions](#). In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, pages 2556–2563.
- Zahar Prasov and Joyce Yue Chai. 2010. [Fusing eye gaze with speech recognition hypotheses to resolve exophoric references in situated dialogue](#). In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 471–481.
- Dhanesh Ramachandram and Graham W Taylor. 2017. [Deep multimodal learning: A survey on recent advances and trends](#). *IEEE Signal Processing Magazine*, 34(6):96–108.
- George Saon, Hagen Soltau, David Nahamoo, and Michael Picheny. 2013. [Speaker adaptation of neural network acoustic models using i-vectors](#). In *2013 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, pages 55–59.
- David Whitney, Miles Eldon, John Oberlin, and Stefanie Tellex. 2016. [Interpreting multimodal referring expressions in real time](#). In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3331–3338.

Learning Question-Guided Video Representation for Multi-Turn Video Question Answering

Guan-Lin Chao
Carnegie Mellon University
guanlinchao@cmu.edu

Abhinav Rastogi
Google AI
abhirast@google.com

Semih Yavuz
University of California, Santa Barbara
syavuz@cs.ucsb.edu

Dilek Hakkani-Tür
Amazon Alexa AI
dilek@ieee.org

Jindong Chen
Google AI
jdchen@google.com

Ian Lane
Carnegie Mellon University
lane@cmu.edu

Abstract

Understanding and conversing about dynamic scenes is one of the key capabilities of AI agents that navigate the environment and convey useful information to humans. Video question answering is a specific scenario of such AI-human interaction where an agent generates a natural language response to a question regarding the video of a dynamic scene. Incorporating features from multiple modalities, which often provide supplementary information, is one of the challenging aspects of video question answering. Furthermore, a question often concerns only a small segment of the video, hence encoding the entire video sequence using a recurrent neural network is not computationally efficient. Our proposed question-guided video representation module efficiently generates the token-level video summary guided by each word in the question. The learned representations are then fused with the question to generate the answer. Through empirical evaluation on the Audio Visual Scene-aware Dialog (AVSD) dataset (Alamri et al., 2019a), our proposed models in single-turn and multi-turn question answering achieve state-of-the-art performance on several automatic natural language generation evaluation metrics.

1 Introduction

Nowadays dialogue systems are becoming more and more ubiquitous in our lives. It is essential for such systems to perceive the environment, gather data and convey useful information to humans in an accessible fashion. Video question answering (VideoQA) systems provide a convenient way for humans to acquire visual information about the environment. If a user wants to obtain information about a dynamic scene, one can simply ask the VideoQA system a question in natural language, and the system generates a natural-language answer. The task of a VideoQA dialogue system in

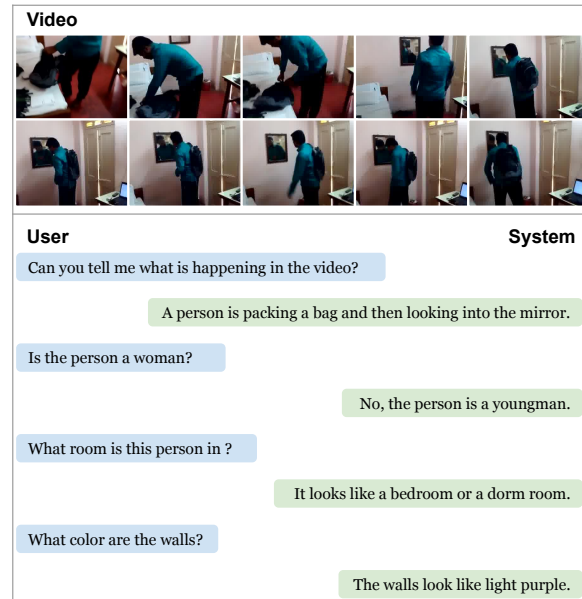


Figure 1: An example from the AVSD dataset. Each example contains a video and its associated question answering dialogue regarding the video scene.

this paper is described as follows. Given a video as grounding evidence, in each dialogue turn, the system is presented a question and is required to generate an answer in natural language. Figure 1 shows an example of multi-turn VideoQA. It is composed of a video clip and a dialogue, where the dialogue contains open-ended question answer pairs regarding the scene in the video. In order to answer the questions correctly, the system needs to be effective at understanding the question, the video and the dialogue context altogether.

Recent work on VideoQA has shown promising performance using multi-modal attention fusion for combination of features from different modalities (Xu et al., 2017; Zeng et al., 2017; Zhao et al., 2018; Gao et al., 2018). However, one of the challenges is that the length of the video sequence can be very long and the question may concern only

a small segment in the video. Therefore, it may be time inefficient to encode the entire video sequence using a recurrent neural network.

In this work, we present the question-guided video representation module which learns 1) to summarize the video frame features efficiently using an attention mechanism and 2) to perform feature selection through a gating mechanism. The learned question-guided video representation is a compact video summary for each token in the question. The video summary and question information are then fused to create multi-modal representations. The multi-modal representations and the dialogue context are then passed as input to a sequence-to-sequence model with attention to generate the answer (Section 3). We empirically demonstrate the effectiveness of the proposed methods using the AVSD dataset (Alamri et al., 2019a) for evaluation (Section 4). The experiments show that our model for single-turn VideoQA achieves state-of-the-art performance, and our multi-turn VideoQA model shows competitive performance, in comparison with existing approaches (Section 5).

2 Related Work

In the recent years, research on visual question answering has accelerated following the release of multiple publicly available datasets. These datasets include COCO-QA (Ren et al., 2015a), VQA (Agrawal et al., 2017), and Visual Madlibs (Yu et al., 2015) for image question answering and MovieQA (Tapaswi et al., 2016), TGIF-QA (Jang et al., 2017), and TVQA (Lei et al., 2018) for video question answering.

2.1 Image Question Answering

The goal of image question answering is to infer the correct answer, given a natural language question related to the visual content of an image. It assesses the system’s capability of multi-modal understanding and reasoning regarding multiple aspects of humans and objects, such as their appearance, counting, relationships and interactions (Lei et al., 2018). State-of-the-art image question answering models make use of spatial attention to obtain a fixed length question-dependent embedded representation of the image, which is then combined with the question feature to predict the answer (Yang et al., 2016; Xu and Saenko, 2016; Kazemi and Elqursh, 2017; Anderson et al., 2018).

Dynamic memory (Kumar et al., 2016; Xiong et al., 2016) and co-attention mechanism (Lu et al., 2016; Ma et al., 2018) are also adopted to model sophisticated cross-modality interactions.

2.2 Video Question Answering

VideoQA is a more complex task. As a video is a sequence of images, it contains not only appearance information but also motion and transitions. Therefore, VideoQA requires spatial and temporal aggregation of image features to encode the video into a question-relevant representation. Hence, temporal frame-level attention is utilized to model the temporal dynamics, where frame-level attribute detection and unified video representation are learned jointly (Ye et al., 2017; Xu et al., 2017; Mun et al., 2017). Similarly, Lei et al. (2018) use Faster R-CNN (Ren et al., 2015b) trained with the Visual Genome (Krishna et al., 2017) dataset to detect object and attribute regions in each frame, which are used as input features to the question answering model. Previous works also adopt various forms of external memory (Sukhbaatar et al., 2015; Kumar et al., 2016; Graves et al., 2016) to store question information, which allows multiple iterations of question-conditioned inference on the video features (Na et al., 2017; Kim et al., 2017; Zeng et al., 2017; Gao et al., 2018; Chenyou Fan, 2019).

2.3 Video Question Answering Dialogue

Recently in DSTC7, Alamri et al. (2019a) introduce the Audio-Visual Scene-aware Dialog (AVSD) dataset for multi-turn VideoQA. In addition to the challenge of integrating the questions and the dynamic scene information, the dialogue system also needs to effectively incorporate the dialogue context for coreference resolution to fully understand the user’s questions across turns. To this end, Alamri et al. (2019b) use two-stream inflated 3D ConvNet (I3D) model (Carreira and Zisserman, 2017) to extract spatiotemporal visual frame features (I3D-RGB features for RGB input and I3D-flow features for optical flow input), and propose the Naïve Fusion method to combine multi-modal inputs based on the hierarchical recurrent encoder (HRE) architecture (Das et al., 2017). Hori et al. (2018) extend the Naïve Fusion approach and propose the Attentional Fusion method which learns multi-modal attention weights to fuse features from different modalities. Zhuang et al. (2019) modify the Attentional Fu-

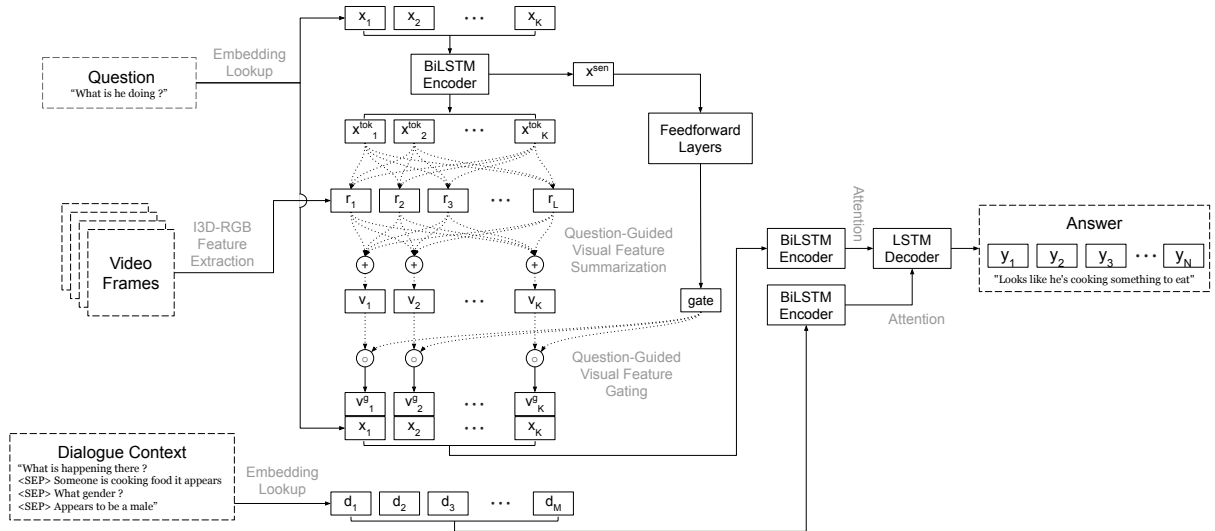


Figure 2: Overview of the proposed approach. First the I3D-RGB frame features are extracted. The question-guided video representation module takes as input the question sentence and the I3D-RGB features, generates a video representation for each token and applies gating using question as guidance. Then the question tokens are augmented by the per-token video representations and encoded by a bidirectional LSTM encoder. Similarly, the dialogue context is encoded by a bidirectional LSTM encoder. Finally, the LSTM answer decoder predicts the answer sequence.

sion method and propose to use Maximum Mutual Information (MMI) (Bahl et al., 1986) as the training objective. Besides the HRE architecture, the multi-source sequence-to-sequence (Multi-Source Seq2Seq) architecture with attention (Zoph and Knight, 2016; Firat et al., 2016) is also commonly applied (Pasunuru and Bansal, 2019; Kumar et al., 2019; Yeh et al., 2019). Previous works (Sanabria et al., 2019; Le et al., 2019; Pasunuru and Bansal, 2019) also explore various attention mechanisms to incorporate the different modal inputs, such as hierarchical attention (Libovický and Helcl, 2017) and cross attention (Seo et al., 2017). For modeling visual features, Lin et al. (2019) propose to use Dynamic memory networks (Kumar et al., 2016) and Nguyen et al. (2019) propose to use feature-wise linear modulation layers (Perez et al., 2018).

3 Approach

We formulate the multi-turn VideoQA task as follows. Given a sequence of raw video frames \mathbf{f} , the embedded question sentence $\mathbf{x} = \{x_1, \dots, x_K\}$ and the single concatenated embedded sentence of the dialogue context $\mathbf{d} = \{d_1, \dots, d_M\}$, the output is an answer sentence $\mathbf{y} = \{y_1, \dots, y_N\}$.

The architecture of our proposed approach is illustrated in Figure 2. First the Video Frame Feature Extraction Module extracts the I3D-

RGB frame features from the video frames (Section 3.1). The Question-Guided Video Representation Module takes as input the embedded question sentence and the I3D-RGB features, and generates a compact video representation for each token in the question sentence (Section 3.2). In the Video-Augmented Question Encoder, the question tokens are first augmented by their corresponding per-token video representations and then encoded by a bidirectional LSTM (Section 3.3). Similarly, in the Dialogue Context Encoder, the dialogue context is encoded by a bidirectional LSTM (Section 3.4). Finally, in the Answer Decoder, the outputs from the Video-Augmented Question Encoder and the Dialogue Context Encoder are used as attention memory for the LSTM decoder to predict the answer sentence (Section 3.5). Our encoders and decoder work in the same way as the multi-source sequence-to-sequence models with attention (Zoph and Knight, 2016; Firat et al., 2016).

3.1 Video Frame Feature Extraction Module

In this work, we make use of the I3D-RGB frame features as the visual modality input, which are pre-extracted and provided in the AVSD dataset (Alamri et al., 2019a). Here we briefly describe the I3D-RGB feature extraction process, and we refer the readers to (Carreira and Zisser-

man, 2017) for more details of the I3D model. Two-stream Inflated 3D ConvNet (I3D) is a state-of-the-art action recognition model which operates on video inputs. The I3D model takes as input two streams of video frames: RGB frames and optical flow frames. The two streams are separately passed to a respective 3D ConvNet, which is inflated from 2D ConvNets to incorporate the temporal dimension. Two sequences of spatiotemporal features are produced by the respective 3D ConvNet, which are jointly used to predict the action class. The I3D-RGB features provided in the AVSD dataset are intermediate spatiotemporal representations from the "Mixed_5c" layer of the RGB stream's 3D ConvNet. The AVSD dataset uses the I3D model parameters pre-trained on the Kinetics dataset (Kay et al., 2017). To reduce the number of parameters in our model, we use a trainable linear projection layer to reduce the dimensionality of I3D-RGB features from 2048 to 256. Extracted from the video frames \mathbf{f} and projected to a lower dimension, the sequence of dimension-reduced I3D-RGB frame features are denoted by $\mathbf{r} = \{r_1, \dots, r_L\}$, where $r_i \in \mathbb{R}^{256}, \forall i$.

3.2 Question-Guided Video Representation Module

We use a bidirectional LSTM network to encode the sequence of question token embedding $\mathbf{x} = \{x_1, \dots, x_K\}$. The token-level intermediate representations are denoted by $\mathbf{x}^{\text{tok}} = \{x_1^{\text{tok}}, \dots, x_K^{\text{tok}}\}$, and the embedded representation of the entire question is denoted by x^{sen} . These outputs will be used to guide the video representation.

$$\vec{h}_0 = \vec{h}_{K+1} = \mathbf{0} \quad (1)$$

$$\vec{h}_k = \text{LSTM}_{\text{guide}}^{\text{forw}}(x_k, \vec{h}_{k-1}) \quad (2)$$

$$\vec{h}_k = \text{LSTM}_{\text{guide}}^{\text{back}}(x_k, \vec{h}_{k+1}) \quad (3)$$

$$x_k^{\text{tok}} = \vec{h}_k \oplus \vec{h}_k \quad (4)$$

$$\forall k \in \{1, \dots, K\}$$

$$x^{\text{sen}} = \vec{h}_K \oplus \vec{h}_1 \quad (5)$$

where \oplus denotes vector concatenation; \vec{h} and \vec{h} represent the local forward and backward LSTM hidden states.

3.2.1 Per-Token Visual Feature Summarization

Generally the sequence length of the video frame features is quite large, as shown in Table 1. There-

fore it is not computationally efficient to encode the video features using a recurrent neural network. We propose to use the attention mechanism to generate a context vector to efficiently summarize the I3D-RGB features. We use the trilinear function (Seo et al., 2017) as a similarity measure to identify the frames most similar to the question tokens. For each question token x_k , we compute the similarity scores of its encoded representation x_k^{tok} with each of the I3D-RGB features \mathbf{r} . The similarity scores \mathbf{s}_k are converted to an attention distribution $\mathbf{w}_k^{\text{att}}$ over the I3D-RGB features by the softmax function. And the video summary v_k corresponding to the question token x_k is defined as the attention weighted linear combination of the I3D-RGB features. We also explored using dot product for computing similarity and empirically found out it yields suboptimal results.

$$s_{k,l} = \text{trilinear}(x_k^{\text{tok}}, r_l) \quad (6)$$

$$= W_{\text{sim}}[x_k^{\text{tok}} \oplus r_l \oplus (x_k^{\text{tok}} \odot r_l)] \quad (7)$$

$$\forall l \in \{1, \dots, L\}$$

$$\mathbf{w}_k^{\text{att}} = \text{softmax}(\mathbf{s}_k) \quad (8)$$

$$v_k = \sum_{l=1}^L w_{k,l}^{\text{att}} r_l \quad (9)$$

$$\forall k \in \{1, \dots, K\}$$

where \odot denotes element-wise multiplication, and W_{sim} is a trainable variable.

3.2.2 Visual Feature Gating

Not all details in the video are important for answering a question. Attention helps in discarding the unimportant frames in the time dimension. We propose a gating mechanism which enables us to perform feature selection within each frame. We project the sentence-level question representation x^{sen} through fully-connected layers with ReLU nonlinearity to generate a gate vector g . For each question token x_k , its corresponding video summary v_k is then multiplied element-wise with the gate vector g to generate a gated visual summary v_k^g . We also experimented applying gating on the dimension-reduced I3D-RGB features \mathbf{r} , prior to the per-token visual feature summarization step, but it resulted in an inferior performance.

$$g = \text{sigmoid}(W_{g,1}(\text{ReLU}(W_{g,2}x^{\text{sen}} + b_{g,2}) + b_{g,1})) \quad (10)$$

$$v_k^g = v_k \odot g \quad (11)$$

$$\forall k \in \{1, \dots, K\}$$

where $W_{g,1}, b_{g,1}, W_{g,2}, b_{g,2}$ are trainable variables.

3.3 Video-Augmented Question Encoder

Given the sequence of per-token gated visual summary $\mathbf{v}^g = \{v_1^g, \dots, v_K^g\}$, we augment the question features by concatenating the embedded question tokens $\mathbf{x} = \{x_1, \dots, x_K\}$ with their associated per-token video summary. The augmented question features are then encoded using a bidirectional LSTM. The token-level video-augmented question features are denoted by $\mathbf{q}^{\text{tok}} = \{q_1^{\text{tok}}, \dots, q_K^{\text{tok}}\}$, and the sentence-level feature is denoted by q^{sen} .

$$\vec{h}_0 = \vec{h}_{K+1} = \mathbf{0} \quad (12)$$

$$\vec{h}_k = \text{LSTM}_{\text{ques}}^{\text{forw}}(x_k \oplus v_k^g, \vec{h}_{k-1}) \quad (13)$$

$$\vec{h}_k = \text{LSTM}_{\text{ques}}^{\text{back}}(x_k \oplus v_k^g, \vec{h}_{k+1}) \quad (14)$$

$$q_k^{\text{tok}} = \vec{h}_k \oplus \vec{h}_k \quad (15)$$

$$\forall k \in \{1, \dots, K\}$$

$$q^{\text{sen}} = \vec{h}_K \oplus \vec{h}_1 \quad (16)$$

where \vec{h} and \vec{h} represent the local forward and backward LSTM hidden states.

3.4 Dialogue Context Encoder

Similar to the video-augmented question encoder, we encode the embedded dialogue context tokens $\mathbf{d} = \{d_1, \dots, d_M\}$ using a bidirectional LSTM. The embedded token-level representations are denoted by $\mathbf{d}^{\text{tok}} = \{d_1^{\text{tok}}, \dots, d_M^{\text{tok}}\}$.

$$\vec{h}_0 = \vec{h}_{M+1} = \mathbf{0} \quad (17)$$

$$\vec{h}_m = \text{LSTM}_{\text{dial}}^{\text{forw}}(d_m, \vec{h}_{m-1}) \quad (18)$$

$$\vec{h}_m = \text{LSTM}_{\text{dial}}^{\text{back}}(d_m, \vec{h}_{m+1}) \quad (19)$$

$$d_m^{\text{tok}} = \vec{h}_m \oplus \vec{h}_m \quad (20)$$

$$\forall m \in \{1, \dots, M\}$$

where \vec{h} and \vec{h} represent the local forward and backward LSTM hidden states.

3.5 Answer Decoder

The final states of the forward and backward LSTM units of the question encoder are used to initialize the state of answer decoder. Let y_n be the output of the decoder at step n , where $1 \leq n \leq N$, y_0 be the special start of sentence token and y_n^{emb} be the embedded representation of y_n . At a decoder step n , the previous decoder hidden state

h_{n-1} is used to attend over \mathbf{q}^{tok} and \mathbf{d}^{tok} to get the attention vectors $h_n^{\text{att},q}$ and $h_n^{\text{att},d}$ respectively. These two vectors retrieve the relevant features from the intermediate representations of the video-augmented question encoder and the dialogue context encoder, both of which are useful for generating the next token of the answer. At each decoder step, the decoder hidden state h_n is used to generate a distribution over the vocabulary. The decoder output y_n^* is defined to be $\text{argmax}_{y_n} p(y_n | y_{\leq n-1})$.

$$h_0 = q^{\text{sen}} \quad (21)$$

$$s_{n,k}^q = v_{\text{ans},q}^\top \tanh(W_{\text{ans},q}[h_{n-1} \oplus q_k^{\text{tok}}]) \quad (22)$$

$$\forall k \in \{1, \dots, K\}$$

$$\mathbf{w}_n^q = \text{softmax}(s_n^q) \quad (23)$$

$$h_n^{\text{att},q} = \sum_{k=1}^K w_{n,k}^q q_k^{\text{tok}} \quad (24)$$

$$s_{n,m}^d = v_{\text{ans},d}^\top \tanh(W_{\text{ans},d}[h_{n-1} \oplus d_m^{\text{tok}}]) \quad (25)$$

$$\forall m \in \{1, \dots, M\}$$

$$\mathbf{w}_n^d = \text{softmax}(s_n^d) \quad (26)$$

$$h_n^{\text{att},d} = \sum_{m=1}^M w_{n,m}^d d_m^{\text{tok}} \quad (27)$$

$$h_n = \text{LSTM}_{\text{ans}}(y_{n-1}^{\text{emb}}, [h_n^{\text{att},q} \oplus h_n^{\text{att},d} \oplus h_{n-1}]) \quad (28)$$

$$p(y_n | y_{\leq n-1}) = \text{softmax}(W_{\text{ans}} h_n + b_{\text{ans}}) \quad (29)$$

$$\forall n \in \{1, \dots, N\}$$

where \mathbf{h} represents the local LSTM hidden states, and $W_{\text{ans},q}, W_{\text{ans},d}, W_{\text{ans}}, b_{\text{ans}}$ are trainable variables.

4 Experiments

4.1 Dataset

We consider the Audio-Visual Scene-aware Dialogue (AVSD) dataset (Alamri et al., 2019a) for evaluating our proposed model in single-turn and multi-turn VideoQA. We use the official release of train set for training, and the public (*i.e.*, prototype) validation and test sets for inference. The AVSD dataset is a collection of text-based human-human question answering dialogues based on the video clips from the CHARADES dataset (Sigurdsson et al., 2016). The CHARADES dataset contains video clips of daily indoor human activities, originally purposed for research in video activity classification and localization. Along with

	Training	Validation	Test
# of dialogues	7659	732	733
# of turns	153,180	14,680	14,660
# of words	1,450,754	138,314	138,790
Avg. length of question (K)	8.5	8.4	8.5
Avg. length of I3D-RGB (L)	179.2	173.0	171.3

Table 1: Data statistics of the AVSD dataset. We use the official training set, and the public (*i.e.*, prototype) validation and test sets. We also present the average length of the question token sequences and the I3D-RGB frame feature sequences to highlight the importance of time efficient video encoding without using a recurrent neural network. The sequence lengths of the questions and I3D-RGB frame features are denoted by K and L respectively in the model description (Section 3).

the video clips and associated question answering dialogues, the AVSD dataset also provides the pre-extracted I3D-RGB visual frame features using a pre-trained two-stream inflated 3D ConvNet (I3D) model (Carreira and Zisserman, 2017). The pre-trained I3D model was trained on the Kinetics dataset (Kay et al., 2017) for human action recognition.

In Table 1, we present the statistics of the AVSD dataset. Given the fact that the lengths of the I3D-RGB frame feature sequences are more than 20 times longer than the questions, using a recurrent neural network to encode the visual feature sequences will be very time consuming, as the visual frames are processed sequentially. Our proposed question-guided video representation module summarizes the video sequence efficiently - aggregating the visual features by question-guided attention and weighted summation and performing gating with a question-guided gate vector, both of which can be done in parallel across all frames.

4.2 Experimental Setup

We implement our models using the Tensor2Tensor framework (Vaswani et al., 2018). The question and dialogue context tokens are both embedded with the same randomly-initialized word embedding matrix, which is also shared with the answer decoder’s output embedding. The dimension of the word embedding is 256, the same dimension to which the I3D-RGB features are transformed. All of our LSTM encoders and decoder

have 1 hidden layer. Bahdanau attention mechanism (Bahdanau et al., 2015) is used in the answer decoder. During training, we apply dropout rate 0.2 in the encoder and decoder cells. We use the ADAM optimizer (Kingma and Ba, 2015) with $\alpha = 2 \times 10^{-4}$, $\beta_1 = 0.85$, $\beta_2 = 0.997$, $\epsilon = 10^{-6}$, and clip the gradient with L2 norm threshold 2.0 (Pascanu et al., 2013). The models are trained up to 100K steps with early stopping on the validation BLEU-4 score using batch size 1024 on a single GPU. During inference, we use beam search decoding with beam width 3. We experimented with word embedding dimension $\{256, 512\}$, dropout rate $\{0, 0.2\}$, Luong and Bahdanau attention mechanisms, $\{1, 2\}$ hidden layer(s) for both encoders and the decoder. We found the aforementioned setting worked best for most models.

5 Results

5.1 Comparison with Existing Methods

We evaluate our proposed approach using the same natural language generation evaluation toolkit NLG Eval (Sharma et al., 2017) as the previous approaches. The corpus-wide scores of the following unsupervised automated metrics are reported, including BLEU-1 through BLEU-4 (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), ROUGE-L (Lin and Och, 2004) and CIDEr (Vedantam et al., 2015). The results of our models in comparison with the previous approaches are shown in Table 2. We report the mean and standard deviation scores of 5 runs using random initialization and early stopping on the public (prototype) validation set. We apply our model in two scenarios: single-turn and multi-turn VideoQA. The only difference is that in single-turn VideoQA, the dialogue context encoder is excluded from the model.

First we observe that our proposed multi-turn VideoQA model significantly outperforms the single-turn VideoQA model. This suggests that the additional dialogue context input can provide supplementary information from the question and visual features, and thus is helpful for generating the correct answer. Secondly, comparing the single-turn VideoQA models, our approach outperforms the existing approaches across all automatic evaluation metrics. This suggests the effectiveness of our proposed question-guided video representations for VideoQA. When comparing

Single-Turn VideoQA Models	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE-L	CIDEr
Naïve Fusion	27.7	17.5	11.8	8.3	11.7	28.8	74.0
Multi-source Seq2Seq	-	-	-	8.83	12.43	34.23	95.54
Ours	29.56±0.75	18.60±0.49	13.16±0.33	9.77±0.21	13.19±0.20	34.29±0.19	101.75±1.03
Multi-Turn VideoQA Models	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE-L	CIDEr
Naïve Fusion	27.7	17.6	12.0	8.5	11.8	29.0	76.5
Attentional Fusion	27.6	17.7	12.2	8.7	11.7	29.3	78.7
Modified Attn. Fusion	27.7	17.6	12.0	8.5	11.8	29.0	76.5
+MMI objective	28.3	18.1	12.4	8.9	12.1	29.6	80.5
Hierarchical Attention	29.1	18.6	12.6	9.0	12.7	30.1	82.4
+pre-trained embedding	30.7	20.4	14.4	10.6	13.6	32.0	99.5
Multi-Source Seq2Seq	-	-	-	10.58	14.13	36.54	105.39
Ours	30.52±0.34	20.00±0.20	14.46±0.14	10.93±0.11	13.87±0.10	36.62±0.23	113.28±1.37

Table 2: Comparison with existing approaches: Naïve Fusion (Alamri et al., 2019b; Zhuang et al., 2019), Attentional Fusion (Hori et al., 2018; Zhuang et al., 2019), Multi-Source Sequence-to-Sequence model (Pasunuru and Bansal, 2019), Modified Attentional Fusion with Maximum Mutual Information objective (Zhuang et al., 2019) and Hierarchical Attention with pre-trained embedding (Le et al., 2019), on the AVSD public test set. For each approach, we report its corpus-wide scores on BLEU-1 through BLEU-4, METEOR, ROUGE-L and CIDEr. We report the mean and standard deviation scores of 5 runs using random initialization and early stopping on the public (prototype) validation set.

Model	BLEU-4	METEOR	ROUGE-L	CIDEr
Ours	10.94	13.73	36.30	111.12
-TokSumm	10.46	13.49	35.81	110.08
-Gating	10.59	13.64	36.11	108.51
-TokSumm-Gating	10.06	13.20	35.35	104.01

Table 3: Ablation study on the AVSD validation set. We observe that the performance degrades when either of both of the question-guided per-token visual feature summarization (TokSumm) and feature gating (Gating) techniques are removed.

with previous multi-turn VideoQA models, our approach that uses the dialogue context (questions and answers in previous turns) yields state-of-the-art performance on the BLEU-3, BLEU-4, ROUGE-L and CIDEr metrics and competitive results on BLEU-1, BLEU-2 and METEOR. It is worth mentioning that our model does not use pre-trained word embedding or audio features as in the previous hierarchical attention approach (Le et al., 2019).

5.2 Ablation Study and Weights Visualization

We perform ablation experiments on the validation set in the multi-turn VideoQA scenario to analyze the effectiveness of the two techniques in the question-guided video representation module. The results are shown in Table 3.

5.2.1 Question-Guided Per-Token Visual Feature Summarization (TokSumm)

Instead of using token-level question representations $\mathbf{x}^{\text{tok}} = \{x_1^{\text{tok}}, \dots, x_K^{\text{tok}}\}$ to generate per-token video summary $\mathbf{v} = \{v_1, \dots, v_K\}$, we experiment with using the sentence-level representation of the question x^{sen} as the query vector to attend over the I3D-RGB visual features to create a visual summary v , and use v to augment each of the question tokens in the video-augmented question encoder.

$$s_l = \text{trilinear}(x^{\text{sen}}, r_l) \quad (30)$$

$$\forall l \in \{1, \dots, L\}$$

$$\mathbf{w}^{\text{att}} = \text{softmax}(\mathbf{s}) \quad (31)$$

$$v = \sum_{l=1}^L w_l^{\text{att}} r_l \quad (32)$$

We observe the performance degrades when the sentence-level video summary is used instead of the token-level video summary.

Figure 3 shows an example of the attention weights in the question-guided per-token visual feature summarization. We can see that for different question tokens, the attention weights are shifted to focus on the different segment in the sequence of the video frame features.

5.2.2 Question-Guided Visual Feature Gating (Gating)

We also experiment with using the non-gated token-level video summary $\mathbf{v} = \{v_1, \dots, v_K\}$ to

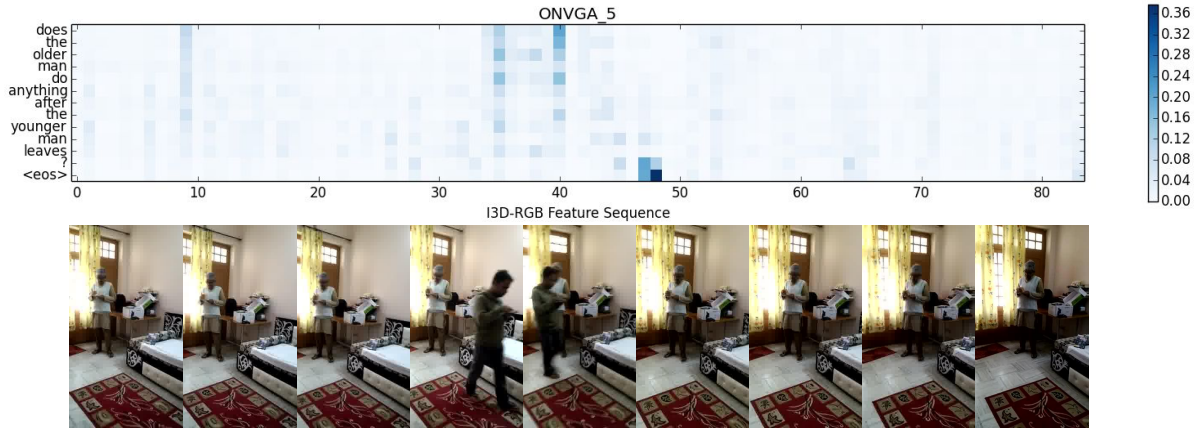


Figure 3: Question-guided per-token visual feature summary weights on a question. Each row represents the attention weights w_k^{att} of the corresponding encoded question token x_k^{tok} over the I3D-RGB visual features. We can observe that the attention weights are shifted to focus on the relevant segment of the visual frame features for the question tokens “after the younger man leaves <eos>?”




Figure 4: Question-guided gate weights g for some example questions. Across the questions about similar subjects, we observe a similar trend of weight distribution over visual feature dimensions. Conversely, questions about different topics show different gate weights patterns.

augment the question information in the video-augmented question encoder. We observe the model’s performance declines when the question-guided gating is not applied on the video summary feature. Removing both the per-token visual feature summarization and the gating mechanism results in further degradation in the model performance.

Figure 4 illustrates the question-guided gate weights g of several example questions. We observe that the gate vectors corresponding to the questions regarding similar subjects assign weights on similar dimensions of the visual feature. Although many of the visual feature dimensions have low weights across different questions, the feature dimensions of higher gate weights still exhibit certain topic-specific patterns.

6 Conclusion and Future Work

In this paper, we present an end-to-end trainable model for single-turn and multi-turn VideoQA.

Our proposed framework takes the question, I3D-RGB video frame features and dialogue context as input. Using the question information as guidance, the video features are summarized as compact representations to augment the question information, which are jointly used with the dialogue context to generate a natural language answer to the question. Specifically, our proposed question-guided video representation module is able to summarize the video features efficiently for each question token using an attention mechanism and perform feature selection through a gating mechanism. In empirical evaluation, our proposed models for single-turn and multi-turn VideoQA outperform existing approaches on several automatic natural language generation evaluation metrics. Detailed analyses are performed, and it is shown that our model effectively attends to relevant frames in the video feature sequence for summarization, and the gating mechanism shows topic-specific patterns in the feature dimension selection within a frame. In future work, we plan

to extend the models to incorporate audio features and experiment with more advanced techniques to incorporate the dialogue context with the question and video information, such as hierarchical attention and co-attention mechanisms. We also plan to employ our model on TVQA, a larger scale VideoQA dataset.

References

- Aishwarya Agrawal, Jiasen Lu, Stanislaw Antol, Margaret Mitchell, C Lawrence Zitnick, Devi Parikh, and Dhruv Batra. 2017. VQA: Visual question answering. *International Journal of Computer Vision (IJCV)*.
- Huda Alamri, Vincent Cartillier, Abhishek Das, Jue Wang, Stefan Lee, Peter Anderson, Irfan Essa, Devi Parikh, Dhruv Batra, Anoop Cherian, Tim K. Marks, and Chiori Hori. 2019a. Audio visual scene-aware dialog. In *Computer Vision and Pattern Recognition (CVPR)*.
- Huda Alamri, Chiori Hori, Tim K. Marks, Dhruv Batra, and Devi Parikh. 2019b. Audio visual scene-aware dialog (avsd) track for natural language generation in dstc7. In *DSTC7 at AAI 2019 Workshop*.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Computer Vision and Pattern Recognition (CVPR)*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations (ICLR)*.
- Lalit R Bahl, Peter F Brown, Peter V De Souza, and Robert L Mercer. 1986. Maximum mutual information estimation of hidden markov model parameters for speech recognition. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *ACL workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*.
- Joao Carreira and Andrew Zisserman. 2017. Quo vadis, action recognition? a new model and the Kinetics dataset. In *Computer Vision and Pattern Recognition (CVPR)*.
- Shu Zhang Wensheng Wang Chi Zhang Heng Huang Chenyou Fan, Xiaofan Zhang. 2019. Heterogeneous memory enhanced multimodal attention model for video question answering. In *Computer Vision and Pattern Recognition (CVPR)*.
- Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José MF Moura, Devi Parikh, and Dhruv Batra. 2017. Visual dialog. In *Computer Vision and Pattern Recognition (CVPR)*.
- Orhan Firat, Kyunghyun Cho, and Yoshua Bengio. 2016. Multi-way, multilingual neural machine translation with a shared attention mechanism. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Jiyang Gao, Runzhou Ge, Kan Chen, and Ram Nevatia. 2018. Motion-appearance co-memory networks for video question answering. In *Computer Vision and Pattern Recognition (CVPR)*.
- Alex Graves, Greg Wayne, Malcolm Reynolds, Tim Harley, Ivo Danihelka, Agnieszka Grabska-Barwińska, Sergio Gómez Colmenarejo, Edward Grefenstette, Tiago Ramalho, John Agapiou, et al. 2016. Hybrid computing using a neural network with dynamic external memory. *Nature*.
- Chiori Hori, Huda Alamri, Jue Wang, Gordon Winchern, Takaaki Hori, Anoop Cherian, Tim K Marks, Vincent Cartillier, Raphael Gontijo Lopes, Abhishek Das, et al. 2018. [End-to-end audio visual scene-aware dialog using multimodal attention-based video features](#). *Computing Research Repository*, arXiv:1806.08409.
- Yunseok Jang, Yale Song, Youngjae Yu, Youngjin Kim, and Gunhee Kim. 2017. TGIF-QA: Toward spatio-temporal reasoning in visual question answering. In *Computer Vision and Pattern Recognition (CVPR)*.
- Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. 2017. [The kinetics human action video dataset](#). *Computing Research Repository*, arXiv:1705.06950.
- Vahid Kazemi and Ali Elqursh. 2017. [Show, ask, attend, and answer: A strong baseline for visual question answering](#). *Computing Research Repository*, arXiv:1704.03162.
- Kyung-Min Kim, Min-Oh Heo, Seong-Ho Choi, and Byoung-Tak Zhang. 2017. DeepStory: video story QA by deep embedded memory networks. In *International Joint Conferences on Artificial Intelligence (IJCAI)*.
- D Kingma and Jimmy Ba. 2015. A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2017. Visual Genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision (IJCV)*.

- Ankit Kumar, Ozan Irsoy, Peter Ondruska, Mohit Iyyer, James Bradbury, Ishaan Gulrajani, Victor Zhong, Romain Paulus, and Richard Socher. 2016. Ask me anything: Dynamic memory networks for natural language processing. In *International Conference on Machine Learning (ICML)*.
- Shachi H Kumar, Eda Okur, Saurav Sahay, Juan Jose Alvarado Leanos, Jonathan Huang, and Lama Nachman. 2019. Context, attention and audio feature explorations for audio visual scene-aware dialog. In *DSTC7 at AAI 2019 workshop*.
- Hung Le, S Hoi, Doyen Sahoo, and N Chen. 2019. End-to-end multimodal dialog systems with hierarchical multimodal attention on video features. In *DSTC7 at AAI 2019 workshop*.
- Jie Lei, Licheng Yu, Mohit Bansal, and Tamara Berg. 2018. TVQA: Localized, compositional video question answering. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Jindřich Libovický and Jindřich Helcl. 2017. Attention strategies for multi-source sequence-to-sequence learning. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Chin-Yew Lin and Franz Josef Och. 2004. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Kuan-Yen Lin, Chao-Chun Hsu, Yun-Nung Chen, and Lun-Wei Ku. 2019. Entropy-enhanced multimodal attention model for scene-aware dialogue generation. In *DSTC7 at AAI 2019 workshop*.
- Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. 2016. Hierarchical question-image co-attention for visual question answering. In *Advances in Neural Information Processing Systems*.
- Chao Ma, Chunhua Shen, Anthony Dick, Qi Wu, Peng Wang, Anton van den Hengel, and Ian Reid. 2018. Visual question answering with memory-augmented networks. In *Computer Vision and Pattern Recognition (CVPR)*.
- Jonghwan Mun, Paul Hongsuck Seo, Ilchae Jung, and Bohyung Han. 2017. MarioQA: Answering questions by watching gameplay videos. In *International Conference on Computer Vision (ICCV)*.
- Seil Na, Sangho Lee, Jisung Kim, and Gunhee Kim. 2017. A read-write memory network for movie story understanding. In *International Conference on Computer Vision (ICCV)*.
- Dat Tien Nguyen, Shikhar Sharma, Hannes Schulz, and Layla El Asri. 2019. From film to video: Multi-turn question answering with multi-modal context. In *DSTC7 at AAI 2019 workshop*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: A method for automatic evaluation of machine translation. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2013. On the difficulty of training recurrent neural networks. In *International Conference on Machine Learning (ICML)*.
- Ramakanth Pasunuru and Mohit Bansal. 2019. Dstc7-avsd: Scene-aware video-dialogue systems with dual attention. In *DSTC7 at AAI 2019 workshop*.
- Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. 2018. FiLM: Visual reasoning with a general conditioning layer. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- Mengye Ren, Ryan Kiros, and Richard Zemel. 2015a. Exploring models and data for image question answering. In *Advances in Neural Information Processing Systems*.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015b. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*.
- Ramon Sanabria, Shruti Palaskar, and Florian Metzger. 2019. Cmu sinbads submission for the dstc7 avsd challenge. In *DSTC7 at AAI 2019 workshop*.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. In *International Conference on Learning Representations (ICLR)*.
- Shikhar Sharma, Layla El Asri, Hannes Schulz, and Jeremie Zumer. 2017. [Relevance of unsupervised metrics in task-oriented dialogue for evaluating natural language generation](#). *Computing Research Repository*, arXiv:1706.09799.
- Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. 2016. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *European Conference on Computer Vision (ECCV)*.
- Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. 2015. End-to-end memory networks. In *Advances in Neural Information Processing Systems*.
- Makarand Tapaswi, Yukun Zhu, Rainer Stiefelhagen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. 2016. MovieQA: Understanding stories in movies through question-answering. In *Computer Vision and Pattern Recognition (CVPR)*.
- Ashish Vaswani, Samy Bengio, Eugene Brevdo, Francois Chollet, Aidan Gomez, Stephan Gouws, Llion Jones, Łukasz Kaiser, Nal Kalchbrenner, Niki Parmar, et al. 2018. Tensor2Tensor for neural machine translation. In *Conference of the Association for Machine Translation in the Americas*.

- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. CIDEr: Consensus-based image description evaluation. In *Computer Vision and Pattern Recognition (CVPR)*.
- Caiming Xiong, Stephen Merity, and Richard Socher. 2016. Dynamic memory networks for visual and textual question answering. In *International Conference on Machine Learning (ICML)*.
- Dejing Xu, Zhou Zhao, Jun Xiao, Fei Wu, Hanwang Zhang, Xiangnan He, and Yueting Zhuang. 2017. Video question answering via gradually refined attention over appearance and motion. In *International Conference on Multimedia*.
- Huijuan Xu and Kate Saenko. 2016. Ask, attend and answer: Exploring question-guided spatial attention for visual question answering. In *European Conference on Computer Vision (ECCV)*.
- Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. 2016. Stacked attention networks for image question answering. In *Computer Vision and Pattern Recognition (CVPR)*.
- Yunan Ye, Zhou Zhao, Yimeng Li, Long Chen, Jun Xiao, and Yueting Zhuang. 2017. Video question answering via attribute-augmented attention network learning. In *SIGIR Conference on Research and Development in Information Retrieval*.
- Yi-Ting Yeh, Tzu-Chuan Lin, Hsiao-Hua Cheng, Yi-Hsuan Deng, Shang-Yu Su, and Yun-Nung Chen. 2019. Reactive multi-stage feature fusion for multimodal dialogue modeling. In *DSTC7 at AACL 2019 Workshop*.
- Licheng Yu, Eunbyung Park, Alexander C Berg, and Tamara L Berg. 2015. Visual Madlibs: Fill in the blank description generation and question answering. In *International Conference on Computer Vision (ICCV)*.
- Kuo-Hao Zeng, Tseng-Hung Chen, Ching-Yao Chuang, Yuan-Hong Liao, Juan Carlos Niebles, and Min Sun. 2017. Leveraging video descriptions to learn video question answering. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- Zhou Zhao, Zhu Zhang, Shuwen Xiao, Zhou Yu, Jun Yu, Deng Cai, Fei Wu, and Yueting Zhuang. 2018. Open-ended long-form video question answering via adaptive hierarchical reinforced networks. In *International Joint Conferences on Artificial Intelligence (IJCAI)*.
- Bairong Zhuang, Wenbo Wang, and Takahiro Shinohara. 2019. Investigation of attention-based multimodal fusion and maximum mutual information objective for dstc7 track3. In *DSTC7 at AACL 2019 workshop*.
- Barret Zoph and Kevin Knight. 2016. Multi-source neural translation. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.

Zero-shot transfer for implicit discourse relation classification

Murathan Kurfali

Department of Linguistics
Stockholm University
murathan.kurfali@ling.su.se

Robert Östling

Department of Linguistics
Stockholm University
robert@ling.su.se

Abstract

Automatically classifying the relation between sentences in a discourse is a challenging task, in particular when there is no overt expression of the relation. It becomes even more challenging by the fact that annotated training data exists only for a small number of languages, such as English and Chinese. We present a new system using zero-shot transfer learning for implicit discourse relation classification, where the only resource used for the target language is unannotated parallel text. This system is evaluated on the discourse-annotated TED-MDB parallel corpus, where it obtains good results for all seven languages using only English training data.

1 Introduction

The difference between a set of randomly selected sentences and a discourse lies in coherence. Among other attempts at defining the elusive nature of coherence, one way is to look at the meaning conveyed between the adjacent pair of sentences. In the current study, we follow the Penn Discourse Treebank (PDTB) framework which regards abstract objects (Asher, 2012) as the units of discourse and views the text as a collection of discourse level predicates, each taking two arguments. Such predicates, called discourse connectives, may (Ex. 1) or may not (Ex. 2) be represented in the surface form:

1. Because **the drought reduced U.S. stock-piles**, *they have more than enough storage space for their new crop*, and that permits them to wait for prices to rise.
2. *But a few funds have taken other defensive steps. Some have raised their cash positions to record levels.* Implicit = BECAUSE **High cash positions help buffer a fund when the market falls.**

where *italics* represents the first and **boldface** the second argument to the underlined discourse connective. The discourse relations which lack an overt discourse connective (Ex. 2) are referred as *implicit discourse relations* and are shown to be the most challenging part of the discourse parsing (e.g. Pitler et al., 2009).

In this paper, we perform implicit discourse relation classification using three recent data sets annotated according to the same guidelines: Penn Discourse Treebank (PDTB) 3.0, the Turkish Discourse Bank (TDB), and the multilingual TED-MDB. To the best of our knowledge, multilingual training and zero-shot transfer has not previously been investigated for this problem. The results suggest that an implicit discourse relation classifier can transfer well across dissimilar languages, and that pooling training data from unrelated languages (English and Turkish) leads to significantly better performance for all languages.

2 Related Work

Implicit discourse relation recognition is often handled as a classification task, where earlier studies focused on using linguistically rich features (Pitler et al., 2009; Zhou et al., 2010; Park and Cardie, 2012; Rutherford and Xue, 2014).

Recently, neural network approaches have become popular. Ji and Eisenstein (2015) use two RNNs on the syntactic trees of the arguments whereas Zhang et al. (2015) use a CNN to perform discourse parsing in a multi-task setting where they consider both explicit and implicit discourse relations.

Rutherford and Xue (2016) use a simple yet robust feedforward network and achieves the highest performance on the out-of-domain blind test in the CoNLL 2016 shared task (Xue et al., 2016).

Lan et al. (2017) apply a multi-task attention-

based neural network model whereas [Bai and Zhao \(2018\)](#) focus on the representation of the sentence pair and take different levels of text, from character to sentence pair, into account to achieve a richer representation.

[Dai and Huang \(2018\)](#) adopt a similar approach and represent discourse units by considering a wider paragraph-level context. The discourse unit representations are created by a Bi-LSTM which takes a sequence of discourse relations in a paragraph which enables capturing the inter-dependencies between discourse relations as well.

3 Data

We use four different data sets: the Penn Discourse Treebank (PDTB) version 2.0 ([Prasad et al., 2008](#)) and version 3.0 ([Prasad et al., 2018](#)), as well as the TED Multilingual Discourse Bank (TED-MDB, [Zeyrek et al. 2018](#)) and the Turkish Discourse Bank (TDB, [Zeyrek and Kurfalı 2017](#)).

The PDTB is built upon the 1 million word Wall Street Journal corpus and is the largest available resource for discourse relations. Most related work uses PDTB 2.0, so we include this for comparing our baseline to previous work.

The recently released PDTB 3.0 adopts a new annotation schema as well as an updated sense hierarchy. PDTB 3.0 includes the annotations of PDTB 2.0 updated according to the new annotation schema, as well as about 13 thousand new annotations, of which about 5K are implicit relations ([Prasad et al., 2018](#)). The distribution of the top level senses of the implicit discourse relations in both PDTB versions is provided in Table 1.

TED-MDB ([Zeyrek et al., 2018](#)) is the first parallel corpus annotated for discourse relations. It closely follows the PDTB 3.0 framework and includes the manual annotations of six TED talks in seven languages (English, Turkish, European Portuguese, Polish, German, Russian) aiming to allow crosslingual comparison of discourse relations¹. It has recently also been updated with Lithuanian ([Oleskeviciene et al., 2019](#)).

Despite the high number of languages covered by TED-MDB, the amount of annotated text per language is limited (see Table 2). Therefore, in the current study, we limit ourselves with the top level senses, namely Expansion, Contingency, Compar-

¹The TED-MDB annotations are available at: <https://github.com/MurathanKurfali/Ted-MDB-Annotations>

ison and Temporal. We only use TED-MDB for evaluation.

Among the TED-MDB languages other than English, only Turkish has another corpus annotated with PDTB 3.0 discourse annotations, namely the Turkish Discourse Bank (TDB). TDB is a multi-genre corpus of 40 000 words, considerably less than the PDTB (see Table 2), but it provides the only directly comparable baseline to assess the performance of zero-shot learning.

	PDTB2			PDTB3		
Sense	Train	Dev	Test	Train	Dev	Test
Comp.	1894	401	146	1828	404	153
Cont.	3281	628	276	5872	1159	527
Exp.	6792	1253	556	7939	1466	643
Temp.	665	93	68	1413	230	148

Table 1: Distribution of top level senses of the implicit discourse relations in PDTB 2.0 and PDTB 3.0 training, development and test sets: comp(arison), cont(ingency), exp(ansion), temp(oral).

4 Model

The main purpose of this study is to assess the performance of transfer learning on the implicit discourse relation classification task. To this end, we use a simple feedforward network fed with multilingual sentence embeddings following the finding of ([Rutherford et al., 2017](#)) which shows that simple discourse models with feedforward layers perform on par or better than those of with surface features or recurrent and convolutional architectures.

We follow the model of ([Rutherford and Xue, 2016](#)) due to its simplicity and robust nature even in the multilingual setting with different argument and discourse relation representations. We represent the arguments of the discourse relation via pre-trained LASER model ([Artetxe and Schwenk, 2018](#)). LASER is chosen as it is the current state-of-the-art model on several Natural Language Inference (NLI) transfer learning tasks, a sentence relation classification problem similar to discourse relation classification.

Given the argument vectors, V_{arg1} and V_{arg2} , the next step is to represent the discourse relation in a way that the interactions between them are captured. To this end, we model the discourse relation vector, V_{dr} , by performing the following pair-wise vector operations following the DisSent model of ([Nie et al., 2017](#)):

Language	Comparison	Contingency	Expansion	Temporal	Total
English	20 (10.31%)	52 (26.80%)	107 (55.15%)	15 (7.73%)	194 (100%)
German	13 (6.07%)	41 (19.16%)	148 (69.16%)	12 (5.61%)	214 (100%)
Lithuanian	26 (10.57%)	53 (21.54%)	154 (62.60%)	13 (5.28%)	246 (100%)
Polish	19 (9.74%)	28 (14.36%)	130 (66.67%)	18 (9.23%)	195 (100%)
Portuguese	23 (9.06%)	47 (18.50%)	169 (66.54%)	15 (5.91%)	254 (100%)
Russian	16 (7.24%)	31 (14.03%)	169 (76.47%)	5 (2.26%)	221 (100%)
Turkish	20 (9.90%)	29 (14.36%)	140 (69.31%)	13 (6.44%)	202 (100%)
TDB (training)	71 (10.94%)	142 (21.88%)	363 (55.93%)	73 (11.25%)	649 (100%)
TDB (dev)	11 (9.82%)	31 (27.68%)	49 (43.75%)	21 (18.75%)	112 (100%)

Table 2: Distribution of top level senses of the implicit discourse relations in the TED-MDB and TDB corpora. The numbers within the parenthesis indicate the ratio. Since there is no official training/dev split for TDB, we arbitrarily chose two sections with different genres for the development set.

$$V_{avg} = \frac{1}{2}(V_{arg1} + V_{arg2})$$

$$V_{sub} = V_{arg1} - V_{arg2}$$

$$V_{mul} = V_{arg1} * V_{arg2}$$

$$V_{dr} = [V_{arg1}, V_{arg2}, V_{avg}, V_{sub}, V_{mul}]$$

The resulting vector is further fed into a hidden layer h_t with d hidden units² to achieve a more abstract representation of the relation and finally the output o is calculated using the sigmoid function. This model is also essentially the same as was used by Artetxe and Schwenk (2018) for NLI transfer learning.

5 Experiments

We formulate the implicit relation classification as four "one vs other" binary classification task. We follow the conventional setting of the first study (Pitler et al., 2009) and split the PDTB 2.0 into training (sections 2-20), development (sections 0-1 and 23-24) and test sets (sections 21-22) to have directly comparable results with the previous work. However, following the PDTB's original distinction but unlike some previous work, we distinguish Entity-based relations from implicit relations. Each classifier is trained on an equal number of positive and negative instances where the negative instances are randomly selected in each epoch to have a better representation of the data during the training. This model is evaluated on the PDTB 2.0 test set to confirm whether our model performs adequately on same-language, same-domain data. These results are directly comparable to previous work.

²We use $d=100$ in the experiments

As TED-MDB is annotated according to the PDTB 3.0 framework, we train separate classifiers on PDTB 3.0 following the same convention as above. We test the trained models on all the implicit discourse relations in the TED-MDB corpus.

The PDTB framework allows annotations to be labelled with more than one label. In such cases we only keep the first label, in line with previous studies (among others Ji and Eisenstein, 2015; Rutherford et al., 2017).

The argument vectors are kept fixed during the training, and we do not update the parameters of the LASER model. We use cross-entropy loss, and AdaGrad as the optimizer. We evaluate using the model which achieved the highest F-score on the development set. As for the regularization, we use a dropout layer between the input and the hidden layer with a dropout probability of 0.3. All models are run 100 times to estimate the variance due to random initialization and stochastic training. All the models are implemented in PyTorch³.

6 Results and Discussion

Table 3 shows the same-language, same-domain performance of our system, in comparison to previous work. All figures refer to PDTB 2.0 test set F-score, when trained on the PDTB 2.0 training set, and are directly comparable. While our model does not achieve state-of-the-art performance in this setting, this experiment shows that it performs adequately for English, and provides a reasonable baseline for the zero-shot experiments presented in Tables 4 and 5. We also include a naive baseline system which always predicts TRUE and is evaluated on the respective (PDTB 2.0 or PDTB 3.0)

³<https://pytorch.org/>

	Comparison	Contingency	Expansion	Temporal
(Pitler et al., 2009)	21.96	47.13	-	16.76
(Zhou et al., 2010)	31.79	47.16	70.11	20.30
(Park and Cardie, 2012)	31.32	49.82	-	26.57
(Rutherford and Xue, 2014)	39.70	54.42	70.23	28.69
(Zhang et al., 2015)	33.22	52.04	69.59	30.54
(Ji and Eisenstein, 2015)	35.93	52.78	-	27.63
(Lan et al., 2017)	40.73	58.96	72.47	38.50
(Bai and Zhao, 2018)	47.85	54.47	70.60	36.87
(Dai and Huang, 2018)	46.79	57.09	70.41	45.61
Baseline	24.49	41.75	69.41	12.20
Our system	28.19 (± 0.83)	50.63 (± 1.00)	64.07 (± 1.90)	29.22 (± 2.53)

Table 3: Comparison of the F scores (%) of binary classifiers on PDTB 2.0 test set. Left out scores refer to the results where EntRel relations are also considered to be Expansion.

Language	Comparison	Contingency	Expansion	Temporal	Average
Baseline (PDTB 3.0)	18.84	52.75	60.83	18.28	37.67
PDTB 3.0	24.90 (± 0.87)	59.18 (± 0.72)	60.10 (± 1.32)	36.73 (± 1.45)	45.23
German	8.62 (± 1.61)	37.34 (± 1.43)	70.81 (± 3.16)	40.11 (± 4.32)	39.22
English	10.18 (± 3.31)	40.92 (± 1.80)	62.28 (± 2.16)	50.45 (± 5.26)	40.96
Lithuanian	23.50 (± 2.33)	34.64 (± 1.43)	62.35 (± 2.65)	36.78 (± 3.28)	39.32
Polish	16.50 (± 3.51)	29.19 (± 1.36)	60.32 (± 2.84)	44.17 (± 3.37)	37.54
Portuguese	19.59 (± 1.99)	33.85 (± 1.27)	66.83 (± 2.57)	37.04 (± 3.43)	39.33
Russian	14.90 (± 2.07)	26.76 (± 1.08)	70.06 (± 3.97)	28.28 (± 4.41)	35.00
Turkish	10.99 (± 3.16)	25.28 (± 1.23)	64.14 (± 2.96)	33.66 (± 4.31)	33.52

Table 4: F scores (%) when the model is trained only on PDTB 3.0. In the table, PDTB 3.0 refers to the test set of the PDTB 3.0 corpus. The remaining rows refer to evaluations using TED-MDB.

test set in our comparisons.

In all zero-shot experiments, evaluation is performed on the available test data with PDTB 3.0 annotations: TED-MDB, and the PDTB 3.0 test set itself. Results in Table 4 use PDTB 3.0 only for training, whereas Table 5 presents the effect of having additional training data from Turkish (a language unrelated to English). Pooling training data from different languages is possible since our model is language-agnostic.

In all zero-shot experiments, we see similar levels of performance across all the evaluated languages in TED-MDB. While not completely comparable numerically since annotations differ slightly between languages, this evaluation set consists of parallel sentences annotated according to the same guidelines. The similarity in scores between the training language(s)—English and/or Turkish—and the remaining languages indicates that little accuracy is lost during transfer.

Comparing the performances with and without additional Turkish data, TDB, reveals that adding

a small amount (relative to the size of PDTB 3.0) of Turkish training data improves the F-scores by a statistically significant amount⁴ for not only Turkish, but for all the languages in TED-MDB Table 5.

7 Conclusion

In the current paper we have presented the (to the best of our knowledge) first study of zero-shot learning in the implicit discourse relation classification task. Our method does not require any discourse level annotation for the target languages, yet still achieves good performance even for those languages where no training data is available. The performance is further increased by pooling training data from multiple languages. Using our published code⁵ and publicly available resources it can be used for implicit discourse classification in

⁴On a sense-wise analysis, we observe that the main increase is in the Expansion relations; however, there is no decrease in any of the other senses.

⁵https://github.com/MurathanKurfali/multilingual_IDRC

Language	TDB	PDTB3	PTDB3+TDB
PDTB3 Test	35.35	45.23	45.62
German	36.93	39.22	41.44
English	38.06	40.96	42.22
Lithuanian	36.92	39.32	41.94
Polish	35.48	37.54	39.65
Portuguese	37.58	39.33	41.04
Russian	30.92	35.00	38.23
Turkish	39.58	33.52	37.14

Table 5: Comparison of average F-scores (%) when the model is trained on different training sets. Bold means significantly higher F-score than the second highest column ($p < 0.001$, Mann-Whitney U test).

nearly a hundred languages.

Acknowledgments

We would like to thank Bonnie Webber for her help in obtaining PDTB 3.0 and Mats Wirén for his useful comments.

References

- Mikel Artetxe and Holger Schwenk. 2018. Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. *arXiv preprint arXiv:1812.10464*.
- Nicholas Asher. 2012. *Reference to abstract objects in discourse*, volume 50. Springer Science & Business Media.
- Hongxiao Bai and Hai Zhao. 2018. Deep enhanced representation for implicit discourse relation recognition. *arXiv preprint arXiv:1807.05154*.
- Zeyu Dai and Ruihong Huang. 2018. Improving implicit discourse relation classification by modeling inter-dependencies of discourse units in a paragraph. *arXiv preprint arXiv:1804.05918*.
- Yangfeng Ji and Jacob Eisenstein. 2015. One vector is not enough: Entity-augmented distributed semantics for discourse relations. *Transactions of the Association for Computational Linguistics*, 3:329–344.
- Man Lan, Jianxiang Wang, Yuanbin Wu, Zheng-Yu Niu, and Haifeng Wang. 2017. Multi-task attention-based neural networks for implicit discourse relationship representation and identification. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1299–1308.
- Allen Nie, Erin D Bennett, and Noah D Goodman. 2017. Dissent: Sentence representation learning from explicit discourse relations. *arXiv preprint arXiv:1710.04334*.
- Giedre Valunaite Oleskeviciene, Deniz Zeyrek, Viktorija Mazeikiene, and Murathan Kurfalı. 2019. Observations on the annotation of discourse relational devices in ted talk transcripts in lithuanian. *Proceedings of the workshop on annotation in digital humanities co-located with ESSLLI 2018*, pages 53–58.
- Joonsuk Park and Claire Cardie. 2012. Improving implicit discourse relation recognition through feature set optimization. In *Proceedings of the 13th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 108–112. Association for Computational Linguistics.
- Emily Pitler, Annie Louis, and Ani Nenkova. 2009. Automatic sense prediction for implicit discourse relations in text. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2-Volume 2*, pages 683–691. Association for Computational Linguistics.
- Rashmi Prasad, Nikhil Dinesh, Alan Lee, Eleni Miltasakaki, Livio Robaldo, Aravind K Joshi, and Bonnie L Webber. 2008. The penn discourse treebank 2.0. In *LREC*. Citeseer.
- Rashmi Prasad, Bonnie Webber, and Alan Lee. 2018. Discourse annotation in the pdtb: The next generation. In *Proceedings 14th Joint ACL-ISO Workshop on Interoperable Semantic Annotation*, pages 87–97.
- Attapol Rutherford, Vera Demberg, and Nianwen Xue. 2017. A systematic study of neural discourse models for implicit discourse relation. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 281–291.
- Attapol Rutherford and Nianwen Xue. 2014. Discovering implicit discourse relations through brown cluster pair representation and coreference patterns. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 645–654.
- Attapol Rutherford and Nianwen Xue. 2016. Robust non-explicit neural discourse parser in english and chinese. *Proceedings of the CoNLL-16 shared task*, pages 55–59.
- Nianwen Xue, Hwee Tou Ng, Sameer Pradhan, Attapol Rutherford, Bonnie Webber, Chuan Wang, and Hongmin Wang. 2016. Conll 2016 shared task on multilingual shallow discourse parsing. *Proceedings of the CoNLL-16 shared task*, pages 1–19.
- Deniz Zeyrek and Murathan Kurfalı. 2017. Tdb 1.1: Extensions on turkish discourse bank. In *Proceedings of the 11th Linguistic Annotation Workshop*, pages 76–81.

- Deniz Zeyrek, Amália Mendes, Yulia Grishina, Murathan Kurfalı, Samuel Gibbon, and Maciej Ogródniczuk. 2018. Ted multilingual discourse bank (tedmdb): a parallel corpus annotated in the pdtb style. *Language Resources and Evaluation*, pages 1–27.
- Biao Zhang, Jinsong Su, Deyi Xiong, Yaojie Lu, Hong Duan, and Junfeng Yao. 2015. Shallow convolutional neural network for implicit discourse relation recognition. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2230–2235.
- Zhi-Min Zhou, Yu Xu, Zheng-Yu Niu, Man Lan, Jian Su, and Chew Lim Tan. 2010. Predicting discourse connectives for implicit discourse relation recognition. In *Proceedings of the 23rd International Conference on Computational Linguistics: Posters*, pages 1507–1514. Association for Computational Linguistics.

A Quantitative Analysis of Patients' Narratives of Heart Failure

Sabita Acharya, Barbara Di Eugenio, Andrew Boyd, Richard Cameron,
Karen Dunn Lopez, Pamela Martyn-Nemeth, Debaleena Chattopadhyay,
Pantea Habibi, Carolyn Dickens, Haleh Vatani, Amer Ardati

University of Illinois at Chicago
Chicago, Illinois

sachar4, bdieugen, boyda, rcameron, kdunn12, pmartyn,
debchatt, phabib4, cdickens, hvatan2, aardati@uic.edu

Abstract

Patients with chronic conditions like heart failure are most likely to be re-hospitalized. One step towards avoiding re-hospitalization is to devise strategies for motivating patients to take care of their own health. In this paper, we perform a quantitative analysis of patients' narratives of their experience with heart failure and explore the different topics that patients talk about. We compare two different groups of patients- those unable to take charge of their illness, and those who make efforts to improve their health. We will use the findings from our analysis to refine and personalize the summaries of hospitalizations that our system automatically generates.

1 Introduction

Patients with heart failure are responsible for around 95% of their chronic illness care and their daily decisions have a huge impact on their quality of life (Funnell, 2000). Studies have shown that the patients' perspective is essential for patient education (Shapiro, 1993) and that engaging the patients in their own care reduces hospitalizations and prevents further deterioration of their health (Riegel et al., 2011; McGinnis et al., 2013).

We are engaged in a large, long-term project that aims to improve patient discharge instructions with a personalized and comprehensible summary of their hospital stay that is informed by the perspectives of the three main stake-holders: doctors, nurses, and patients. Over the last few years, we have developed and implemented a framework for summarizing heterogeneous information (textual discharge notes from the doctor, structured information from the nurses) and providing explanations for difficult medical terms (Di Eugenio et al., 2014; Acharya et al., 2016, 2018, 2019). Figure 1 shows a part of a summary that is generated by our system.

You were admitted for acute subcortical cerebrovascular accident. During your hospitalization, you were monitored for chances of ineffective cerebral tissue perfusion, risk for falls, problem in verbal communication and walking. We treated difficulty walking related to nervous system disorder with body mechanics promotion. [...] As a result, fall prevention behavior and [...] improved slightly. With your nurse and doctors, you learned about disease process and medication. Follow-up: Can follow-up with General Neurology clinic and Medicine clinic as outpatient if desired.

Figure 1: Portion of a summary generated by our system. Underlined terms provide a lay language definition when clicked.

A high level flowchart of the algorithm is shown in Appendix A. At this point, most inputs in pink are available for and used by the algorithm other than *strengths and concerns* of the patient (uncovering which is part of the focus of this paper). Specifically, the current version of our system uses the following personalization features:

A) **Participation in self-care:** The Patient Activation Measure (PAM) (Hibbard et al., 2004) quantifies how motivated patients are in taking care of their health. Based on the responses to a set of 13 questions, PAM assigns a level between 1 and 4. Level 1 indicates that the patient is overwhelmed, while level 4 indicates that the patient is motivated to participate in self-care.

B) **Familiarity with the health issue :** This feature takes into account different factors that contribute towards a patient's understanding of their health - i) The health literacy of the patient, which represents the ability of a patient to read and understand general health information, as measured by the REALM metric (Davis et al., 1993) ; ii) Patient's prior experience - either because of their own sufferings or because someone in their family had the same issue; iii) Patient's self-assessment of their health knowledge, as obtained from some of the questions in the PAM.

While the developers of the PAM provide some instructions on how to address patients at different PAM levels¹, no systematic study of these patients and their views of their illness exist. Under the assumption that such a study can reveal features useful to personalize our summaries, we conducted interviews with 26 patients, who were also asked to answer the PAM questionnaire. In this paper, we describe the quantitative analyses that we have performed on those patient interviews. We start with differences in the terms that those patients use for recounting their health experiences. We found that in spite of using fewer medical terms, patients with high PAM levels speak a higher proportion of unique medical terms. Our analysis of the patients’ use of pronouns suggests that patients with low PAM levels tend to self-focus, which is associated with negative effects and low self-confidence (Duval and Wicklund, 1972; Pyszczynski and Greenberg, 1987). Finally, we discuss themes that emerge from those conversations, with the goal of highlighting aspects that hold significance in the patients’ lives: for example, patients with high PAM focus more on activities they are interested in, patients with low PAM on their own feelings (confirming the finding about pronouns just discussed).

2 Related Work

While several systems exist that summarize medical content (Scott et al., 2013; Pauws et al., 2019), only a few of them produce personalized summaries (Mahamood and Reiter, 2011). Unlike these systems that focus on data-to-text summarization, our personalized summary generation system combines the information from physician and nursing documents and provides hospitalization information to patients in a form that they can understand. Even though a lot of studies have focused on verifying the reliability of the PAM metric (Fowles et al., 2009), no work uses it to produce personalized content for patients. Most of the existing qualitative studies on the narratives of heart failure patients (Jeon et al., 2010; Seah et al., 2016) focus on identifying the factors that impact the patient’s self-care and self-management skills. However, none of these studies has looked into the relationship between the content spoken by the pa-

¹ <https://participatorymedicine.org/epatients/2011/10/the-patient-activation-measure-pam-a-framework-for-developing-patient-engagement.html>

tients and their motivation to participate in self-care. Our quantitative analyses are inspired by Pennebaker (2003) and are similar to the studies that predict the empathy of the counselor based on the words used during the session (Althoff et al., 2016; Pérez-Rosas et al., 2017; Xiao et al., 2014).

3 Interview Collection

Category	Values
Avg. number of words in an interview(P)	1655
Avg. number of words in an interview (I)	1104
Avg. number of words/utterance (P)	8
Avg. number of words/utterance (I)	6
Number of low PAM patients	14
Number of high PAM patients	12

Table 1: Distributional analysis of the interviews (P: Patient, I: Interviewer)

Since there are no existing publicly available data sets that provide information on the experiences of heart-failure patients, we proceeded to collect one.² We interviewed 26 patients (age range 20-70 years, 58% females) who were hospitalized because of heart issues (snippets of an interview can be found in Appendix A). These 50 minutes long open-ended interview sessions were led by a sociolinguist and were later transcribed by professional transcribers. In general, each interview consists of the following stages: 1) Patients are asked to provide their demographic information; 2) Patients are asked to recount their first experience with heart issues, which often leads to them talking about many other issues related to life-style or family; 3) Patients are asked about the recent hospitalization and their experiences; 4) Patients answer the PAM questions; 5) If interested, patients talk about their interests or have a general conversation with the interviewer. For our analyses, we group the patients with PAM level 1 or 2 and refer to them as *low* PAM patients, while we refer to the group of patients with PAM levels 3 or 4 as *high* PAM patients. The general statistics on the interviews is shown in Table 1.

4 Distinguishing Low from High PAM Patients

We extracted several features from the transcripts, including the counts of different part-of-speech

²The data is not sharable because of human subject protection constraints, especially as dictated by HIPAA (the Health Insurance Portability and Accountability Act of 1996), United States legislation regarding the safeguard of private health care information.

tags, total positive and negative words, and medical and non-medical type-token ratio (TTR) (i.e. ratio of the number of unique words to the number of words). We then used a Random Forest based approach³ for determining the importance of each feature in predicting the PAM level. This process identified some significant features, which are further analyzed in Section 4.1 and Section 4.2.

4.1 Usage of medical terms

We extracted the average number of words spoken by the group of *low* and *high* PAM patients separately, including the number of medical terms (extracted using cTAKES tool (Savova et al., 2010)), TTR, and medical TTR. We found that patients with *low* PAM speak more but have lower value for TTR. Similarly, patients with *low* PAM use more medical terms, which account for 8% of their words; for patients with *high* PAM, medical terms constitute 6% of their total words, but they were found to have higher medical TTR. Although none of these differences is statistically significant, they still suggest that there is a difference in the lexical diversity (both general and medical) of the two groups of patients.

4.2 Patient outlook

4.2.1 Reference to self

Researchers on human psychology mention that when individuals start to focus their attention on themselves, they step into a self-evaluative process, where they compare their present to where they aspire to be. For those cases where the present lags behind the aspired standard, self-focus produces a negative effect (Duval and Wicklund, 1972; Pyszczynski and Greenberg, 1987). Similarly, the PAM metric characterizes a patient with *low* PAM score as an individual who is overwhelmed and weighed down by negative emotions. On the other hand, patients with *high* PAM are ready to take on challenges and make efforts to improve their health. Hence, in order to verify whether patients with *low* PAM focus more on themselves, we compared the relative amount of first person singular pronoun vs second and third person pronouns (both singular and plural) used by *low* and *high* PAM patients. We split each patient transcript into five parts and observe the trend in the reference to self. As seen in Figure 2, a greater

amount of self-focus is indeed associated with patients with *low* PAM. The differences across the five parts are statistically significant with a sign test (z -value= 2.23607, $p = .02535$).




Figure 2: Relative use of first person (singular) vs second and third person pronouns by the patients.

4.2.2 Sentiment of the patient




Figure 3: Relative fraction of positive sentences used by patients.

In order to determine how the sentiment of the patients change throughout the conversation, we used the VADER⁴ tool for performing sentiment analysis. Figure 3 shows the relative fraction of positive sentences that are spoken by patients. Interestingly, we can see that the curve for *high* PAM patients drops during 20-40% of the conversation, while the plot for *low* PAM patients drops during 40-60% of the conversation. One reason behind this is because at around 20-40% of the conversation, patients are asked to describe their first encounter with heart issues, while at around 40-60% of the conversation, patients are asked about their current reason behind hospitalization (as was mentioned in Section 3). We can also see that *high* PAM patients are fairly constant as concerns the fraction of positive content spoken after the first 20-40% of conversation, while the curve for *low* PAM patients has more rises and falls. This further supports the observation made by the developers of PAM that *low* PAM patients are overwhelmed.

³<https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.ExtraTreesClassifier.html>

⁴<https://github.com/cjhutto/vaderSentiment>

4.3 Topics Discussed by Patients and Their Relation to PAM Level

In order to identify the possible themes that arise from the conversations, we extracted only the nouns and adjectives from the patient utterances that occur in at least 10% of the transcripts and created a document matrix. We then performed Principal Component Analysis (PCA), along with Varimax rotation⁵ on the document matrix. We opted for 9 components (also known as factor loadings) that were able to explain around 75% of the total variance in the document matrix. Similar to (Wilson et al., 2016), we consider any word with a factor loading of at least 0.2 for a particular component as a positive contributor and words whose factor loadings are less than -0.2 as negative contributors. For each component, we calculate the normalized count of the words from each document that are positive contributors minus the number of negative contributors, and find the averages for *low* and *high* PAM patients. Finally, we calculate a score for each component, which is the ratio of average normalized count for *high* PAM to *low* PAM patients. A score > 1 indicates the prevalence of the category in *high* PAM, while a score < 1 indicates its prevalence in *low* PAM patients.

Topic	Sample words	Score
Activities of interest	ball, game, park, love, physical, recipe, swimming	2.59
Technical medical terms	lasix, murmur, supplement, admission, sign, diet, specialist	2.55
Family and support	husband, family, kind, everybody, father, parent, support	1.55
Life experiences	plan, jump, vacation, swimming, talk, breath, experience	1.02
Family and beliefs	niece, church, grandkid, grandma, honest, truth, folk	0.95
Life experiences	shower, shop, sugar, experience, contact, downtown, longtime	0.78
Feelings	terrible, difficult, horrible, difference, teaching, dizzy, force	0.77
Health	muscle, workout, lunch, healthy, vegetable, chicken, information	0.62
Food	bake, turkey, potato, vegetable, hot, green, meat, taste	0.21

Table 2: Sample words that reflect the 9 categories/themes, along with the topic scores.

Table 2 shows the sample words that represent the 9 categories/themes and the corresponding scores. The topics in the first column of the table are manual interpretations of what the sam-

⁵Varimax rotation causes the weights in the principal components to be closely associated to only one component, which makes it easier to interpret the results of PCA.

ple words refer to. The topics with scores greater than 1 are prominent for *high* PAM patients, while the ones with scores less than one are prominent for *low* PAM patients. Some interesting observations can be made from this table. First, *high* PAM patients seem to make more use of technical medical terms, which complements our finding in Section 4.1 that they are less repetitive in their usage of medical terms. Second, patients with *low* PAM seem to talk about their feelings, most of which relate to the negative effects of their health conditions, while *high* PAM patients talk about the activities they are interested in. This supports our findings from Section 4.2.1 and Section 4.2.2, which showed that patients with *low* PAM focus more on the negative changes in their lives.

5 Conclusion and Future Work

In this paper, we presented a quantitative study on the interviews with heart failure patients we collected. We analyzed the difference between patients with *low* and *high* PAM levels based on their reference to self, usage of medical terms, and the change in their sentiment through the conversation. We also identified the key topics that the patients from both groups talk about. The findings from these analyses have provided additional insights into the characteristics of heart failure patients and will be used for tuning different aspects of our personalized summary generation system.

Incorporating personalization features: Currently, our personalization algorithm provides different levels of details to patients depending upon their familiarity with their health issues. Similarly, for patients with a *low* PAM level, we show empathy with sentences like “Dealing with this issue must have been tough for you”, while *high* PAM patients are provided encouragement with sentences like “Keep up the good work”.

From the analyses of the usage of medical terms (Section 4.1) and the topics discussed by patients (Section 4.3), we found that the PAM level of a patient is also an indicator of the type and amount of medical terms that patients use while recounting their health experience. This suggests that in addition to health literacy (as we do currently), the PAM level should also be taken into account for deciding on the details that will be provided to the patient. Based on the findings in Section 4.2.1 and Section 4.2.2, we plan to divert *low* PAM patients from self-focus and its potential negative ef-

fects, for example by focusing more on positive outcomes and improvements in their health status. We also plan to use some of the topics that were discovered in Section 4.3 as multiple choice questions that will be shown to patients in real time. Based on the values that are selected, some generic sentences that motivate the patients to get better will be included in the summary.

Acknowledgments

This work was supported by award R01 CA225446-01 from the National Cancer Institute.

References

- Sabita Acharya, Andrew Boyd, Richard Cameron, Karen Dunn Lopez, Pamela Martyn-Nemeth, Carolyn Dickens, Amer Ardati, and Barbara Di Eugenio. 2019. Incorporating personalization features in a hospital-stay summary generation system. In *Proceedings of the 52nd Hawaii International Conference on System Sciences*.
- Sabita Acharya, Andrew D Boyd, Richard Cameron, Karen Dunn Lopez, Pamela Martyn-Nemeth, Carolyn Dickens, Amer Ardati, Jose D Flores, Matt Baumann, Betty Welland, et al. 2018. What happened to me while i was in the hospital? challenges and opportunities for generating patient-friendly hospitalization summaries. *Journal of Healthcare Informatics Research*, pages 1–17.
- Sabita Acharya, Barbara Di Eugenio, Andrew D Boyd, Karen Dunn Lopez, Richard Cameron, and Gail M Keenan. 2016. Generating summaries of hospitalizations: A new metric to assess the complexity of medical terms and their definitions. In *The 9th International Natural Language Generation conference*, page 26.
- Tim Althoff, Kevin Clark, and Jure Leskovec. 2016. Large-scale analysis of counseling conversations: An application of Natural Language Processing to mental health. *Transactions of the Association for Computational Linguistics*, 4:463–476.
- Terry C Davis, Sandra W Long, Robert H Jackson, EJ Mayeaux, Ronald B George, Peggy W Murphy, and Michael A Crouch. 1993. Rapid estimate of adult literacy in medicine: a shortened screening instrument. *Family medicine*, 25(6):391–395.
- Barbara Di Eugenio, Andrew D Boyd, Camillo Lugaresi, Abhinaya Balasubramanian, Gail M Keenan, Mike Burton, Tamara G Rezende Macieira, Karen Dunn Lopez, Carol Friedman, Jianrong Li, et al. 2014. Patientnarr: Towards generating patient-centric summaries of hospital stays. *INLG 2014*, page 6.
- Shelley Duval and Robert A. Wicklund. 1972. A theory of objective self awareness.
- Jinnet Briggs Fowles, Paul Terry, Min Xi, Judith Hibbard, Christine Taddy Bloom, and Lisa Harvey. 2009. Measuring self-management of patients and employees health: further validation of the patient activation measure (pam) based on its relation to employee characteristics. *Patient education and counseling*, 77(1):116–122.
- Martha M. Funnell. 2000. Helping patients take charge of their chronic illnesses. *Family practice management*, 7(3):47.
- Judith H Hibbard, Jean Stockard, Eldon R Mahoney, and Martin Tusler. 2004. Development of the patient activation measure (pam): conceptualizing and measuring activation in patients and consumers. *Health services research*, 39(4p1):1005–1026.
- Yun-Hee Jeon, Stefan G Kraus, Tanisha Jowsey, and Nicholas J Glasgow. 2010. The experience of living with chronic heart failure: a narrative review of qualitative studies. *BMC health services research*, 10(1):77.
- Saad Mahamood and Ehud Reiter. 2011. Generating affective natural language for parents of neonatal infants. In *Proceedings of the 13th European Workshop on Natural Language Generation*, pages 12–21. Association for Computational Linguistics.
- J Michael McGinnis, Leigh Stuckhardt, Robert Saunders, Mark Smith, et al. 2013. *Best care at lower cost: the path to continuously learning health care in America*. National Academies Press.
- Steffen Pauws, Albert Gatt, Emiel Kraemer, and Ehud Reiter. 2019. Making effective use of healthcare data using data-to-text technology. In *Data Science for Healthcare*, pages 119–145. Springer.
- James W Pennebaker, Matthias R Mehl, and Kate G Niederhoffer. 2003. Psychological aspects of natural language use: Our words, our selves. *Annual review of psychology*, 54(1):547–577.
- Verónica Pérez-Rosas, Rada Mihalcea, Kenneth Resnicow, Satinder Singh, and Lawrence An. 2017. Understanding and predicting empathic behavior in counseling therapy. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1426–1435.
- Tom Pyszczynski and Jeff Greenberg. 1987. Self-regulatory perseveration and the depressive self-focusing style: a self-awareness theory of reactive depression. *Psychological bulletin*, 102(1):122.
- Barbara Riegel, Christopher S Lee, Nancy Albert, Terry Lennie, Misook Chung, Eun Kyeong Song, Brooke Bentley, Seongkum Heo, Linda Worrall-Carter, and Debra K Moser. 2011. From novice to expert: confidence and activity status determine

- heart failure self-care performance. *Nursing research*, 60(2):132–138.
- Guergana K Savova, James J Masanz, Philip V Ogren, Jiaping Zheng, Sunghwan Sohn, Karin C Kipper-Schuler, and Christopher G Chute. 2010. Mayo clinical text analysis and knowledge extraction system (ctakes): architecture, component evaluation and applications. *Journal of the American Medical Informatics Association*, 17(5):507–513.
- Donia Scott, Catalina Hallett, and Rachel Fettiplace. 2013. Data-to-text summarisation of patient records: Using computer-generated summaries to access patient histories. *Patient education and counseling*, 92(2):153–159.
- Alvin Chuen Wei Seah, Khoon Kiat Tan, Juvena Chew Huang Gan, and Wenru Wang. 2016. Experiences of patients living with heart failure: a descriptive qualitative study. *Journal of Transcultural Nursing*, 27(4):392–399.
- Johanna Shapiro. 1993. The use of narrative in the doctor-patient encounter. *Family Systems Medicine*, 11(1):47.
- Steven R Wilson, Rada Mihalcea, Ryan L Boyd, and James W Pennebaker. 2016. Cultural influences on the measurement of personal values through words. In *2016 AAAI Spring Symposium Series*.
- Bo Xiao, Daniel Bone, Maarten Van Segbroeck, Zac E Imel, David C Atkins, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2014. Modeling therapist empathy through prosody in drug addiction counseling. In *Fifteenth Annual Conference of the International Speech Communication Association*.

A Supplemental Material

Patient: Quality of life is more valuable to me than working on the next project
Interviewer: Alright. Okay. So self care...
Patient: matters
Interviewer: Is primary. Yeah it matters. Yeah definitely.
Patient: Yeah
Interviewer: Now do you have some support at home? You girlfriend's around? She is going to help you out or?
Patient: Yeah. You know I am probably the most fortunate guy you are gonna meet today. I am fortunate in that even aside from my family, there are friends that genuinely care about me.
Interviewer: Alright
Patient: I mean really genuinely..genuinely care about me
Interviewer: Ah
Patient: Care about what I do..care about my well being..care about who I am..what I do..I have friends that will not allow me to fail.

Figure 4: A portion of an interview where the patient talks about the things that matter to him .

Patient: Ahm hm. Yeah. So they give you this Lasix they try to get that fluid and stuff off your lungs that makes your legs swell too.
Interviewer: Ah hm. And is that water part of the heart thing? The heart problem or?
Patient: Right. Ah hm
Interviewer: What is this Lasix thing? What is that?
Patient: it is a shot they give you. In fact they give you pills too... So they will bring water..you know..out of your body

Figure 5: A portion of an interview where the patient explains about her health issue.

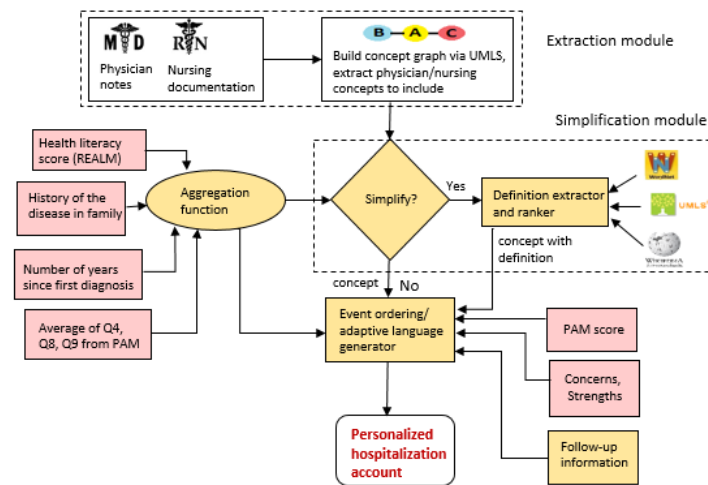


Figure 6: Work-flow of the algorithm for generating personalized summaries. *Extraction module* is responsible for exploring the relationship between the medical terms from the input documents. *Simplification module* identifies difficult medical terms and provides explanations to them. The boxes in pink represent the features that guide the personalization process.

TDDiscourse: A Dataset for Discourse-Level Temporal Ordering of Events

Aakanksha Naik

Luke Breitfeller

Carolyn Rose

Language Technologies Institute, Carnegie Mellon University

{anaik, mbreitfe, cprose}@cs.cmu.edu

Abstract

Prior work on temporal relation classification has focused extensively on event pairs in the same or adjacent sentences (local), paying scant attention to discourse-level (global) pairs. This restricts the ability of systems to learn temporal links between global pairs, since reliance on local syntactic features suffices to achieve reasonable performance on existing datasets. However, systems should be capable of incorporating cues from document-level structure to assign temporal relations. In this work, we take a first step towards discourse-level temporal ordering by creating **TDDiscourse**, the first dataset focusing specifically on temporal links between event pairs which are more than one sentence apart. We create TDDiscourse by augmenting TimeBank-Dense, a corpus of English news articles, manually annotating global pairs that cannot be inferred automatically from existing annotations. Our annotations double the number of temporal links in TimeBank-Dense, while possessing several desirable properties such as focusing on long-distance pairs and not being automatically inferable. We adapt and benchmark the performance of three state-of-the-art models on TDDiscourse and observe that existing systems indeed find discourse-level temporal ordering harder.

1 Introduction

Temporal ordering of events is a crucial problem in automated text analysis. Systems capable of performing this task find widespread applicability in areas such as time-aware summarization, temporal information extraction or event timeline construction. Prior work has focused extensively on creating annotated corpora for temporal ordering, some notable efforts being the development of the TimeML annotation schema (Pustejovsky et al., 2003), TimeBank (Pustejovsky et al.) and

TimeBank-Dense (Cassidy et al., 2014). However, most work has focused mainly on local ordering, i.e., events present in the same or adjacent sentences. This leads to a major drawback, also pointed out by Reimers et al. (2016). Low prevalence of global discourse-level temporal ordering annotation in existing datasets allows systems to achieve moderate performance simply using local syntactic cues. Having more global annotations would require systems to incorporate global consistency and assimilate features from document-level structure and flow to achieve high performance, thus presenting a more challenging task. In this work, we present TDDiscourse, a dataset focused on discourse-level temporal ordering.

We create TDDiscourse by augmenting TimeBank-Dense (Cassidy et al., 2014), a corpus of English news articles, with more long-distance event pair annotations. Our work makes the first attempt to *explicitly* annotate relations between event pairs that are more than one sentence apart, a more difficult annotation task than previous datasets. In addition to facing similar challenges as prior work (eg: hypothetical/negated events (Cassidy et al., 2014)), we tackle new *global discourse-level* issues such as incorporating event coreference and causality/prerequisite links arising from world knowledge. To handle these, we design a careful coding scheme that achieves high inter-annotator agreement (Cohen’s Kappa of 0.69 on the test set). However, getting expert manual annotation for all possible long-distance event pairs is expensive. Moreover, it is possible to leverage annotations from existing datasets to automatically infer temporal relations for certain event pairs. To make optimal use of expert annotation, we develop a heuristic algorithm for automatic inference of temporal relations using EventTime (Reimers et al., 2016) and apply this

to all documents.¹ We then randomly subsample the unannotated event pairs and source expert annotations for those. At 6150 pairs, our manually annotated subset is of the same size as TimeBank-Dense. Adding the automatic subset makes our dataset 7x larger (§6). Finally, we perform a principled comparison between manual and automatic pairs by annotating 3 test documents (107 manual and 110 automatic event pairs) with phenomena required to reason correctly about the pair. These annotations suggest that our manual subset exhibits a high proportion of global discourse-level phenomena such as reasoning about chains of events.

In addition to developing TDDiscourse, we adapt three state-of-the-art models on TimeBank-Dense for discourse-level temporal ordering and benchmark their performance on our data, separating scores on manual and automatic subsets. We observe that models perform worse on average on TDDiscourse, with none beating a majority class baseline on the manual subset. A manual analysis of model errors reveals key shortcomings of current temporal ordering techniques. We offer our dataset² as a challenging new resource for the temporal ordering community and hope that insights from our analysis will spark interest in the development of more global discourse-aware models.

2 Related Work

2.1 Prior Work on Temporal Annotation

The development of TimeML (Pustejovsky et al., 2003) and TimeBank (Pustejovsky et al.) marked the first attempt towards creating a corpus for temporal ordering of events. TimeML uses temporal links (TLINKs) (Setzer, 2002), to represent ordering. A TLINK expresses the temporal relation between two events. For example, an event $e1$ can occur *before* another event $e2$. TimeBank is annotated using TLINKs, but the number of possible TLINKs in a document is large (quadratic in number of events). So annotation is restricted to a subset of TLINKs, leading to sparsity. To combat this, several works attempted to create denser corpora (Bramsen et al., 2006; Kolomiyets et al., 2012; Do et al., 2012; Cassidy et al., 2014), but still focused largely on local TLINKs.

¹We validate our algorithm by obtaining human annotations for a subset of 100 examples and observing agreement with the generated label in 99% cases

²<https://github.com/aakanksha19/TDDiscourse>

Reimers et al. (2016) addressed high annotation cost by proposing a new scheme in which events were associated with explicit time expressions. Annotation effort now scaled linearly with number of events, making it feasible to annotate all of them. Using this scheme, they created EventTime, which had some discourse-level temporal annotation. However this dataset had one major drawback: events which could not be associated with a time expression were ignored. We observed that it may not always be possible to determine specific times for an event, but ordering it with respect to other events is often possible based on world knowledge. For example, consider the snippet: “Police discover body of *kidnapped* man. Police found the man’s *dismembered* body wrapped in garbage bags”. In this text, *dismembered* cannot be associated with a time. But the temporal relation between *dismembered* and *kidnapped* is clear because the kidnapping should have happened *before* dismembering. Based on this, we address the drawback in EventTime, by using TLINK-based annotation, which is expensive but allows more expressive power. Following TimeML, we augment TimeBank-Dense (Cassidy et al., 2014) with global discourse-level TLINKs. To optimize manual effort, we automatically generate all TLINKs that can be inferred from EventTime. Then, we manually annotate a large subset of missing TLINKs involving events not associated with specific dates.

Most recently, Ning et al. (2018b) proposed a new scheme, which labels TLINKs based only on event start time. This improved inter-annotator agreement allowing for crowdsourcing of long-distance annotations at lower cost. However, they focused only on verb events, whereas our work is broader in scope and poses no such restrictions.

2.2 Prior Temporal Ordering Systems

TimeBank and the TempEval tasks (Verhagen et al., 2007, 2010; UzZaman et al., 2013) spurred the development of many temporal ordering systems (UzZaman and Allen, 2010; Llorens et al., 2010; Strötgen and Gertz, 2010; Chang and Manning, 2012; Chambers, 2013; Bethard, 2013). More recently, TimeBank-Dense and EventTime prompted development of newer models (Chambers et al., 2014; Mirza and Tonelli, 2016; Cheng and Miyao, 2017; Reimers et al., 2018). Most systems built for TimeBank/ TimeBank-Dense focus

on TLINKs between events in the same or adjacent sentences, relying on local features rather than document-level structure, with some exceptions. Chambers and Jurafsky (2008); Denis and Muller (2011); Ning et al. (2017) introduce document-level consistency via integer linear programming constraints. Bramsen et al. (2006); Do et al. (2012) also incorporate document-level structure, but focus on different corpora. Reimers et al. (2018) develop a model for EventTime, which uses a decision tree of CNNs to associate each event from a document with a time. Several works have explored techniques to incorporate document-level cues such as event coreference (Do et al., 2012; Llorens et al., 2015) and causality (Do et al., 2012; Ning et al., 2018a) in temporal ordering systems. However, due to a lack of standard datasets focusing on global discourse-level links, most work has been evaluated on datasets of their own creation or standard datasets with mainly local TLINKs. This further stresses the need for a standardized benchmarking effort, which we address by evaluating adaptations of three state-of-the-art systems on our dataset (§8).

3 Constructing TDDiscourse

To emphasize the need for a global discourse-level focus in temporal ordering, we develop TDDiscourse, the first dataset which focuses *explicitly* on TLINK annotations between event pairs that are more than one sentence apart. To create TDDiscourse, we augment a subset of documents from TimeBank with global TLINKs. We use the same set of 36 documents as TimeBank-Dense (Cassidy et al., 2014) and EventTime (Reimers et al., 2016) to facilitate comparison with previous work. We also utilize the same set of temporal relations as TimeBank-Dense.³ Table 1 gives a brief summary of these relations. To add global links, we use two approaches:

- **Manual annotation:** We manually label a subset of global TLINKs using document cues, world knowledge and causality (§4). To optimize human effort, we ensure that these TLINKs are not automatically inferable.
- **Automatic inference:** We use a heuristic algorithm to automatically label global TLINKs using EventTime (§5) annotations, to generate a large number of links at low

³We discard the “vague” label since we do not require annotators to label all event pairs

Symbol	Relation
a	$e1$ occurs after $e2$
b	$e1$ occurs before $e2$
s	$e1$ and $e2$ are simultaneous
i	$e1$ includes $e2$
ii	$e1$ is included in $e2$

Table 1: Temporal relation set used in TDDiscourse. All relations are mutually exclusive.

cost.

4 Manual Annotation

In this phase, we ask experts⁴ to label discourse-level TLINKs that cannot be inferred automatically.⁵ Getting expert annotation for all missing TLINKs is expensive. Hence, we randomly subsample TLINKs not annotated by TimeBank-Dense or automatic inference. This subsample is as large as TimeBank-Dense, thus doubling the data size while making the overall task harder (see §8). Note that TLINKs annotated in this phase may involve events for which a specific time of occurrence cannot be determined, which were ignored in EventTime. We refer to this subset as **TDD-Man**.

Since TLINKs are not restricted to the same or adjacent sentences, our annotation task becomes harder, requiring cues from the entire document. Many TLINKs also require the use of causal links and world knowledge to label the relation. Based on our observations, we develop a coding scheme. To ensure high inter-annotator agreement, we refine our scheme over multiple rounds of annotation and discussion of disagreements.

4.1 Coding Scheme

Our scheme reduces the task of labeling a TLINK to a set of concrete decision steps:

1. Using textual cues
2. Using world knowledge
3. Using narrative ordering

A TLINK may be assigned a label at any step. If it cannot be assigned a label, it moves on to the next step. Information from previous steps is retained, making it possible to combine multiple sources of evidence. For example, textual cues may not suffice, but they can be used in conjunction with world knowledge to label a pair. We

⁴Expert annotators are the authors of the paper, with a background in computational linguistics

⁵The automatic inference algorithm is explained in §5

Snippet
Atlanta nineteen ninety-six. A bomb blast shocks the Olympic games. One person is killed . January nineteen ninety-seven. Atlanta again. This time a bomb at an abortion clinic. More people are hurt .
Event pair: <i>blast, hurt</i>
Relation: before
Textual cues: Event <i>blast</i> occurred in 1996. Event <i>hurt</i> occurred because of second bomb blast in 1997.

Table 2: Sample document-level textual cues used during temporal annotation

choose to organize our coding scheme as mentioned above, to make the process of gathering evidence about an event pair systematic, and ensure that experts do not miss important cues. The final step is guaranteed to assign a label. We choose not to allow annotators to leave event pairs unlabeled or label them “vague”, to keep them from overusing this option. Owing to this decision, we need to develop mechanisms for handling TLINKs containing events which have not actually occurred (eg: negated, hypothetical or conditional events). Drawing from prior work, we interpret these events using a *possible worlds* analysis, in which the event is treated as if it has occurred. We refer interested readers to (Chambers et al., 2014) for a more detailed discussion.

4.1.1 Using textual cues

In this step, we use document-level textual cues to label a TLINK. The cues used are similar to those used in previous datasets (Cassidy et al., 2014). Table 2 gives an example of the types of cues used.

A key textual cue we use here is event coreference. Event coreference has not been used for annotation because the occurrence of coreferent events in adjacent sentences is rare. However, this cue is crucial for global discourse-level annotation. Since TimeBank does not contain event coreference annotation, we develop a procedure to annotate our document subset. Our procedure is based on the ERE (Entities, Relations, and Events) scheme (Song et al., 2015), which cannot be directly used for TimeBank due to differing notions of what constitutes an event and different metadata. In our procedure, events are considered coreferent iff they share the following:

- Entities involved in the event
- Temporal attributes
- Location attributes
- Realis (whether event is real or hypothetical)

Events which are synonymous in context are also

considered coreferent (for instance, in “...held an interview Monday. The segment covered...”, *interview* and *segment* are synonymous). These attributes (barring temporal) are not provided in TimeBank and must be inferred. Often, an event may only have partial information about these attributes - here we use human judgment. Our definition of coreference is closer to the strict notion of “event identity” in Light ERE than the relaxed definition in Rich ERE.⁶ To test our procedure, we select all “simultaneous” TLINKs from TimeBank-Dense to ensure that our sample contains a sizeable proportion of *possibly* coreferent event pairs. The corpus contains 179 “simultaneous” links, of which 93 are event pair TLINKs. Our first annotation pass achieves high agreement between two annotators, with a Kappa of 0.70. We refine our guidelines through an adjudication step, reaching perfect agreement on this sample. Post-adjudication guidelines are used to annotate event coreference for all documents. Resulting annotations are used as textual cues in our scheme. Based on textual cues, an appropriate label is assigned to a TLINK. Coreferent TLINKs are labeled “simultaneous”. Unlabeled links move on to the next decision step.

4.1.2 Using world knowledge

This step uses real world knowledge to determine causal/prerequisite links which are used to label a TLINK. We consider both events in the TLINK and determine whether they possess one or both of the following:

- **Causal Link:** Two events have a causal link if the occurrence of one event results in the other event coming about. For example, in the sentence “The paper got wet when I spilled water on it”, the event pair (spilled, wet) have a causal link.
- **Prerequisite Link:** Two events have a prerequisite link if one event *must* occur before the other can happen. For example, in the sentence “We cooked dinner and ate it”, the event pair (cooked, ate) have a prerequisite link. Note that we use the knowledge that a meal must be cooked before it can be eaten, though it is not explicitly mentioned.

We examine the event pair in the context of the entire document to detect causal/prerequisite links, also allowing weak or transitive links. For in-

⁶Examples in the appendix

Rule	Label
TLINK=(A, B), A=P	Before
TLINK=(A, B), A=I	Includes
TLINK=(B, A), A=P	After
TLINK=(B, A), A=I	Is Included

Table 3: Labels assigned to event pairs based on event and TLINK metadata

stance, in the text “Diplomacy is making headway in resolving the UN’s standoff with Iraq. One major sticking point has been Iraq’s proposal...”, *proposal* causes *standoff*, which is a prerequisite for *resolving*. Hence, the pair (proposal, resolving) is considered causal/prerequisite. Our assignment of causal/prerequisite links is unordered. For example, reverse event pairs (wet, spilled), (ate, cooked), and (resolving, proposal) are also considered causal/prerequisite. Link order is taken into consideration while assigning a temporal relation.

If two events contain a causal/prerequisite link, we identify the event in the pair that causes or is a prerequisite for the other. We call this event “A” and the other “B”. For example, (spilled, wet) is expressed as (A, B), while (wet, spilled) is expressed as (B, A). To label the TLINK, we determine whether A is a point (P) or interval (I) event using existing date annotations from EventTime (Reimers et al., 2016). This helps us catch cases where A is a long-lasting interval and the time span for B is completely included in A. For instance, in “the war forced civilians to evacuate”, (war, evacuate) has a causal/prerequisite link with *war* being event A. Though *war* caused *evacuation*, it is reasonable to expect that the war started *before* and ended *after* evacuation. If A is not present in EventTime (i.e it cannot be assigned a specific time), we use our judgment to determine event length. We then assign a label as per table 3. Unlabeled links are passed to the next step.

4.1.3 Using narrative ordering

This step uses a heuristic based on the intuition that events in a narrative are often presented in chronological order. To label a TLINK, we determine which event appeared first in the document. This event is called “A”, and the other is “B”. We then detect whether A is a point (P) or interval (I) from EventTime, falling back to our own judgment if it is not present. Finally, a label is assigned following table 3. This step is guaranteed to assign a label since every pair will have a narrative-based order.

Dataset	Kappa
TimeBank	0.71
TimeBank-Dense	0.56-0.64
TDD-Man	0.69

Table 4: Inter-annotator agreement (Cohen’s Kappa) on temporal ordering datasets. Kappa scores for TDD-Man are reported on the test set containing 1500 links.

	a	b	s	i	ii
a	137	22	0	12	22
b	30	311	1	72	23
s	0	0	42	5	4
i	9	36	3	462	35
ii	12	32	0	21	209

Table 5: Relation agreement between annotators on the TDD-Man test set containing 1500 links.

4.2 Inter-annotator agreement

Our annotation scheme was developed over multiple rounds of coding and discussion between two experts. In each round, experts separately annotated a set of 10-15 TLINKs.⁷ Cohen’s Kappa was computed and disagreements were discussed. TLINKs were changed in every round to ensure exposure to diverse event pair types. Inter-annotator agreement in preliminary rounds ranged from 0.48-0.69. The final coding scheme resulted in an agreement of 0.69 on the test set. Table 4 shows that our agreement is comparable to prior work. Table 5 presents a class-wise distribution of agreements between pairs of annotators. Disagreements mainly include cases where one annotator chose after/before while the second chose includes/is included (64%). This indicates that determining precise end-points for an interval event is difficult, as corroborated by Ning et al. (2018b).

5 Automatic Inference

This approach uses automatic inference to derive new TLINKs at low cost from EventTime (Reimers et al., 2016), which assigns specific times to events. EventTime divides events into two types: SingleDay and MultiDay. SingleDay events are assigned dates, while MultiDay events are assigned intervals. Possible event pairs can be divided into three categories: SS (both events are SingleDay), SM (one event is SingleDay while the other is MultiDay) and MM (both events are MultiDay). Not all assigned dates and intervals are exact. EventTime relies heavily on under-specified

⁷chosen from documents in the development set

temporal expressions (such as “after1998-06-08”), making automatic inference non-trivial.

We follow separate algorithms to infer TLINKs for each pair type (SS, SM and MM). For SS pairs, both events are associated with dates, which may be expressed in one of four ways⁸, resulting in 16 date combinations for SS links. We develop heuristics⁹ for each combination, which generate a temporal relation based on date values. Our heuristics were developed with a focus on precision to avoid adding incorrect links. Often, a relation cannot be generated. For example, consider two events associated with the same date “after02-01-1999”. We know that both events occur after 02-01-1999, but we cannot infer their order with respect to *each other*. In such cases, we do not label the pair. For SM pairs, one event is associated with a time interval having begin and end dates. Here we use the SS pair inference algorithm to generate relations between the SingleDay event date and the MultiDay event begin and end dates. These relations are compared to infer the label for the pair. For MM pairs, both events have begin and end dates. We infer relations between begin and end points using SS link inference and use these to infer the pair label. After inference, we perform temporal closure, according to [Chambers et al. \(2014\)](#). To evaluate validity of generated TLINKs, we randomly sample a subset of 100 TLINKs and ask three annotators¹⁰ to determine the correctness of the labels. All annotators unanimously agree with the assigned label in 99% cases. We call this subset **TDD-Auto**.

6 Dataset Statistics

Our data construction pipeline produces the first dataset focused on temporal links between global discourse-level event pairs (**TDDiscourse**), consisting of two subsets **TDD-Man** and **TDD-Auto**. Table 6 presents train, dev and test set sizes for both subsets, Timebank-Dense as well as an augmented version of TimeBank-Dense with additional links inferred via temporal closure. Our complete dataset is 7x larger than both, indicating that our construction adds valuable new TLINKs. **TDD-Man** itself is as large as TimeBank-Dense

⁸MM-DD-YYYY, afterMM-DD-YYYY, beforeMM-DD-YYYY, afterMM-DD-YYYYbeforeMM-DD-YYYY (MM-DD-YYYY stands for a specific date value)

⁹Sample heuristics provided in the appendix

¹⁰Annotators were volunteers with no vested interest in the corpus

Dataset	Train	Dev	Test
TB-Dense	4032	629	1427
TB-Dense + Closure	4399	722	1575
TDD-Man	4000	650	1500
TDD-Auto	32609	1435	4258

Table 6: Dataset sizes for TimeBank-Dense and our dataset. Note that we only count event-event TLINKs

and can be used in isolation, however incorporating **TDD-Auto** provides a large amount of training data making the task more amenable to deep neural net approaches.

Table 7 presents class distributions for TDD-Man and TDD-Auto test sets. Though there is a clear majority class, both sets are more balanced than TimeBank-Dense, in which 40% event pairs are labeled “vague”. To evaluate the presence of long-distance TLINKs, we present the distribution of distance between event pairs from annotated TLINKs in table 8 which shows that nearly 53% TLINKs in our dataset comprise of event pairs which are more than 5 sentences apart. Further, to gain deeper insight into global discourse-level phenomena exhibited by our dataset, we augment 3 documents from the test set (107 manual and 110 automated event pairs) with additional annotations about phenomena required to label them correctly. We consider the following phenomena:

- **SingleSent (SS):** Textual cues from sentences containing the events suffice to predict the relation (irrespective of distance).
- **Chain Reasoning (CR):** Correct relation prediction requires reasoning about other events from the document.
- **Tense Indicator (TI):** For verb events, tense information indicates the correct relation.
- **Future Events (FE):** One or both events from the pair will occur in the future.
- **Hypothetical/ Negated (HN):** One or both events are hypothetical or negated.
- **Event Coreference (EC):** Event coreference resolution is needed to predict relation.
- **Causal/ Prereq (CP):** Causal/ prerequisite links must be identified to predict relation.
- **World Knowledge (WK):** Real world knowledge is needed to identify the relation.

Table 9 shows the distribution of these phenomena in TDD-Man and TDD-Auto. TDD-Man shows a higher percentage of difficult phenomena (CR, CP). On the other hand, TDD-Auto shows high prevalence of SS, indicating that local information

Dataset	a	b	s	i	ii
TB-Dense	0.18	0.22	0.02	0.05	0.06
TDD-Man	0.13	0.27	0.03	0.38	0.19
TDD-Auto	0.28	0.32	0.16	0.11	0.13

Table 7: Class distributions for our test sets and TimeBank-Dense. Note that the distribution for TimeBank-Dense does not sum to 1, since it includes a vague class.

Dataset	<5	<10	<15	<20	>20
TDD-Man	0.40	0.40	0.15	0.04	0.01
TDD-Auto	0.50	0.32	0.12	0.05	0.01

Table 8: Distribution of distance between events for all TLINKs in our test sets (in terms of #sentences)

may be sufficient to label many long-distance links in this subset correctly. This principled comparison of both subsets leads us to hypothesize that models which perform well on TimeBank-Dense, should achieve similar scores on TDD-Auto but perform much worse on TDD-Man.

7 Experiments

To statistically evaluate the difficulty of TDDiscourse, we adapt and benchmark three SOTA models on our data. Our results reveal interesting insights about model drawbacks, highlighting the need to shift focus to handling global discourse-level phenomena such as chain reasoning.

7.1 Adapting State-of-the-Art Models for Benchmarking

As most state-of-the-art temporal ordering models are built on datasets containing mainly local TLINKs, they are not well-equipped to handle global TLINKs. Hence, we adapt these models to ensure fair evaluation. We focus on the following: **CAEVO** (Chambers et al., 2014): This system

Phenomenon	TDDMan	TDDAuto
SS	25.23%	90.91%
CR	58.88%	9.09%
TI	12.10%	46.36%
FE	36.45%	29.09%
HN	14.02%	19.09%
EC	16.82%	4.55%
CP	64.49%	29.09%
WK	16.82%	0.91%

Table 9: Distribution of various phenomena in the annotated test subset. These phenomena were labeled manually.

consists of specialized learners (sieves) which include heuristic rules and trained models. For each document, sieves run in decreasing order of precision. Decisions made by earlier sieves constrain following ones. This framework integrates transitive reasoning, but decisions made by earlier sieves cannot be overturned, causing error cascades. To extend CAEVO, we increase window sizes and remove the AllVague sieve.¹¹

BiLSTM (Cheng and Miyao, 2017): Inspired by Xu et al. (2015), this model uses a BiLSTM classifier. For each pair, dependency paths from source and target events to the sentence root are fed to a BiLSTM. For events in adjacent sentences, source and target event sentences are assumed to be connected to a "common root". We follow the same framework to build a BiLSTM.

SP+ILP (Ning et al., 2017): CAEVO and BiLSTM make separate local decisions for each TLINK, which may result in global inconsistency. For example, for events A, B and C, if A occurs before B and B occurs before C, transitivity implies that A occurs before C. Models classifying each pair independently may assign a different relation to A-C. To correct this, Ning et al. (2017) proposed SP+ILP, which uses a structured perceptron with ILP constraints, explicitly enforcing global consistency. This model was trained on TimeBank-Dense which contains fewer TLINKs per document, making joint learning tractable with loose transitivity constraints. But loose transitivity is an issue for our data with 7x more TLINKs, since the number of constraints increases tremendously. To improve tractability, we define a stricter transitivity constraint. Let E , R and P be sets of events, temporal relations and event pairs respectively ($P = \{(e_i, e_j) \in E \times E | e_i, e_j \in E, i \neq j\}$). We define an array of binary indicator variables y , where $y_{\langle r, i, j \rangle}$ indicates whether the relation r holds between events e_i and e_j . Our objective function is defined as:

$$\arg \min_y \sum_{\langle e_i, e_j \rangle \in P} \sum_{r \in R} -y_{\langle r, i, j \rangle} \log p_{\langle r, i, j \rangle} \quad (1)$$

subject to the following constraints:

$$y_{\langle r, i, j \rangle} \in \{0, 1\}, \forall (e_i, e_j) \in P, \forall r \in R \quad (2)$$

$$\sum_{r \in R} y_{\langle r, i, j \rangle} = 1, \forall (e_i, e_j) \in P \quad (3)$$

¹¹since our data does not include the vague class. We also remove the WordNet sieve and add MLEventEventDiffSent. For more details on these sieves, we refer interested readers to Chambers et al. (2014)

System	TB-Dense			TDD-Auto			TDD-Man		
	P	R	F1	P	R	F1	P	R	F1
MAJOR	40.5	40.5	40.5	34.2	32.3	33.2	37.8	36.3	37.1
CAEVO	49.9	46.6	48.2	61.1	32.6	42.5	32.3	10.7	16.1
BiLSTM	63.9	38.9	48.4	55.7	48.3	51.8	24.9	23.8	24.3
SP	37.7	37.8	37.7	43.2	43.2	43.2	22.7	22.7	22.7
SP+ILP	58.4	58.4	58.4	46.4	45.9	46.1	23.9	23.8	23.8

Table 10: Performance of SOTA models on TB-Dense, TDD-Auto and TDD-Man. MAJOR represents a majority-class baseline. We report performance on non-vague event-event links for TB-Dense to ensure fair comparison.

$$y_{\langle r1,i,j \rangle} + y_{\langle r2,j,k \rangle} - y_{\langle r3,i,k \rangle} \leq 1, \quad (4)$$

$\forall (e_i, e_j), (e_j, e_k), (e_i, e_k) \in P, \forall (r1, r2, r3) \in TC$
where $p_{\langle r,i,j \rangle}$ is the probability that event pair (e_i, e_j) has label r . (2) ensures that indicator variables are binary, (3) forces event pairs to be assigned a unique label and (4) imposes transitivity. TC denotes the set of transitive relation triples.¹² Relation probabilities $(p_{\langle r,i,j \rangle})$ come from the structured perceptron. In addition to this model, we also evaluate the structured perceptron (SP) in isolation, which lets us study the effect of introducing global consistency via ILP.

8 Results and Analysis

We benchmark 4 adapted SOTA models (CAEVO, BiLSTM, SP and SP+ILP) on TDD-Auto and TDD-Man. SP is a local perceptron-based classifier, while SP+ILP introduces transitivity via ILP into the perceptron. This . For tractability, we limit all models to using event pairs which are 15 or fewer sentences apart. This discards only 5% of our data (table 8). Table 10 presents the benchmarking results. We also benchmark models on TimeBank-Dense (TB-Dense) to demonstrate that our modifications do not affect performance on local TLINKs.

All models perform better than a majority class baseline on TDD-Auto. The BiLSTM and SP perform particularly well, achieving a higher F1 than TB-Dense, while CAEVO and SP+ILP show slight degradation in comparison to TB-Dense. This corroborates our hypothesis that many long-distance TLINKs in TDD-Auto can be handled with local information. However, all models show a significant drop on TDD-Man, with none outperforming a majority class baseline. Further analysis of model errors offers valuable insights into which phenomena are not handled by models, posing in-

¹²(“before”, “before”, “before”) form a transitive relation triple as A before B and B before C implies A before C

teresting challenges for future work.

Maintaining global consistency: Most SOTA models make separate local decisions for each pair and are not globally consistent. Adding global consistency improves the performance of a local classifier, as evinced by a 3-point F1 gain observed on adding ILP to SP. We validate this observation by performing a transitivity analysis of BiLSTM and SP+ILP on TDD-Auto. We go through all event triples (e_1, e_2, e_3) . For each model, if (e_1, e_2) , (e_2, e_3) and (e_1, e_3) are all assigned labels, we check whether labels are consistent. For example, e_1 after e_2 , e_2 after e_3 and e_1 after e_3 is a consistent assignment. We observe that though the BiLSTM has higher F1, it maintains transitivity in 41.9% cases, while SP+ILP enforces transitivity in 53.6% cases, a 12% increase. We believe that incorporating such constraints into neural models can help, which we delegate to future work.

Incorporating real world knowledge: To examine the dismal performance of all models on TDD-Man, we manually look at 100 pairs on which all models made mistakes. 40% of these cases require real world knowledge. Some examples include determining that “military actions” refers to the same event as “air strikes” (strikes would have to be carried out by the military which cannot be inferred from text), or knowing that certain events (eg: “war”) are long-term. No SOTA model currently has this ability.

Using event coreference and structure: Our analysis reveals another source of errors arising from models’ inability to handle event coreference and event structure such as sub-events or aspectual predication, a grammatical device which focuses on different facets of event history (eg: using “begin” to indicate initiation) (Pustejovsky et al., 2003). This inability causes models to fail in 22% cases indicating that exploiting rich event structure information is a promising direction.

Dealing with hypothetical or negated events:

We observe that SOTA models do not possess the ability to handle these, causing 31% of errors.

9 Conclusion and Future Work

In this work, we created TDDiscourse, the first dataset focused on global discourse-level temporal ordering. Our annotation scheme for TDDiscourse handled several issues which have not been explicitly addressed in prior work. We further adapted and benchmarked 3 SOTA models. All models, on average, performed worse on TDDiscourse, validating the difficulty of the task. Our error analysis reveals key phenomena not handled by current systems, such as hypothetical/negated events, event coreference, aspectual predication, real world knowledge and global consistency. Future work in temporal ordering must address these issues, and we suggest several avenues for exploration, such as a BiLSTM-ILP joint learning framework which has the advantage of combining representational power of neural models with key linguistic insights, and introducing event coreference information via ILP into a structured learning approach similar to [Ning et al. \(2017\)](#). Finally, we hope that our dataset offers a challenging testbed for the development of more global discourse-aware models for temporal ordering.

Acknowledgements

This work was supported by the University of Pittsburgh Medical Center (UPMC) and Abridge AI Inc through the Center for Machine Learning and Health at Carnegie Mellon University. It was also funded in part through NSF IIS 1723454. The authors would like to thank Lisa Carey Lohmueller, Shivani Poddar and Michael Miller Yoder for assisting in human evaluation of TDD-Auto, Evangelia Spiliopoulou for help in implementing parts of the SP+ILP system and the anonymous reviewers for their helpful feedback on this work.

References

Steven Bethard. 2013. [Cleartk-timeml: A minimalist approach to tempeval 2013](#). In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 10–14, Atlanta, Georgia, USA. Association for Computational Linguistics.

Philip Bramsen, Pawan Deshpande, Yoong Keok Lee, and Regina Barzilay. 2006. [Inducing temporal graphs](#). In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, pages 189–198, Sydney, Australia. Association for Computational Linguistics.

Taylor Cassidy, Bill McDowell, Nathanel Chambers, and Steven Bethard. 2014. An annotation framework for dense event ordering. Technical report, CARNEGIE-MELLON UNIV PITTSBURGH PA.

Nathanael Chambers. 2013. Navytime: Event and time ordering from raw text. Technical report, Naval Academy Annapolis MD.

Nathanael Chambers, Taylor Cassidy, Bill McDowell, and Steven Bethard. 2014. Dense event ordering with a multi-pass architecture. *Transactions of the Association for Computational Linguistics*, 2:273–284.

Nathanael Chambers and Daniel Jurafsky. 2008. [Jointly combining implicit constraints improves temporal ordering](#). In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 698–706, Honolulu, Hawaii. Association for Computational Linguistics.

Angel X Chang and Christopher D Manning. 2012. Suntime: A library for recognizing and normalizing time expressions. In *Lrec*, volume 2012, pages 3735–3740.

Fei Cheng and Yusuke Miyao. 2017. [Classifying temporal relations by bidirectional lstm over dependency paths](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1–6, Vancouver, Canada. Association for Computational Linguistics.

Pascal Denis and Philippe Muller. 2011. Predicting globally-coherent temporal structures from texts via endpoint inference and graph decomposition. In *Twenty-Second International Joint Conference on Artificial Intelligence*.

Quang Do, Wei Lu, and Dan Roth. 2012. [Joint inference for event timeline construction](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 677–687, Jeju Island, Korea. Association for Computational Linguistics.

Oleksandr Kolomiyets, Steven Bethard, and Marie-Francine Moens. 2012. Extracting narrative timelines as temporal dependency structures. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pages 88–97. Association for Computational Linguistics.

- Hector Llorens, Nathanael Chambers, Naushad UzZaman, Nasrin Mostafazadeh, James Allen, and James Pustejovsky. 2015. [Semeval-2015 task 5: Qa tempeval - evaluating temporal information understanding with question answering](#). In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 792–800, Denver, Colorado. Association for Computational Linguistics.
- Hector Llorens, Estela Saquete, and Borja Navarro. 2010. [Tipsem \(english and spanish\): Evaluating crfs and semantic roles in tempeval-2](#). In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pages 284–291, Uppsala, Sweden. Association for Computational Linguistics.
- Paramita Mirza and Sara Tonelli. 2016. [On the contribution of word embeddings to temporal relation classification](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2818–2828, Osaka, Japan. The COLING 2016 Organizing Committee.
- Qiang Ning, Zhili Feng, and Dan Roth. 2017. [A structured learning approach to temporal relation extraction](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1027–1037, Copenhagen, Denmark. Association for Computational Linguistics.
- Qiang Ning, Zhili Feng, Hao Wu, and Dan Roth. 2018a. [Joint reasoning for temporal and causal relations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2278–2288, Melbourne, Australia. Association for Computational Linguistics.
- Qiang Ning, Hao Wu, and Dan Roth. 2018b. [A multi-axis annotation scheme for event temporal relations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1318–1328, Melbourne, Australia. Association for Computational Linguistics.
- James Pustejovsky, José M Castano, Robert Ingria, Roser Sauri, Robert J Gaizauskas, Andrea Setzer, Graham Katz, and Dragomir R Radev. 2003. [Timeml: Robust specification of event and temporal expressions in text](#). *New directions in question answering*, 3:28–34.
- James Pustejovsky, Patrick Hanks, Roser Sauri, Andrew See, Robert Gaizauskas, Andrea Setzer, Dragomir Radev, Beth Sundheim, David Day, Lisa Ferro, et al. The timebank corpus.
- Nils Reimers, Nazanin Dehghani, and Iryna Gurevych. 2016. [Temporal anchoring of events for the timebank corpus](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2195–2204, Berlin, Germany. Association for Computational Linguistics.
- Nils Reimers, Nazanin Dehghani, and Iryna Gurevych. 2018. [Event time extraction with a decision tree of neural classifiers](#). *Transactions of the Association for Computational Linguistics*, 6:77–89.
- Andrea Setzer. 2002. *Temporal information in newswire articles: an annotation scheme and corpus study*. Ph.D. thesis, University of Sheffield.
- Zhiyi Song, Ann Bies, Stephanie Strassel, Tom Riese, Justin Mott, Joe Ellis, Jonathan Wright, Seth Kulick, Neville Ryant, and Xiaoyi Ma. 2015. [From light to rich ere: Annotation of entities, relations, and events](#). In *Proceedings of the 3rd Workshop on EVENTS at the NAACL-HLT 2015*, pages 89–98, Denver, Colorado. Association for Computational Linguistics.
- Jannik Strötgen and Michael Gertz. 2010. [HeidelTime: High quality rule-based extraction and normalization of temporal expressions](#). In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pages 321–324, Uppsala, Sweden. Association for Computational Linguistics.
- Naushad UzZaman and James Allen. 2010. [TRIPS and TRIOS system for TempEval-2: Extracting temporal information from text](#). In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pages 276–283, Uppsala, Sweden. Association for Computational Linguistics.
- Naushad UzZaman, Hector Llorens, Leon Derczynski, James Allen, Marc Verhagen, and James Pustejovsky. 2013. [Semeval-2013 task 1: Tempeval-3: Evaluating time expressions, events, and temporal relations](#). In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 1–9, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Marc Verhagen, Robert Gaizauskas, Frank Schilder, Mark Hepple, Graham Katz, and James Pustejovsky. 2007. [Semeval-2007 task 15: Tempeval temporal relation identification](#). In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, pages 75–80, Prague, Czech Republic. Association for Computational Linguistics.
- Marc Verhagen, Roser Sauri, Tommaso Caselli, and James Pustejovsky. 2010. [Semeval-2010 task 13: Tempeval-2](#). In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pages 57–62, Uppsala, Sweden. Association for Computational Linguistics.
- Yan Xu, Lili Mou, Ge Li, Yunchuan Chen, Hao Peng, and Zhi Jin. 2015. [Classifying relations via long short term memory networks along shortest dependency paths](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1785–1794, Lisbon, Portugal. Association for Computational Linguistics.

Appendix

A Coreference Examples

- **Coreferent example:** In the example "their disputes have been **bedeviled** by a number of **disputes**", the event "disputes" is itself the entity enacting the event "bedeviled". The events take place over the same time period and location, and are both real events. Thus, we can conclude the events are coreferent.
- **Non-coreferent example:** In "lower rates have **helped** invigorate housing by **making** loans more affordable", though the events share an agent ("lower rates") and realises states, they act on different patient entities and thus are not coreferent.

B Sample heuristic rules from SS link inference procedure:

Assume S1 and S2 indicate the points associated with events 1 and 2 which are to be linked. Following subsections provide a brief sample of some of the heuristic rules we develop to infer the temporal link based on the values of S1 and S2.

B.1 S1 is of type MM-DD-YYYY and S2 is of type afterMM-DD-YYYY

- Get the relation (rel) between the date values from S1 and S2
- If rel is simultaneous or before, the SS link value is before
- Else skip this link

B.2 S1 is of type MM-DD-YYYY and S2 is of type beforeMM-DD-YYYY

- Get the relation (rel) between the date values from S1 and S2
- If rel is simultaneous or after, the SS link value is after
- Else skip this link

B.3 S1 is of type MM-DD-YYYY and S2 is of type afterMM-DD-YYYY beforeMM-DD-YYYY

- From S2, the date associated with after is named date1 and the date associated with before is named date2
- Get the relation (rel1) between date value from S1 and date1 from S2

- If rel1 is simultaneous or before, the SS link value is before
- Get the relation (rel2) between date value from S1 and date2 from S2
- If rel2 is simultaneous or after, the SS link value is after
- Else skip this link

We develop similar rules for the remaining 13 cases. We also develop rule-based inference procedures for SM and MM links. Please refer to the autogeneration code for the complete set of rules.

Real Life Application of a Question Answering System Using BERT Language Model

Francesca Alloatti^{1,2}, Luigi Di Caro², Gianpiero Sportelli¹

¹CELI - Language Technology, Italy

²Department of Computer Science - Università degli Studi di Torino, Italy
{francesca.alloatti, gianpiero.sportelli}@celi.it
luigi.dicaro@unito.it

Abstract

Real life scenarios are often left untouched by the newest advances in research. They usually require the resolution of some specific task applied to a restricted domain, all the while providing small amounts of data to begin with. In this study we apply one of the newest innovations in Deep Learning to a task of text classification. The goal is to create a question answering system in Italian that provides information about a specific subject, e-invoicing and digital billing. Italy recently introduced a new legislation about e-invoicing and people have some legit doubts, therefore a large share of professionals could benefit from this tool. We gathered few pairs of question and answers; afterwards, we expanded the data, using it as a training corpus for BERT language model. Through a separate test corpus we evaluated the accuracy of the answer provided. Values show that the automatic system alone performs surprisingly well. The demo interface is hosted on Telegram, which makes the system immediately available to test.

1 Introduction

Pre-trained models have proven to be of great help in accomplishing many NLP tasks, such as natural language inference, text classification and question-answering. All of these paradigms contain a semi-supervised language model trained on large corpora of data; they are later fine-tuned to work on downstream tasks (Peters et al., 2018; Howard and Ruder, 2018; Radford et al., 2018). However, real life applications can't often benefit from these advances, for many reasons: lack of data, lack of time and resources to reach a sufficient accuracy level, or the need to address some very specific domain that elude the scope of a general-purpose architecture. As a result, many concrete scenarios of applications are left untouched by the scientific progress, even though

these obstacles are far from impossible to overcome.

The goal of this study is to build a question-answering systems using only BERT (Bidirectional Encoder Representations from Transformers) language model (Devlin et al., 2018), without exploiting any rule-based refinement system or any other proprietary algorithm. This process allows to prevail over the obstacles previously listed: scarce original data was expanded mostly by using generative grammars; the whole project (data expansion plus the various training phases) took no more than eight days to complete, and the computational resources required were fairly affordable¹. Moreover, the application domain is very specific, such that the fine-tuning of the linguistic model significantly increased the performances². The architecture is simple yet effective (as shown in Figure 1) and the output of the system can be tested immediately through a Telegram bot.

2 Related Works

Since its first appearance, BERT has gained a lot of popularity in the academic community. It has been applied to various NLP tasks, including text classification for question answering. The original work by Devlin et al. (2018) contained results on BERT's performance over the Stanford Question Answering Dataset task (Rajpurkar et al., 2016), where the system had to predict the answer span for a specific question in a Wikipedia passage. Yang et al. (2019) went further, creating a question answering system deployed as a chatbot. However, both these studies tackled the task of open-domain question answering, while we focus on cases where BERT was exploited to develop systems for real life applications. For instance,

¹CPU 8 core, GPU 28 GB, RAM 32 GB

²See the Results section for details on the performance.




Figure 1: Architecture of our question answering system

Lee et al. (2019) created a new BERT language model pre-trained on the biomedical field to solve domain-specific text mining tasks (BioBERT). Its results are impressive, but BioBERT is capable to perform well on domain specific knowledge because of its large pre-training process. While the pre-training surely yields better performances, it is highly expensive with regard to computational costs and time consumption. Our results show good performances even without any pre-training.

JESSI, a Joint Encoders for Stable Suggestion Inference (Park et al., 2019), was created upon the knowledge that BERT is severely unstable for out-of-domain samples. This is true for every system that does not implement any other tool other than the language model, such as ours. To solve this problem, Park et al. (2019) combined BERT with a non-BERT encoder and used a RNN classifier on top of BERT. In our case, an heuristic could be applied to the answers given by the system. It would allow to maximize the probability that the output is the correct match and not solely the one with a higher confidence score.

Other studies focus on generating pre-trained embeddings for specific domains (Beltagy et al., 2019; Alsentzer et al., 2019), but they do not test them on specific tasks.

3 BERT's Head start as a Language Model

BERT's architecture is built as a multi-layer bidirectional encoder and it is based on the Transformer model originally proposed by Vaswani et al. (2017). Although BERT has been widely used in the past year, it is not the only tool available to automatically build a working question-answering system. Attention based RNN models, especially with the addition of a LSTM or GRU module, have yielded good results on a variety of tasks (Wang et al., 2016; Zhou et al., 2016). The use of a recur-

rent neural network for our work was eventually ruled out for two main reasons: first, BERT encoder architecture is already trained to work as a language model on more than 104 languages (including Italian) and needs to be refined only for the specific task of text classification. The training of a RNN needs to be done for both the language model creation and the fine-tuning part, which requires a higher volume of data.

Second, the RNN training activities cannot be carried out simultaneously due to network constraints. This causes a more time-consuming and costly process.

4 Data and Fine-tuning Process

Since this aims to be a real life application, the chosen domain was e-invoicing and all the new regulations revolving around the theme of digital billing that was recently introduced by the Italian legislation. The field is very technical and specific; the data needed to reflect the features of the language employed to discuss such subject.

We first gathered pairs of clauses coherent to the domain. The data was cleaned from duplicated questions and badly written sentences, resulting in a corpus of approximately 300 pairs of sentences (a question and an answer). Half of the questions was expanded manually, while the other half - that presented recurrent linguistic patterns - was expanded using generative grammars. A grammar is written as follows:

```
{vb_might} {vb_collect} an
{n_invoice} ?
```

Resulting in sentences such as *Is it possible to collect an e-invoice?* together with all its meaningful variations. The two expansion methods created a corpus of more than 210.000 sentence pairs. No expansion was operated on the answers, since the goal is to match the correct answer to any possible expression of a question, and not to produce variegated answers.

Separately, a different corpus of 200 questions was obtained on a voluntary base from people who did not take part in the expansion process (otherwise, they would have had knowledge of the existing sentences in the training corpus). This distinct, unbiased corpus served as a test set.

4.1 The Fine-tuning

During the fine-tuning process the goal is to expand the network architecture and to train it to-

wards a specific task. A new layer is created while the weights of the underlying original layer are modified according to the text classification job. The final training corpus consisted of 2300 sentences (obtained from the 210.000 previously mentioned). This number resulted from balancing the total of manually expanded questions with the automatically expanded ones. Otherwise we would have had an overfitting problem, since the automatic expansion generates way more sentences than the manual one. Afterwards, we used the test corpus to verify the output of the new network.

Values such as accuracy, precision and recall are not taken into consideration during the training process. Instead, the goal is to optimize, i.e. minimize, a loss function. For this study the loss function is a Cosine Proximity (1). To compute it we created a One Hot Vector that represented the 300 original sentences - each one of them as a label. The loss function takes into account two values: the One Hot Vector and the logarithm of the network output's softmax.

$$L = -\frac{y \cdot \hat{y}}{\|y\|_2 \cdot \|\hat{y}\|_2} = -\frac{\sum_{i=1}^N y_i \cdot \hat{y}_i}{\sqrt{\sum_{i=1}^N y_i^2} \cdot \sqrt{\sum_{i=1}^N \hat{y}_i^2}} \quad (1)$$

Cosine Proximity Loss function

Each experimental round takes approximately one hour.

5 Results

To assess the performance, different experiments were conducted in a subsequent way to evaluate the accuracy of the test set. The first attempts were considered baseline for the following ones. When BERT model was used without applying any fine-tuning the accuracy reached 3,6 % for 40 epochs. Fine-tuning proved to be essential: accuracy on the first answer selected by the system is 86%. When considering the first three answers, the value rises up to 93,6%. The most recent experiment operates on the pre-trained language model too see if further improvement could be reached on that front. The language model was trained with new data extracted from reliable sources (operational handbooks from the *Italian Fiscal Agency*) and

later fine-tuned with the same data of the previous trial. Accuracy gained +2 points, achieving 88 % on first answer.

We also compared our results to other intent matching systems such as Google DialogFlow. Using external API for intent detection accuracy reached 84%, which is slightly lower to our first experiment.

An example of the matched question (and its answer) is presented in Table 1. The user can give a feedback on each answer received, and the positive or negative feedback will add up to constantly improve the performance for the next questions. The average time to obtain a single answer is 0.2 seconds on a CPU architecture. It is therefore perfectly viable for a real time employ as a question answering system.

Unfortunately, it is impossible to compare these results with those obtained from other studies, because of the specificity of this domain, which has never been considered in this kind of experiments (at least for the Italian language).

6 Conclusion and Next Steps

We have demonstrated that it is possible to create a question answering system in a few days. The human effort was minimized - no rule was handwritten and no other algorithm was implemented - and overall the computational cost was bearable. We also showed that scarce data is not always an insurmountable obstacle, since the expansion effort can be split between manual work and automatic one. The results show that such a system can already be used with a decent degree of success. In the next future some improvements are going to be made regarding the context management and the comparison between BERT and other tools.

To improve the spectrum of questions that are correctly matched, we propose two ways to manage the dialog context using BERT:

- **External operation.** The context is given as an external factor to the model through the writing of specific rules. It modifies the labeling, i.e. the probability assigned to a label that selects a matching question.
- **Internal operation.** In this case, BERT needs to be trained towards two inputs, where one is always the context. The network changes its way of calculating the probability from $p(lt)$ (l being the label and t the text

Type of Sentence	Content
<i>Question posed</i>	Hello, I have a question: do I have to issue an invoice also for private clients? Even though they don't refer to any VAT number?
<i>Question matched</i>	Does an invoice need to be issued also towards people without a VAT number?
<i>Answer provided</i>	Yes, the electronic document has to be issued towards private clients without a VAT number

Table 1: Given a certain question posed by a user, the system matches one of the example in his knowledge bases and sends out the correct answer. The sentences have been translated from Italian into English for the purpose of this paper.

of the sentence) to $p(I \cap c)$.

Regarding the other tools, our goal is to verify if other models could perform equally or better given the same dataset. Many platforms are currently under review, such as Amazon Lex.

References

- Emily Alsentzer, John R. Murphy, Willie Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, and Matthew B. A. McDermott. 2019. [Publicly available clinical bert embeddings](#). *arXiv preprint*, arXiv:1904.03323.
- Iz Beltagy, Arman Cohan, and Kyle Lo. 2019. [Scibert: Pretrained contextualized embeddings for scientific text](#). *arXiv preprint*, arXiv:1903.10676.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). *arXiv preprint*, arXiv:1810.04805v1.
- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, pages 328–339.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2019. [Biobert: a pre-trained biomedical language representation model for biomedical text mining](#). *arXiv preprint*, arXiv:1901.08746v3.
- Cheoneum Park, Juae Kim, Hyeon gu Lee, Reinald Kim Amplayo, Harksoo Kim, Jungyun Seo, and Changki Lee. 2019. [This is competition at semeval-2019 task 9: Bert is unstable for out-of-domain samples](#). *arXiv preprint*, arXiv:1904.03339v1.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, , and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). *NAACL*.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. *OpenAI preprint*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [Squad: 100, 000+ questions for machine comprehension of text](#). *CoRR*, abs/1606.05250.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Lukasz Kaiser Aidan N Gomez, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 6000–6010.
- Bingning Wang, Kang Liu, and Jun Zhao. 2016. Inner attention based recurrent neural networks for answer selection. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 1288–1297. Association for Computational Linguistics.
- Wei Yang, Yuqing Xie, Aileen Lin, Xingyu Li, Luchen Tan, Kun Xiong, Ming Li, and Jimmy Lin. 2019. [End-to-end open-domain question answering with bertserini](#). *CoRR*, abs/1902.01718.
- Xinjie Zhou, Xiaojun Wanand, and Jianguo Xiao. 2016. Attention-based lstm network for cross-lingual sentiment classification. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 247–256. Association for Computational Linguistics.

Hierarchical Multi-Task Natural Language Understanding for Cross-domain Conversational AI: HERMIT NLU

Andrea Vanzo
Interaction Lab
Heriot-Watt University
a.vanzo@hw.ac.uk

Emanuele Bastianelli
Interaction Lab
Heriot-Watt University
e.bastianelli@hw.ac.uk

Oliver Lemon
Interaction Lab
Heriot-Watt University
o.lemon@hw.ac.uk

Abstract

We present a new neural architecture for wide-coverage Natural Language Understanding in Spoken Dialogue Systems. We develop a hierarchical multi-task architecture, which delivers a multi-layer representation of sentence meaning (i.e., Dialogue Acts and Frame-like structures). The architecture is a hierarchy of self-attention mechanisms and BiLSTM encoders followed by CRF tagging layers. We describe a variety of experiments, showing that our approach obtains promising results on a dataset annotated with Dialogue Acts and Frame Semantics. Moreover, we demonstrate its applicability to a different, publicly available NLU dataset annotated with domain-specific intents and corresponding semantic roles, providing overall performance higher than state-of-the-art tools such as RASA, Dialogflow, LUIS, and Watson. For example, we show an average 4.45% improvement in entity tagging F-score over Rasa, Dialogflow and LUIS.

1 Introduction

Research in Conversational AI (also known as Spoken Dialogue Systems) has applications ranging from home devices to robotics, and has a growing presence in industry. A key problem in real-world Dialogue Systems is Natural Language Understanding (NLU) – the process of extracting structured representations of meaning from user utterances. In fact, the effective extraction of semantics is an essential feature, being the entry point of any Natural Language interaction system. Apart from challenges given by the inherent complexity and ambiguity of human language, other challenges arise whenever the NLU has to operate over multiple domains. In fact, interaction patterns, domain, and language vary depending on the device the user is interacting with. For example, chit-chatting and instruction-giving for executing

an action are different processes in terms of language, domain, syntax and interaction schemes involved. And what if the user combines two interaction domains: “*play some music, but first what’s the weather tomorrow*”?

In this work, we present HERMIT, a HiERarchical Multi-Task Natural Language Understanding architecture¹, designed for effective semantic parsing of domain-independent user utterances, extracting meaning representations in terms of high-level intents and frame-like semantic structures. With respect to previous approaches to NLU for SDS, HERMIT stands out for being a cross-domain, multi-task architecture, capable of recognising multiple intents/frames in an utterance. HERMIT also shows better performance with respect to current state-of-the-art commercial systems. Such a novel combination of requirements is discussed below.

Cross-domain NLU A cross-domain dialogue agent must be able to handle heterogeneous types of conversation, such as chit-chatting, giving directions, entertaining, and triggering domain/task actions. A domain-independent and rich meaning representation is thus required to properly capture the intent of the user. Meaning is modelled here through three layers of knowledge: dialogue acts, frames, and frame arguments. Frames and arguments can be in turn mapped to domain-dependent intents and slots, or to Frame Semantics’ (Fillmore, 1976) structures (i.e. semantic frames and frame elements, respectively), which allow handling of heterogeneous domains and language.

Multi-task NLU Deriving such a multi-layered meaning representation can be approached through a multi-task learning approach. Multi-task learning has found success in several NLP

¹<https://gitlab.com/hwu-ilab/hermit-nlu>

problems (Hashimoto et al., 2017; Strubell et al., 2018), especially with the recent rise of Deep Learning. Thanks to the possibility of building complex networks, handling more tasks at once has been proven to be a successful solution, provided that some degree of dependence holds between the tasks. Moreover, multi-task learning allows the use of different datasets to train sub-parts of the network (Sanh et al., 2018). Following the same trend, HERMIT is a hierarchical multi-task neural architecture which is able to deal with the three tasks of tagging dialogue acts, frame-like structures, and their arguments in parallel. The network, based on self-attention mechanisms, seq2seq bi-directional Long-Short Term Memory (BiLSTM) encoders, and CRF tagging layers, is hierarchical in the sense that information output from earlier layers flows through the network, feeding following layers to solve downstream dependent tasks.

Multi-dialogue act and -intent NLU Another degree of complexity in NLU is represented by the granularity of knowledge that can be extracted from an utterance. Utterance semantics is often rich and expressive: approximating meaning to a single user intent is often not enough to convey the required information. As opposed to the traditional single-dialogue act and single-intent view in previous work (Guo et al., 2014; Liu and Lane, 2016; Hakkani-Tur et al., 2016), HERMIT operates on a meaning representation that is multi-dialogue act and multi-intent. In fact, it is possible to model an utterance’s meaning through multiple dialogue acts and intents at the same time. For example, the user would be able both to request tomorrow’s weather and listen to his/her favourite music with just a single utterance.

A further requirement is that for practical application the system should be **competitive with state-of-the-art**: we evaluate HERMIT’s effectiveness by running several empirical investigations. We perform a robust test on a publicly available NLU-Benchmark (NLU-BM) (Liu et al., 2019) containing 25K cross-domain utterances with a conversational agent. The results obtained show a performance higher than well-known off-the-shelf tools (i.e., Rasa, DialogueFlow, LUIS, and Watson). The contribution of the different network components is then highlighted through an ablation study. We also test HERMIT on the smaller

Robotics-Oriented MULTITask Language Understanding (ROMULUS) corpus, annotated with Dialogue Acts and Frame Semantics. HERMIT produces promising results for the application in a real scenario.

2 Related Work

Much research on Natural (or Spoken, depending on the input) Language Understanding has been carried out in the area of Spoken Dialogue Systems (Chen et al., 2017), where the advent of statistical learning has led to the application of many data-driven approaches (Lemon and Pietquin, 2012). In recent years, the rise of deep learning models has further improved the state-of-the-art. Recurrent Neural Networks (RNNs) have proven to be particularly successful, especially uni- and bi-directional LSTMs and Gated Recurrent Units (GRUs). The use of such deep architectures has also fostered the development of joint classification models of intents and slots. Bi-directional GRUs are applied in (Zhang and Wang, 2016), where the hidden state of each time step is used for slot tagging in a seq2seq fashion, while the final state of the GRU is used for intent classification. The application of attention mechanisms in a BiLSTM architecture is investigated in (Liu and Lane, 2016), while the work of (Chen et al., 2016) explores the use of memory networks (Sukhbaatar et al., 2015) to exploit encoding of historical user utterances to improve the slot-filling task. Seq2seq with self-attention is applied in (Li et al., 2018), where the classified intent is also used to guide a special gated unit that contributes to the slot classification of each token.

One of the first attempts to jointly detect domains in addition to intent-slot tagging is the work of (Guo et al., 2014). An utterance syntax is encoded through a Recursive NN, and it is used to predict the joined domain-intent classes. Syntactic features extracted from the same network are used in the per-word slot classifier. The work of (Hakkani-Tur et al., 2016) applies the same idea of (Zhang and Wang, 2016), this time using a context-augmented BiLSTM, and performing domain-intent classification as a single joint task. As in (Chen et al., 2016), the history of user utterances is also considered in (Bapna et al., 2017), in combination with a dialogue context encoder. A two-layer hierarchical structure made of a combination of BiLSTM and BiGRU is used

for joint classification of domains and intents, together with slot tagging. (Rastogi et al., 2018) apply multi-task learning to the dialogue domain. Dialogue state tracking, dialogue act and intent classification, and slot tagging are jointly learned. Dialogue states and user utterances are encoded to provide hidden representations, which jointly affect all the other tasks.

Many previous systems are trained and compared over the ATIS (Airline Travel Information Systems) dataset (Price, 1990), which covers only the flight-booking domain. Some of them also use bigger, not publicly available datasets, which appear to be similar to the NLU-BM in terms of number of intents and slots, but they cover no more than three or four domains. Our work stands out for its more challenging NLU setting, since we are dealing with a higher number of domains/scenarios (18), intents (64) and slots (54) in the NLU-BM dataset, and dialogue acts (11), frames (58) and frame elements (84) in the ROMULUS dataset. Moreover, we propose a multi-task hierarchical architecture, where each layer is trained to solve one of the three tasks. Each of these is tackled with a seq2seq classification using a CRF output layer, as in (Sanh et al., 2018).

The NLU problem has been studied also on the Interactive Robotics front, mostly to support basic dialogue systems, with few dialogue states and tailored for specific tasks, such as semantic mapping (Kruijff et al., 2007), navigation (Kollar et al., 2010; Bothe et al., 2018), or grounded language learning (Chai et al., 2016). However, the designed approaches, either based on formal languages or data-driven, have never been shown to scale to real world scenarios. The work of (Hatori et al., 2018) makes a step forward in this direction. Their model still deals with the single ‘pick and place’ domain, covering no more than two intents, but it is trained on several thousands of examples, making it able to manage more unstructured language. An attempt to manage a higher number of intents, as well as more variable language, is represented by the work of (Bastianelli et al., 2016) where the sole Frame Semantics is applied to represent user intents, with no Dialogue Acts.

3 Jointly parsing dialogue acts and frame-like structures

The identification of Dialogue Acts (henceforth DAs) is required to drive the dialogue manager

to the next dialogue state. General frame structures (FRs) provide a reference framework to capture user intents, in terms of required or desired actions that a conversational agent has to perform. Depending on the level of abstraction required by an application, these can be interpreted as more domain-dependent paradigms like *intent*, or to shallower representations, such as *semantic frames*, as conceived in FrameNet (Baker et al., 1998). From this perspective, semantic frames represent a versatile abstraction that can be mapped over an agent’s capabilities, allowing also the system to be easily extended with new functionalities without requiring the definition of new ad-hoc structures. Similarly, frame arguments (ARs) act as *slots* in a traditional intent-slots scheme, or to *frame elements* for semantic frames.

In our work, the whole process of extracting a complete semantic interpretation as required by the system is tackled with a multi-task learning approach across DAs, FRs, and ARs. Each of these tasks is modelled as a seq2seq problem, where a task-specific label is assigned to each token of the sentence according to the IOB2 notation (Sang and Veenstra, 1999), with “B-” marking the Beginning of the chunk, “I-” the tokens Inside the chunk while “O-” is assigned to any token that does not belong to any chunk. Task labels are drawn from the set of classes defined for DAs, FRs, and ARs. Figure 1 shows an example of the tagging layers over the sentence *Where can I find Starbucks?*, where Frame Semantics has been selected as underlying reference theory.

3.1 Architecture description

The central motivation behind the proposed architecture is that there is a dependence among the three tasks of identifying DAs, FRs, and ARs. The relationship between tagging frame and arguments appears more evident, as also developed in theories like Frame Semantics – although it is defined independently by each theory. However, some degree of dependence also holds between the DAs and FRs. For example, the FrameNet semantic frame *Desiring*, expressing a desire of the user for an event to occur, is more likely to be used in the context of an INFORM DA, which indicates the state of notifying the agent with an information, other than in an INSTRUCTION. This is clearly visible in interactions like “*I’d like a cup of hot chocolate*” or “*I’d like to find a shoe shop*”, where

	<i>Where</i>	<i>can</i>	<i>I</i>	<i>find</i>	<i>Starbucks</i>	<i>?</i>
DAs	B-REQ_INFO	I-REQ_INFO	I-REQ_INFO	I-REQ_INFO	I-REQ_INFO	O
FRs	B-Locating	I-Locating	I-Locating	I-Locating	I-Locating	O
ARs	O	O	B-COGNIZER	B-LEXICAL_UNIT	B-ENTITY	O

Figure 1: Dialogue Acts (DAs), Frames (FRs – here semantic frames) and Arguments (ARs – here frame elements) IOB2 tagging for the sentence *Where can I find Starbucks?*

the user is actually notifying the agent about a desire of hers/his.

In order to reflect such inter-task dependence, the classification process is tackled here through a hierarchical multi-task learning approach. We designed a multi-layer neural network, whose architecture is shown in Figure 2, where each layer is trained to solve one of the three tasks, namely labelling dialogue acts (*DA* layer), semantic frames (*FR* layer), and frame elements (*AR* layer). The layers are arranged in a hierarchical structure that allows the information produced by earlier layers to be fed to downstream tasks.

The network is mainly composed of three BiLSTM (Schuster and Paliwal, 1997) encoding layers. A sequence of input words is initially converted into an embedded representation through an ELMo embeddings layer (Peters et al., 2018), and is fed to the *DA* layer. The embedded representation is also passed over through shortcut connections (Hashimoto et al., 2017), and concatenated with both the outputs of the *DA* and *FR* layers. Self-attention layers (Zheng et al., 2018) are placed after the *DA* and *FR* BiLSTM encoders. Where w_t is the input word at time step t of the sentence $\mathbf{w} = (w_1, \dots, w_T)$, the architecture can be formalised by:

$$\begin{aligned}
e_t &= ELMo(w_t), \quad s_t^{DA} = BiLSTM(e_t) \\
a_t^{DA} &= SelfAtt(s_t^{DA}, \mathbf{s}^{DA}), \\
s_t^{FR} &= BiLSTM(e_t \oplus a_t^{DA}), \\
a_t^{FR} &= SelfAtt(s_t^{FR}, \mathbf{s}^{FR}), \\
s_t^{AR} &= BiLSTM(e_t \oplus a_t^{FR})
\end{aligned}$$

where \oplus represents the vector concatenation operator, e_t is the embedding of the word at time t , and $\mathbf{s}^L = (s_1^L, \dots, s_T^L)$ is the embedded sequence output of each L layer, with $L = \{DA, FR, AR\}$. Given an input sentence, the final sequence of labels \mathbf{y}^L for each task is computed through a CRF tagging layer, which operates on the output of the *DA* and *FR* self-attention, and of the *AR* BiLSTM em-

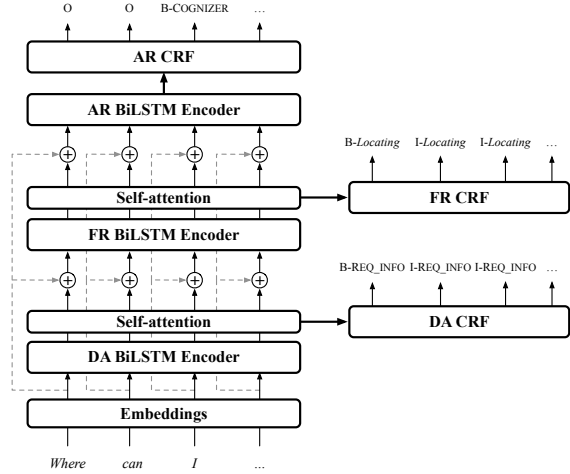


Figure 2: HERMIT Network topology

bedding, so that:

$$\begin{aligned}
\mathbf{y}^{DA} &= CRF^{DA}(\mathbf{a}^{DA}), \quad \mathbf{y}^{FR} = CRF^{FR}(\mathbf{a}^{FR}) \\
\mathbf{y}^{AR} &= CRF^{AR}(\mathbf{s}^{AR}),
\end{aligned}$$

where \mathbf{a}^{DA} , \mathbf{a}^{FR} are attended embedded sequences. Due to shortcut connections, layers in the upper levels of the architecture can rely both on direct word embeddings as well as the hidden representation a_t^L computed by a previous layer. Operationally, the latter carries task specific information which, combined with the input embeddings, helps in stabilising the classification of each CRF layer, as shown by our experiments. The network is trained by minimising the sum of the individual negative log-likelihoods of the three CRF layers, while at test time the most likely sequence is obtained through the Viterbi decoding over the output scores of the CRF layer.

4 Experimental Evaluation

In order to assess the effectiveness of the proposed architecture and compare against existing off-the-shelf tools, we run several empirical evaluations.

4.1 Datasets

We tested the system on two datasets, different in size and complexity of the addressed language.

NLU-Benchmark dataset The first (publicly available) dataset, NLU-Benchmark (NLU-BM), contains 25,716 utterances annotated with targeted *Scenario*, *Action*, and involved *Entities*. For example, “*schedule a call with Lisa on Monday morning*” is labelled to contain a `calendar_scenario`, where the `set_event` action is instantiated through the entities [`event_name: a call with Lisa`] and [`date: Monday morning`]. The Intent is then obtained by concatenating scenario and action labels (e.g., `calendar_set_event`). This dataset consists of multiple home assistant task domains (e.g., scheduling, playing music), chit-chat, and commands to a robot (Liu et al., 2019).²

	NLU-BM	NLU-BM (reduced)
Sentences	25715	11020
Sentences length	7.06	6.84
Scenario labels set	18	18
Action labels set	54	51
Intent labels set	68	64
Entity labels set	56	54
Number of intent	25715	11020
Number of entities	20597	9130
Intents/sentence	1	1
Entities/sentence	0.8	0.83

Table 1: Statistics of the NLU-Benchmark dataset (Liu et al., 2019).

ROMULUS dataset The second dataset, ROMULUS, is composed of 1,431 sentences, for each of which dialogue acts, semantic frames, and corresponding frame elements are provided. This dataset is being developed for modelling user utterances to open-domain conversational systems for robotic platforms that are expected to handle different interaction situations/patterns – e.g., chit-chat, command interpretation. The corpus is composed of different subsections, addressing heterogeneous linguistic phenomena, ranging from imperative instructions (e.g., “*enter the bedroom slowly, turn left and turn the lights off*”) to complex requests for information (e.g., “*good morning I want to buy a new mobile phone is there any shop nearby?*”) or open-domain chit-chat (e.g., “*nope thanks let’s talk about cinema*”). A considerable number of utterances in the dataset is collected through Human-Human Interaction studies in robotic domain ($\approx 70\%$), though a small portion has been synthetically generated for balancing the frame distribution.

²Available at <https://github.com/xliuhw/NLU-Evaluation-Data>.

ROMULUS dataset	
Sentences	1431
Sentences length	7.24
Dialogue act labels set	11
Frame labels set	58
Frame element labels set	84
Number of dialogue acts	1906
Number of frames	2013
Number of frame elements	5059
Dialogue act/sentence	1.33
Frames/sentence	1.41
Frame elements/sentence	3.54

Table 2: Statistics of the ROMULUS dataset.

Note that while the NLU-BM is designed to have at most one intent per utterance, sentences are here tagged following the IOB2 sequence labelling scheme (see example of Figure 1), so that multiple dialogue acts, frames, and frame elements can be defined at the same time for the same utterance. For example, three dialogue acts are identified within the sentence [*good morning*]_{OPENING} [*I want to buy a new mobile phone*]_{INFORM} [*is there any shop nearby?*]_{REQ-INFO}. As a result, though smaller, the ROMULUS dataset provides a richer representation of the sentence’s semantics, making the tasks more complex and challenging. These observations are highlighted by the statistics in Table 2, that show an average number of dialogue acts, frames and frame elements always greater than 1 (i.e., 1.33, 1.41 and 3.54, respectively).

4.2 Experimental setup

All the models are implemented with Keras (Chollet et al., 2015) and Tensorflow (Abadi et al., 2015) as backend, and run on a Titan Xp. Experiments are performed in a 10-fold setting, using one fold for tuning and one for testing. However, since HERMIT is designed to operate on dialogue acts, semantic frames and frame elements, the best hyperparameters are obtained over the ROMULUS dataset via a grid search using early stopping, and are applied also to the NLU-BM models.³ This guarantees fairness towards other systems, that do not perform any fine-tuning on the training data. We make use of pre-trained 1024-dim ELMo embeddings (Peters et al., 2018) as word vector representations without re-training the weights.

³Notice that in the NLU-BM experiments only the number of epochs is tuned, using 10% of the training data.

4.3 Experiments on the NLU-Benchmark

This section shows the results obtained on the NLU-Benchmark (NLU-BM) dataset provided by (Liu et al., 2019), by comparing HERMIT to off-the-shelf NLU services, namely: **Rasa**⁴, **Dialogflow**⁵, **LUIS**⁶ and **Watson**⁷. In order to apply HERMIT to NLU-BM annotations, these have been aligned so that Scenarios are treated as DAs, Actions as FRs and Entities as ARs.

To make our model comparable against other approaches, we reproduced the same folds as in (Liu et al., 2019), where a resized version of the original dataset is used. Table 1 shows some statistics of the NLU-BM and its reduced version. Moreover, micro-averaged Precision, Recall and F1 are computed following the original paper to assure consistency. TP, FP and FN of intent labels are obtained as in any other multi-class task. An entity is instead counted as TP if there is an overlap between the predicted and the gold span, and their labels match.

Experimental results are reported in Table 3. The statistical significance is evaluated through the Wilcoxon signed-rank test. When looking at the intent F1, HERMIT performs significantly better than Rasa [$Z = -2.701, p = .007$] and LUIS [$Z = -2.807, p = .005$]. On the contrary, the improvements w.r.t. Dialogflow [$Z = -1.173, p = .241$] do not seem to be significant. This is probably due to the high variance obtained by Dialogflow across the 10 folds. Watson is by a significant margin the most accurate system in recognising intents [$Z = -2.191, p = .028$], especially due to its Precision score.

The hierarchical multi-task architecture of HERMIT seems to contribute strongly to entity tagging accuracy. In fact, in this task it performs significantly better than Rasa [$Z = -2.803, p = .005$], Dialogflow [$Z = -2.803, p = .005$], LUIS [$Z = -2.803, p = .005$] and Watson [$Z = -2.805, p = .005$], with improvements from 7.08 to 35.92 of F1.⁹

⁴<https://rasa.com/>

⁵<https://dialogflow.com/>

⁶<https://www.luis.ai/>

⁷<https://www.ibm.com/watson>

⁹Results for Watson are shown for the non-contextual training. Due to Watson limitations, i.e. 2000 training examples for contextual training, we could not run the whole test in such configuration. For fairness, we report results made on 8 random samplings of 2000/1000 train/test examples a each (F1): Intent= 72.64 ± 7.46 , Slots= 77.01 ± 10.65 , Combined= 74.85 ± 7.54

Following (Liu et al., 2019), we then evaluated a metric that combines intent and entities, computed by simply summing up the two confusion matrices (Table 4). Results highlight the contribution of the entity tagging task, where HERMIT outperforms the other approaches. Paired-samples t-tests were conducted to compare the HERMIT combined F1 against the other systems. The statistical analysis shows a significant improvement over Rasa [$Z = -2.803, p = .005$], Dialogflow [$Z = -2.803, p = .005$], LUIS [$Z = -2.803, p = .005$] and Watson [$Z = -2.803, p = .005$].

4.3.1 Ablation study

In order to assess the contributions of the HERMIT’s components, we performed an ablation study. The results are obtained on the NLU-BM, following the same setup as in Section 4.3.

Results are shown in Table 5. The first row refers to the complete architecture, while -SA shows the results of HERMIT without the self-attention mechanism. Then, from this latter we further remove shortcut connections (- SA/CN) and CRF taggers (- SA/CRF). The last row (- SA/CN/CRF) shows the results of a simple architecture, without self-attention, shortcuts, and CRF. Though not significant, the contribution of the several architectural components can be observed. The contribution of self-attention is distributed across all the tasks, with a small inclination towards the upstream ones. This means that while the entity tagging task is mostly lexicon independent, it is easier to identify pivoting keywords for predicting the intent, e.g. the verb “*schedule*” triggering the `calendar_set_event` intent. The impact of shortcut connections is more evident on entity tagging. In fact, the effect provided by shortcut connections is that the information flowing throughout the hierarchical architecture allows higher layers to encode richer representations (i.e., original word embeddings + latent semantics from the previous task). Conversely, the presence of the CRF tagger affects mainly the lower levels of the hierarchical architecture. This is not probably due to their position in the hierarchy, but to the way the tasks have been designed. In fact, while the span of an entity is expected to cover few tokens, in intent recognition (i.e., a combination of Scenario and Action recognition) the span always covers all the tokens of an utterance. CRF therefore preserves consistency of IOB2 sequences structure. However, HERMIT seems to be the most stable ar-

	Intent			Entity		
	P	R	F1	P	R	F1
Rasa	86.31±1.07	86.31±1.07	86.31±1.07	85.93±1.05	69.40±1.66	76.78±1.27
Dialogflow	86.97±2.02	85.87±2.33	86.42±2.18	78.21±3.35	70.85±4.70	74.30±3.74
LUIS	85.53±1.14	85.51±1.15	85.52±1.15	83.69±1.31	72.46±2.05	77.66±1.45
Watson ⁸	88.41±0.68	88.08±0.74	88.24±0.70	35.39±0.93	78.70±2.01	48.82±1.14
HERMIT	87.41±0.63	87.70±0.64	87.55±0.63	87.65±0.98	82.04±2.12	84.74±1.18

Table 3: Comparison of HERMIT with the results obtained in (Liu et al., 2019) for Intents and Entity Types.

	P	Combined R	F1
	Rasa	86.16±0.90	78.66±1.28
Dialogflow	83.19±2.43	79.07±3.10	81.07±2.64
LUIS	84.76±0.67	79.61±1.25	82.1±0.90
Watson	54.02±0.75	83.83±1.02	65.7±0.75
HERMIT	87.52±0.61	85.03±1.11	86.25±0.66

Table 4: Comparison of HERMIT with the results in (Liu et al., 2019) by combining Intent and Entity.

	Intent	Entity	Combined
	HERMIT	87.55±0.63	84.74±1.18
- SA	87.03±0.74	84.35±1.15	85.81±0.81
- SA/CN	87.09±0.78	82.43±1.42	84.97±0.72
- SA/CRF	83.57±0.75	84.77±1.06	84.09±0.79
- SA/CN/CRF	83.78±1.10	82.22±1.41	83.10±1.06

Table 5: Ablation study of HERMIT on the NLU-BM.

chitecture, both in terms of standard deviation and task performance, with a good balance between intent and entity recognition.

4.4 Experiments on the ROMULUS dataset

In this section we report the experiments performed on the ROMULUS dataset (Table 6). Together with the evaluation metrics used in (Liu et al., 2019), we report the span F1, computed using the CoNLL-2000 shared task evaluation script, and the Exact Match (EM) accuracy of the entire sequence of labels. It is worth noticing that the EM Combined score is computed as the conjunction of the three individual predictions – e.g., a match is when all the three sequences are correct.

Results in terms of EM reflect the complexity of the different tasks, motivating their position within the hierarchy. Specifically, dialogue act identification is the easiest task (89.31%) with respect to frame (82.60%) and frame element (79.73%), due to the shallow semantics it aims to catch. However, when looking at the span F1, its score (89.42%) is lower than the frame element identification task (92.26%). What happens is that

even though the label set is smaller, dialogue act spans are supposed to be longer than frame element ones, sometimes covering the whole sentence. Frame elements, instead, are often one or two tokens long, that contribute in increasing span based metrics. Frame identification is the most complex task for several reasons. First, lots of frame spans are interlaced or even nested; this contributes to increasing the network entropy. Second, while the dialogue act label is highly related to syntactic structures, frame identification is often subject to the inherent ambiguity of language (e.g., *get* can evoke both *Commerce.buy* and *Arriving*). We also report the metrics in (Liu et al., 2019) for consistency. For dialogue act and frame tasks, scores provide just the extent to which the network is able to detect those labels. In fact, the metrics do not consider any span information, essential to solve and evaluate our tasks. However, the frame element scores are comparable to the benchmark, since the task is very similar.

Overall, getting back to the combined EM accuracy, HERMIT seems to be promising, with the network being able to reproduce all the three gold sequences for almost 70% of the cases. The importance of this result provides an idea of the architecture behaviour over the entire pipeline.

4.5 Discussion

The experimental evaluation reported in this section provides different insights. The proposed architecture addresses the problem of NLU in wide-coverage conversational systems, modelling semantics through multiple Dialogue Acts and Frame-like structures in an end-to-end fashion. In addition, its hierarchical structure, which reflects the complexity of the single tasks, allows providing rich representations across the whole network. In this respect, we can affirm that the architecture successfully tackles the multi-task problem, with results that are promising in terms of usability and applicability of the system in real scenarios.

	P	R	F1	span F1	EM
<i>Dialogue act</i>	96.49±0.98	95.95±1.41	96.21±1.13	89.42±3.74	89.31±3.28
<i>Frame</i>	95.26±0.95	94.02±1.20	94.64±1.09	84.40±2.99	82.60±2.68
<i>Frame element</i>	95.62±0.61	93.98±0.76	94.79±0.56	92.26±1.22	79.73±2.03
Combined	93.90±0.89	92.95±0.86	93.42±0.83	–	69.53±2.50

Table 6: HERMIT performance over the ROMULUS dataset. P,R and F1 are evaluated following (Liu et al., 2019) metrics

However, a thorough evaluation in the wild must be carried out, to assess to what extent the system is able to handle complex spoken language phenomena, such as repetitions, disfluencies, etc. To this end, a real scenario evaluation may open new research directions, by addressing new tasks to be included in the multi-task architecture. This is supported by the scalable nature of the proposed approach. Moreover, following (Sanh et al., 2018), corpora providing different annotations can be exploited within the same multi-task network.

We also empirically showed how the same architectural design could be applied to a dataset addressing similar problems. In fact, a comparison with off-the-shelf tools shows the benefits provided by the hierarchical structure, with better overall performance better than any current solution. An ablation study has been performed, assessing the contribution provided by the different components of the network. The results show how the shortcut connections help in the more fine-grained tasks, successfully encoding richer representations. CRFs help when longer spans are being predicted, more present in the upstream tasks.

Finally, the seq2seq design allowed obtaining a multi-label approach, enabling the identification of multiple spans in the same utterance that might evoke different dialogue acts/frames. This represents a novelty for NLU in conversational systems, as such a problem has always been tackled as a single-intent detection. However, the seq2seq approach carries also some limitations, especially on the Frame Semantics side. In fact, label sequences are linear structures, not suitable for representing nested predicates, a tough and common problem in Natural Language. For example, in the sentence “*I want to buy a new mobile phone*”, the [to buy a new mobile phone] span represents both the DESIRED_EVENT frame element of the *Desiring* frame and a *Commerce.buy* frame at the same time. At the moment of writing, we are working on modeling nested predicates through the application of bilinear models.

5 Future Work

We have started integrating a corpus of 5M sentences of real users chit-chatting with our conversational agent, though at the time of writing they represent only 16% of the current dataset.

As already pointed out in Section 4.5, there are some limitations in the current approach that need to be addressed. First, we have to assess the network’s capability in handling typical phenomena of spontaneous spoken language input, such as repetitions and disfluencies (Shalyminov et al., 2018). This may open new research directions, by including new tasks to identify/remove any kind of noise from the spoken input. Second, the seq2seq scheme does not deal with nested predicates, a common aspect of Natural Language. To the best of our knowledge, there is no architecture that implements an end-to-end network for FrameNet based semantic parsing. Following previous work (Strubell et al., 2018), one of our future goals is to tackle such problems through hierarchical multi-task architectures that rely on bilinear models.

6 Conclusion

In this paper we presented HERMIT NLU, a hierarchical multi-task architecture for semantic parsing sentences for cross-domain spoken dialogue systems. The problem is addressed using a seq2seq model employing BiLSTM encoders and self-attention mechanisms and followed by CRF tagging layers. We evaluated HERMIT on a 25K sentences NLU-Benchmark and outperform state-of-the-art NLU tools such as Rasa, Dialogflow, LUIS and Watson, even without specific fine-tuning of the model.

Acknowledgement

This research was partially supported by the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 688147 (MuMMER project¹⁰).

¹⁰<http://mummer-project.eu/>

References

- Martín Abadi et al. 2015. **TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems**. Software available from tensorflow.org.
- Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet Project. In *Proceedings of ACL and COLING*, Association for Computational Linguistics, pages 86–90.
- Ankur Bapna, Gokhan Tur, Dilek Hakkani-Tur, and Larry Heck. 2017. Sequential dialogue context modeling for spoken language understanding. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 103–114. Association for Computational Linguistics.
- Emanuele Bastianelli, Danilo Croce, Andrea Vanzo, Roberto Basili, and Daniele Nardi. 2016. A discriminative approach to grounded spoken language understanding in interactive robotics. In *Proceedings of the 2016 International Joint Conference on Artificial Intelligence (IJCAI)*, New York, USA.
- Chandrakant Bothe, Fernando García, Arturo Cruz-Maya, Amit Kumar Pandey, and Stefan Wermter. 2018. Towards dialogue-based navigation with multivariate adaptation driven by intention and politeness for social robots. In *ICSR*, volume 11357 of *Lecture Notes in Computer Science*, pages 230–240. Springer.
- Joyce Y Chai, Rui Fang, Changsong Liu, and Lanbo She. 2016. Collaborative language grounding toward situated human-robot dialogue. *AI Magazine*, 37(4):32–45.
- Hongshen Chen, Xiaorui Liu, Dawei Yin, and Jiliang Tang. 2017. A survey on dialogue systems: Recent advances and new frontiers. *SIGKDD Explor. Newsl.*, 19(2):25–35.
- Yun-Nung Chen, Dilek Hakkani-Tür, Gökhan Tür, Jianfeng Gao, and Li Deng. 2016. End-to-end memory networks with knowledge carryover for multi-turn spoken language understanding. In *INTER-SPEECH*, pages 3245–3249. ISCA.
- François Chollet et al. 2015. Keras. <https://keras.io>.
- Charles J Fillmore. 1976. Frame semantics and the nature of language. *Annals of the New York Academy of Sciences*, 280(1):20–32.
- Daniel Guo, Gökhan Tür, Wen-tau Yih, and Geoffrey Zweig. 2014. Joint semantic utterance classification and slot filling with recursive neural networks. In *2014 IEEE Spoken Language Technology Workshop, SLT 2014, South Lake Tahoe, NV, USA, December 7-10, 2014*, pages 554–559.
- Dilek Hakkani-Tur, Gokhan Tur, Asli Celikyilmaz, Yun-Nung Chen, Jianfeng Gao, Li Deng, and Ye-Yi Wang. 2016. Multi-Domain Joint Semantic Frame Parsing using Bi-directional RNN-LSTM. In *Proceedings of Interspeech*.
- Kazuma Hashimoto, Caiming Xiong, Yoshimasa Tsunokawa, and Richard Socher. 2017. **A Joint Many-Task Model: Growing a Neural Network for Multiple NLP Tasks**. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1923–1933. Association for Computational Linguistics.
- Jun Hatori, Yuta Kikuchi, Sosuke Kobayashi, Kuniyuki Takahashi, Yuta Tsuboi, Yuya Unno, Wilson Ko, and Jethro Tan. 2018. Interactively picking real-world objects with unconstrained spoken language instructions. In *2018 IEEE International Conference on Robotics and Automation, ICRA 2018, Brisbane, Australia, May 21-25, 2018*, pages 3774–3781.
- Thomas Kollar, Stefanie Tellex, Deb Roy, and Nicholas Roy. 2010. Toward understanding natural language directions. In *Proceedings of the 5th ACM/IEEE International Conference on Human-robot Interaction, HRI '10*, pages 259–266, Piscataway, NJ, USA. IEEE Press.
- Geert-Jan M. Kruijff, H. Zender, P. Jensfelt, and Henrik I. Christensen. 2007. Situated dialogue and spatial organization: What, where... and why? *International Journal of Advanced Robotic Systems*, 4(2).
- Oliver Lemon and Olivier Pietquin. 2012. *Data-Driven Methods for Adaptive Spoken Dialogue Systems: Computational Learning for Conversational Interfaces*. Springer Publishing Company, Incorporated.
- Changliang Li, Liang Li, and Ji Qi. 2018. **A self-attentive model with gate mechanism for spoken language understanding**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3824–3833. Association for Computational Linguistics.
- Bing Liu and Ian Lane. 2016. Attention-Based Recurrent Neural Network Models for Joint Intent Detection and Slot Filling. In *Interspeech 2016, 17th Annual Conference of the International Speech Communication Association, San Francisco, CA, USA, September 8-12, 2016*, pages 685–689.
- Xingkun Liu, Arash Eshghi, Pawel Swietojanski, and Verena Rieser. 2019. Benchmarking Natural Language Understanding Services for building Conversational Agents. In *Proceedings of the International Workshop on Spoken Dialogue System*, page to appear, Siracusa, Sicily, Italy.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. **Deep Contextualized Word Representations**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237. Association for Computational Linguistics.

- P. J. Price. 1990. [Evaluation of spoken language systems: The atis domain](#). In *Proceedings of the Workshop on Speech and Natural Language, HLT '90*, pages 91–95, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Abhinav Rastogi, Raghav Gupta, and Dilek Hakkani-Tur. 2018. [Multi-task Learning for Joint Language Understanding and Dialogue State Tracking](#). In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue*, pages 376–384. Association for Computational Linguistics.
- Erik F. Tjong Kim Sang and Jorn Veenstra. 1999. [Representing text chunks](#). In *Proceedings of the Ninth Conference on European Chapter of the Association for Computational Linguistics, EACL '99*, pages 173–179, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Victor Sanh, Thomas Wolf, and Sebastian Ruder. 2018. A hierarchical multi-task approach for learning embeddings from semantic tasks. *arXiv preprint arXiv:1811.06031*.
- M. Schuster and K.K. Paliwal. 1997. Bidirectional Recurrent Neural Networks. *IEEE Transactions on Signal Processing*, 45(11):2673–2681.
- Igor Shalyminov, Arash Eshghi, and Oliver Lemon. 2018. Multi-task learning for domain-general spoken disfluency detection in dialogue systems. In *Proceedings of SemDIAL 2018 (AixDial)*.
- Emma Strubell, Patrick Verga, Daniel Andor, David Weiss, and Andrew McCallum. 2018. [Linguistically-Informed Self-Attention for Semantic Role Labeling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5027–5038. Association for Computational Linguistics.
- Sainbayar Sukhbaatar, Arthur Szlam, Jason Weston, and Rob Fergus. 2015. End-to-end memory networks. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 2440–2448. Curran Associates, Inc.
- Xiaodong Zhang and Houfeng Wang. 2016. A Joint Model of Intent Determination and Slot Filling for Spoken Language Understanding. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI'16*. AAAI Press.
- Guineng Zheng, Subhabrata Mukherjee, Xin Luna Dong, and Feifei Li. 2018. [Opentag: Open attribute value extraction from product profiles](#). In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '18*, pages 1049–1058, New York, NY, USA. ACM.

Dialog State Tracking: A Neural Reading Comprehension Approach

Shuyang Gao*, Abhishek Sethi*, Sanchit Agarwal*
Tagyoung Chung, Dilek Hakkani-Tur

Amazon Alexa AI

{shuyag, abhsethi, agsanchi, tagyoung, hakkanit}@amazon.com

Abstract

Dialog state tracking is used to estimate the current belief state of a dialog given all the preceding conversation. Machine reading comprehension, on the other hand, focuses on building systems that read passages of text and answer questions that require some understanding of passages. We formulate dialog state tracking as a reading comprehension task to answer the question *what is the state of the current dialog?* after reading conversational context. In contrast to traditional state tracking methods where the dialog state is often predicted as a distribution over a closed set of all the possible slot values within an ontology, our method uses a simple attention-based neural network to point to the slot values within the conversation. Experiments on MultiWOZ-2.0 cross-domain dialog dataset show that our simple system can obtain similar accuracies compared to the previous more complex methods. By exploiting recent advances in contextual word embeddings, adding a model that explicitly tracks whether a slot value should be carried over to the next turn, and combining our method with a traditional joint state tracking method that relies on closed set vocabulary, we can obtain a joint-goal accuracy of 47.33% on the standard test split, exceeding current state-of-the-art by 11.75%**.

1 Introduction

A task-oriented spoken dialog system involves continuous interaction with a machine agent and a human who wants to accomplish a predefined task through speech. Broadly speaking, the system has

four components, the Automatic Speech Recognition (ASR) module, the Natural Language Understanding (NLU) module, the Natural Language Generation (NLG) module, and the Dialog Manager. The dialog manager has two primary missions: dialog state tracking (DST) and decision making. At each dialog turn, the state tracker updates the belief state based on the information received from the ASR and the NLU modules. Subsequently, the dialog manager chooses the action based on the dialog state, the dialog policy and the backend results produced from previously executed actions.

Table 1 shows an example conversation with the associated dialog state. Typical dialog state tracking system combines user speech, NLU output, and context from previous turns to track what has happened in a dialog. More specifically, the dialog state at each turn is defined as a distribution over a set of predefined variables (Williams et al., 2005). The distributions output by a dialog state tracker are sometimes referred to as the tracker’s belief or the belief state. Typically, the tracker has complete access to the history of the dialog up to the current turn.

Traditional machine learning approaches to dialog state tracking have two forms, generative and discriminative. In generative approaches, a dialog is modeled as a dynamic Bayesian network where true dialog state and true user action are unobserved random variables (Williams and Young, 2007); whereas the discriminative approaches are directly modeling the distribution over the dialog state given arbitrary input features.

Despite the popularity of these approaches, they often suffer from a common yet overlooked problem — relying on fixed ontologies. These systems, therefore, have trouble handling previously unseen mentions. On the other hand, reading comprehension tasks (Rajpurkar et al., 2016; Chen et al.,

*Authors contributed equally.

**We note that after publication, a new state-of-the-art can now be obtained with a similar attention mechanism followed by an encoder-decoder architecture (Wu et al., 2019).

2017; Reddy et al., 2019) require us to find the answer spans within the given passage and hence state-of-the-art models are developed in such a way that a fixed vocabulary for an answer is usually not required. Motivated by the limitations of previous dialog state tracking methods and the recent advances in reading comprehension (Chen, 2018), we propose a reading comprehension based approach to dialog state tracking. In our approach, we view the dialog as a *passage* and ask the question *what is the state of the current dialog?* We use a simple attention-based neural network model to find answer spans by directly pointing to the tokens within the dialog, which is similar to Chen et al. (2017). In addition to this attentive reading model, we also introduce two simple models into our dialog state tracking pipeline, a *slot carryover* model to help the tracker make a binary decision whether the slot values from the previous turn should be used; a *slot type* model to predict whether the answer is {Yes, No, DontCare, Span}, which is similar to Zhu et al. (2018). To summarize our contributions:

- We formulate dialog state tracking as a reading comprehension task and propose a simple attention-based neural network to find the state answer as a span over tokens within the dialog. Our approach overcomes the limitations of fixed-vocabulary issue in previous approaches and can generalize to unseen state values.
- We present the task of dialog state tracking as making three sequential decisions: i) a binary *carryover* decision by a simple slot carryover model ii) a *slot type* decision by a slot type model iii) a *slot span* decision by an attentive reading comprehension model. We show effectiveness of this approach.
- We adopt recent progress in large pre-trained contextual word embeddings, i.e., BERT (Devlin et al., 2018) into dialog state tracking, and get considerable improvement.
- We show our proposed model outperforms more complex previously published methods on the recently released MultiWOZ-2.0 corpus (Budzianowski et al., 2018; Ramadan et al., 2018). Our approach achieves a joint-goal accuracy of 42.12%, resulting in a 6.5% absolute improvement over previous state-of-

```

User: I need to book a hotel in the east that has 4 stars.
Hotel area=east, stars=4
Agent: I can help you with that. What is your price range?
User: That doesn't matter if it has free wifi and parking.
Hotel parking=yes, internet=yes
price=dontcare, stars=4, area=east
Agent: If you'd like something cheap,
I recommend Allenbell
User: That sounds good, I would also like a
taxi to the hotel from cambridge
Hotel parking=yes, internet=yes
price=dontcare, area=east, stars=4
Taxi departure=Cambridge
destination=Allenbell

```

Table 1: An example conversation in MultiWOZ-2.0 with dialog states after each turn.

the-art. Furthermore, if we combine our results with the traditional joint state tracking method in Liu and Lane (2017), we achieve a joint-goal accuracy of 47.33%, further advancing the state-of-the-art by 11.75%.

- We provide an in-depth error analysis of our methods on the MultiWOZ-2.0 dataset and explain to what extent an attention-based reading comprehension model can be effective for dialog state tracking and inspire future improvements on this model.

2 Related Work

Dialog State Tracking Traditionally, dialog state tracking methods assume a *fixed ontology*, wherein the output space of a slot is constrained by the predefined set of possible values (Liu and Lane, 2017). However, these approaches are not applicable for unseen values and do not scale for large or potentially unbounded vocabulary (Nouri and Hosseini-Asl, 2018). To address these concerns, a class of methods employing scoring mechanisms to predict the slot value from a endogenously defined set of candidates have been proposed (Rastogi et al., 2017; Goel et al., 2018). In these methods, the candidates are derived from either a predefined ontology or by extraction of a word or n -grams in the prior dialog context. Previously, Perez and Liu (2017) also formulated state tracking as a machine reading comprehension problem. However, their model architecture used a memory network which is relatively complex and still assumes a fixed-set vocabulary. Perhaps, the most similar technique to our work is the pointer networks proposed by Xu and Hu (2018) wherein an attention-based mechanism is

employed to *point* the start and end token of a slot value. However, their formulation does not incorporate a *slot carryover* component and outlines an encoder-decoder architecture in which the slot type embeddings are derived from the last state of the RNN.

Reading Comprehension A reading comprehension task is commonly formulated as a supervised learning problem where for a given training dataset, the goal is to learn a predictor, which takes a passage p and a corresponding question q as inputs and gives the answer a as output. In these tasks, an answer type can be cloze-style as in CNN/Daily Mail (Hermann et al., 2015), multiple choice as in MCTest (Richardson et al., 2013), span prediction as in SQuAD (Rajpurkar et al., 2016), and free-form answer as in NarrativeQA (Kočiskỳ et al., 2018). In span prediction tasks, most models encode a question into an embedding and generate an embedding for each token in the passage and then a similarity function employing attention mechanism between the question and words in the passage to decide the starting and ending positions of the answer spans (Chen et al., 2017; Chen, 2018). This approach is fairly generic and can be extended to multiple choice questions by employing bilinear product for different types (Lai et al., 2017) or to free-form text by employing seq-to-seq models (Sutskever et al., 2014).

Deep Contextual Word Embeddings The recent advancements in the neural representation of words includes using character embeddings (Seo et al., 2016) and more recently using contextualized embeddings such as ELMO (Peters et al., 2018) and BERT (Devlin et al., 2018). These methods are usually trained on a very large corpus using a language model objective and show superior results across a variety of tasks. Given their wide applicability (Liu et al., 2019), we employ these architectures in our dialog state tracking task.

3 Our Approach

3.1 DST as Reading Comprehension

Let us denote a sub-dialog D_t of a dialog D as prefix of a full dialog ending with user’s t th utterance, then state of the dialog D_t is defined by the values of constituent slots $s_j(t)$, i.e., $S_t = \{s_1(t), s_2(t), \dots, s_M(t)\}$.

Using the terminology in reading comprehension tasks, we can treat D_t as a *passage*, and for each slot i , we formulate a question q_i : *what is the value for slot i ?* The dialog state tracking task then becomes understanding a sub-dialog D_t and to answer the question q_i for each slot i .

3.2 Encoding

Dialog Encoding For a given dialog D_t at turn t , we first concatenate user utterances and agent utterances $\{\mathbf{u}_1, \mathbf{a}_1, \mathbf{u}_2, \mathbf{a}_2, \dots, \mathbf{u}_t\}$. To differentiate between user utterance and agent utterance, we add symbol [U] before each user utterance and [A] before each agent utterance. Then, we use pre-trained word vectors to form \mathbf{p}_i for each token in the dialog sequence and pass them as input into a recurrent neural network, i.e.,

$$\{\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_L\} = RNN(\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_L) \quad (1)$$

where L is the total length of the concatenated dialog sequence and \mathbf{d}_i is the output of RNN for each token, which is expected to encode context-aware information of the token. In particular, for pre-trained word vectors \mathbf{p}_i , we experiment with using deep contextualized word embeddings using BERT (Devlin et al., 2018). For RNN, we use a one layer bidirectional long short-term memory network (LSTM) and each \mathbf{d}_i is the concatenation of two LSTMs from both directions, i.e., $\mathbf{d}_i = (\overleftarrow{\mathbf{d}}_i; \overrightarrow{\mathbf{d}}_i)$. Furthermore, we denote $\mathbf{e}(t)$ as our dialog embedding at turn t as follows:

$$\mathbf{e}(t) = (\overleftarrow{\mathbf{d}}_1; \overrightarrow{\mathbf{d}}_L) \quad (2)$$

Question Encoding In our methodology, we formulate questions q_i defined earlier as *what is the value for slot i ?* For each dialog, there are M similar questions corresponding to M slots, therefore, we represent each question q_i as a fixed-dimension vector \mathbf{q}_i to learn.

3.3 Models

Overview In our full model set up, three different model components are used to make a sequence of predictions: first, we use a *slot carryover* model for deciding whether to carryover a slot value from the last turn. If the first model decided not to carry over, a *slot type* model is executed to predict type of the answer from a set of {Yes, No, DontCare, Span}. If the *slot type* model predicts span, *slot span* model will finally be predicting the slot value as a span of tokens

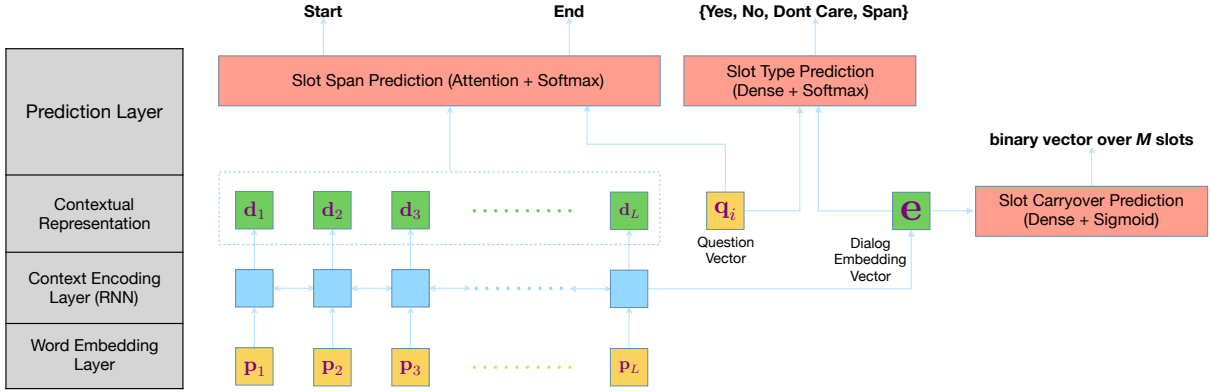


Figure 1: Our attentive reading comprehension system for dialog state tracking. There are three prediction components on top (from right to left): 1) slot carryover model to predict whether a particular slot needs to be updated from previous turn 2) slot type model to predict the type of slot values from $\{\text{Yes, No, DontCare, Span}\}$ 3) slot span model to predict the start and end span of the value within the dialog.

within the dialog. The full model architecture is shown in Figure 1.

Slot Carryover Model To model dynamic nature of dialog state, we introduce a model whose purpose is to decide whether to carry over a slot value from the previous turn. For a given slot s_j , $C_j(t) = 1$ if $s_j(t) \neq s_j(t-1)$ and 0 if they are equal. We multiply the dialog embedding $\mathbf{e}(t)$ with a fully connected layer \mathbf{W}_i to predict the change for slot i as:

$$P(C_i(t)) = \text{sigmoid}(\mathbf{e}(t) \cdot \mathbf{W}_i) \quad (3)$$

The network architecture is shown in Figure 1. In our implementation, the weights \mathbf{W}_i for each slot are trained together, i.e., the neural network would predict the slot carryover change $C_i(t)$ jointly for all M slots.

Slot Type Model A typical dialog state comprises of slots that can have both categorical and named entities within the context of conversation. To adopt a flexible approach and inspired by the state-of-the-art reading comprehension approaches, we propose a classifier that predicts the type of slot value at each turn. In our setting, we prescribe the output space to be $\{\text{Yes, No, DontCare, Span}\}$ where *Span* indicates the slot value is a named entity which can be found within the dialog. As shown in Figure 1, we concatenate the dialog embedding $\mathbf{e}(t)$ with the question encoding \mathbf{q}_i for slot i as the input to the affine layer \mathbf{A} to predict the slot type $T_i(t)$ as:

$$P(T_i(t)) \propto \exp(\mathbf{A} \cdot (\mathbf{e}(t); \mathbf{q}_i)) \quad (4)$$

Slot Span Model We map our slot values into a span with start and end position in our flattened conversation D_t . We then use the dialog encoding vectors $\{\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_L\}$ and the question vector \mathbf{q}_i to compute the bilinear product and train two classifiers to predict the start position and end position of the slot value. More specifically, for slot j ,

$$P_j^{(start)}(x) = \frac{\exp(\mathbf{d}_x \Theta^{(start)} \mathbf{q}_j)}{\sum_{x'} \exp(\mathbf{d}_{x'} \Theta^{(start)} \mathbf{q}_j)} \quad (5)$$

Similarly, we define $P_j^{(end)}(x)$ with $\Theta^{(end)}$. During span inference, we choose the best span from word i to word i' such that $i \leq i'$ and $P_j^{(start)}(i) \times P_j^{(end)}(i')$ is maximized, in line with the approach by Chen et al. (2017).

4 Experiments

4.1 Data

We use the recently-released MultiWOZ-2.0 dataset (Budzianowski et al., 2018; Ramadan et al., 2018) to test our approach. This dataset consists of multi-domain conversations from seven domains with a total of 37 slots across domains. Many of these slot types such as *day* and *people* are shared across multiple domains. In our experiments, we process each slot independently by considering the concatenation of slot domain, slot category, and slot name, e.g., $\{\text{bus.book.people}\}$, $\{\text{restaurant.semi.food}\}$. An example of conversation is shown in Table 1. We use standard training/development/test present in the data set.

It is worth-noting that the dataset in the current form has certain annotation errors. First, there is

Method	Accuracy
MultiWOZ Benchmark	25.83%
GLAD (Zhong et al., 2018)	35.57%
GCE (Nouri and Hosseini-Asl, 2018)	35.58%
Our approach (single)	39.41%
Our approach (ensemble)	42.12%
HyST (ensemble) (Goel et al., 2019)	44.22%
Our approach + JST (ensemble)	47.33%

Table 2: Joint goal accuracy on MultiWOZ-2.0. We present both single and ensemble results for our approach.

lack of consistency between the slot values in the ontology and the ground truth in the context of the dialog. For example, the ontology has *moderate* but the dialog context has *moderately*. Second, there are erroneous delay in the state updates, sometimes extending turns in the dialog. This error negatively impacts the performance of the slot carryover model.

4.2 Experimental Setup

We train our three models independently without sharing the dialog context. For all the three models, we encode the word tokens with BERT (Devlin et al., 2018) followed by an affine layer with 200 hidden units. This output is then fed into a one-layer bi-directional LSTM with 50 hidden units to obtain the contextual representation as show in Figure 1. In all our experiments, we keep the parameters of the BERT embeddings frozen.

For slot carryover model, we predict a binary vector over 37 slots jointly to get the decisions of whether to carry over values for each slot. For slot type and slot span models, we treat dialog-question pairs (D_t, q_i) as separate prediction tasks for each slot.

We use the learning rate of 0.001 with ADAM optimizer and batch size equal to 32 for all three models. We stop training our models when the loss on the development set has not been decreasing for ten epochs.

5 Results

Table 2 presents our results on MultiWOZ-2.0 test dataset. We compare our methods with global-local self-attention model (GLAD) (Zhong et al., 2018), global-conditioned encoder model (GCE) (Nouri and Hosseini-Asl, 2018), and hybrid joint state tracking model (OV ST+JST) (Liu

and Lane, 2017; Goel et al., 2019). As in previous work, we report joint goal accuracy as our metric. For each user turn, joint goal accuracy checks whether all predicted states exactly matches the ground truth state for all slots. We can see that our system with single model can achieve 39.41% joint goal accuracy, and with the ensemble model we can achieve 42.12% joint goal accuracy.

Table 3 shows the accuracy for each slot type for both our method and the joint state tracking approach with fix vocabulary in Goel et al. (2019). We can see our approach tends to have higher accuracy on some of the slots that have larger set of possible values such as `attraction.semi.name` and `taxi.semi.destination`. However, it is worth-noting that even for slots with smaller vocabulary sizes such as `hotel.book.day` and `hotel.semi.pricerange`, our approach achieves better accuracy than using closed vocabulary approach. Our hypothesis for difference is that such information appear more frequently in user utterance thus our model is able to learn it more easily from the dialog context.

We also reported the result for a hybrid model by combining our approach with the JST approach in (Goel et al., 2019). Our combination strategy is as follows: first we calculated the slot type accuracy for each model on the development dataset; then for each slot type, we choose to use the predictions from either our model or JST model based on the accuracy calculated on the development set, whichever is higher. With this approach, we achieve the joint-goal accuracy of 46.28%. We hypothesize that this is because our method uses an open vocabulary, where all the possible values can only be obtained from the conversation; the joint state tracking method uses closed ontology, we can get the best of both the worlds by combining two methods.

5.1 Ablation Analysis

Table 4 illustrates the ablation studies for our model on development set. The contextual embedding BERT (Devlin et al., 2018) can give us around 2% gains. As for the oracle models, we can see that even if using all the oracle results (ground truth), our development set accuracy is only 73.12%. This is because our approach is only considering the values within the conversation, if values are not present in the dialog, the

Slot Name	Ours	JST	Vocab Size
attraction.semi.area	0.9637	0.9719	16
attraction.semi.name	0.9213	0.9013	137
attraction.semi.type	0.9205	0.9637	37
bus.book.people	1.0000	1.0000	1
bus.semi.arriveBy	1.0000	1.0000	1
bus.semi.day	1.0000	1.0000	2
bus.semi.departure	1.0000	1.0000	2
bus.semi.destination	1.0000	1.0000	5
bus.semi.leaveAt	1.0000	1.0000	2
hospital.semi.department	0.9991	0.9988	52
hotel.book.day	0.9863	0.9784	11
hotel.book.people	0.9714	0.9847	9
hotel.book.stay	0.9736	0.9809	9
hotel.semi.area	0.9679	0.9570	24
hotel.semi.internet	0.9713	0.9718	8
hotel.semi.name	0.9147	0.9056	89
hotel.semi.parking	0.9563	0.9657	8
hotel.semi.pricerange	0.9679	0.9666	9
hotel.semi.stars	0.9627	0.9759	13
hotel.semi.type	0.9140	0.9261	18
restaurant.book.day	0.9874	0.9871	10
restaurant.book.people	0.9787	0.9881	9
restaurant.book.time	0.9882	0.9578	61
restaurant.semi.area	0.9607	0.9654	19
restaurant.semi.food	0.9741	0.9691	104
restaurant.semi.name	0.9113	0.8781	183
restaurant.semi.pricerange	0.9662	0.9626	11
taxi.semi.arriveBy	0.9893	0.9719	101
taxi.semi.departure	0.9665	0.9304	261
taxi.semi.destination	0.9634	0.9288	277
taxi.semi.leaveAt	0.9821	0.9524	119
train.book.people	0.9586	0.9718	13
train.semi.arriveBy	0.9738	0.9491	107
train.semi.day	0.9854	0.9783	11
train.semi.departure	0.9599	0.9710	35
train.semi.destination	0.9538	0.9699	29
train.semi.leaveAt	0.9595	0.9478	134

Table 3: Slot accuracy breakdown for our approach versus joint state tracking method. Bolded slots are the ones have better performance using our attentive reading comprehension approach.

oracle models would fail. It is interesting to see that if we replace our slot carryover model with an oracle one, the accuracy improves significantly to 60.18% (+19.08%) compared to replacing other two models (41.43% and 45.77%). This is because our span-based reading comprehension approach model already gives us accuracy as high as 96% per slot on development data, there is not much room for improvement. Whereas our binary slot carryover model only achieve an accuracy of 72% per turn. We hypothesis that for slot carryover problem is imbalanced, i.e., there are significantly more slot carryovers than slot updates, making the

Ablation	Dev Accuracy
Oracle Models	73.12%
Our approach	41.10%
- BERT	39.19%
+ Oracle Slot Type Model	41.43%
+ Oracle Slot Span Model	45.77%
+ Oracle Slot Carryover Model	60.18%

Table 4: Ablation study on our model components for MultiWOZ-2.0 on development set for joint goal accuracy.

model training and predictions harder. This suggest further improvements are needed for *slot carryover* model to make overall state tracking accuracy higher.

5.2 Error Analysis

In Table 5, we conduct an error analysis of our models and investigate its performance for different use cases. Since we formulate the problem to be an open-vocabulary state tracking approach wherein the slot values are extracted in the dialog context, we divide the errors into following categories:

- **Unanswerable Slot Error** This category contains two type of errors: (1) Ground truth slot is a not *None* value, but our prediction is *None*; (2) Ground truth slot is *None*, but our prediction is a not *None* value. This type of error can be attributed to the incorrect predictions made by our slot carryover model.
- **Imprecise Slot Reference** where multiple potential candidates in the context exists. The model refers to the incorrect *entity* in the conversation. This error can be largely attributed to following reasons: (1) the model overfits to the set of tokens that it has seen more frequently in the training set; (2) the model does not generalize well for scenarios where the user corrects the previous entity; (3) the model incorrectly overfits to the order or position of the entity in the context. These reasons motivate future research in incorporating more neural reading comprehension approaches for dialog state tracking.
- **Imprecise Slot Resolution** In this type of errors, we cannot find the exact match of ground truth value in the dialog context.

Category	Hypothesis	Reference	Context	(%)
Unanswerable	not <i>None</i>	<i>None</i>	...	42.4
Slot Error	<i>None</i>	not <i>None</i>	...	23.1
Incorrect slot Reference	4	8	...3 nights, and 4 people . Thank You! [A] Booking was unsuccessful ... I'd like to book there Monday for 1 night with 8 people	19.1
Incorrect Slot Resolution	3:30	15:30	...you like to arrive at the Cinema? [U] I want to leave the hotel by 3:30 [A] Your taxi reservation departing ...	12.9
Imprecise Slot Boundary	nandos city centre	nandos	...number is 01223902168 [U] Great I am also looking for a restaurant called nandos city centre ...	2.5

Table 5: Error categorization and percentage distribution: representative example from each category and an estimate breakdown of the error types on development set, based on the analysis of 200 error samples produced by our model. Numbers of the first category is exact because we are able to summarize this error category statistically.

However, our predicted model span is a paraphrase or has very close meaning to the ground truth. This error is inherent in approaches that do not extract the slot value from an ontology but rather the dialog context. On similar lines, we also observe cases where the slot value in the dialog context is *resolved* (or canonicalized) to a different surface-form entity that is perhaps more amenable for downstream applications.

- **Imprecise Slot Boundary** In this category of errors, our model chooses a span that is either a superset or subset of the correct reference. This error is especially frequent for proper nouns where the model has a weaker signal to outline the slot boundary precisely.

Table 5 provides us the error examples and estimated percentage from each category. "Unanswerable Slot" accounts for 65.5% errors for our model, this indicates further attention may be needed to the slot carryover model, otherwise it would become a barrier even if we have a perfect span model. This finding is in alignment with our ablation studies in Table 4, where oracle slot carryover model would give us the most boost in joint goal accuracy. Additionally, 12.9% of errors are due to imprecise slot resolution, this suggests future directions of resolving the context words to the ontology.

5.3 Evaluating Different Context Encoders for Slot Carryover Model

As shown in oracle ablation studies in Table 4, slot carryover model plays a significant role in our

pipeline. Therefore we explore the different types of context encoders for slot carryover model to see whether if it improves the performance in table 6. In addition to use a flat dialog context of user and agent turns [U] and [A] to predict carryover for every slot in the state, we explored hierarchical context encoder with an utterance-level LSTM over each user and agent utterance and a dialog-level LSTM over the whole dialog with both constrained and unconstrained context window, similar to Liu and Lane (2017). However, we did not witness any significant performance change across the two variants as show in Table 6. Lastly, we employed self-attention over the flattened dialog context in line with Vaswani et al. (2017). However, we can see from Table 6 that this strategy slightly hurts the model performance. One hypothesis for sub par slot carryover model performance is due to the inherent noise in the annotated data for state updates. Through a preliminary analysis on the development set, we encountered few erroneous delay in the state updates sometimes extending to over multiple turns. Nevertheless, these experimental results motivate future research in *slot carryover* models for multi-domain conversations.

5.4 Analyzing Conversation Depth

In Table 7, we explore the relationship between the depth of a conversation and the performance of our models. More precisely, we segment a given set of dialogs into individual *turns* and measure the state accuracy for each of these segments. We mark a turn correct only if all the slots in its state are predicted correctly. We observe that the model perfor-

Context Feature	Per Turn Carryover Accuracy
Flat Context (LSTM)	75.10%
Hierarchical Context (all turns)	75.98%
Hierarchical Context (≤ 3 turns)	75.60%
Flat Context (Self-Attention)	74.75%

Table 6: Analyzing the different types of context features for Slot Carryover Model

Conversation Depth t	Total Turns	% Incorrect Turns
1	1000	23.90
2	1000	38.30
3	997	50.85
4	959	61.52
5	892	71.52
6	811	76.82
7	656	82.77
8	475	87.37
9	280	89.64
10	153	94.77

Table 7: Analyzing the overall model robustness for conversation depth for MultiWOZ-2.0

mance degrades as the number of turns increase. The primary reason for this behavior is that an error committed earlier in the conversation can be carried over for later turns. This results in a strictly higher probability for a later turn to be incorrect as compared to the turns earlier in the conversation. These results motivate future research in formulating models for state tracking that are more robust to the depth of the conversation.

6 Conclusion

The problem of tracking user’s belief state in a dialog is a historically significant endeavor. In that context, research on dialog state tracking has been geared towards discriminative methods, where these methods are usually estimating the distribution of user state over a fixed vocabulary. However, modern dialog systems presents us with problems requiring a large scale perspective. It is not unusual to have thousands of slot values in the vocabulary which could have millions variations of dialogs. So we need a vocabulary-free way to pick out the slot values.

How can we pick the slot values given an in-

finite amount of vocabulary size? Some methods adopt a candidate generation mechanism to generate slot values and make a binary decision with the dialog context. Attention-based neural network gives a clear and general basis for selecting the slot values by direct pointing to the context spans. While this type of methods has already been proposed recently, we explored this type of idea furthermore on MultiWOZ-2.0 dataset.

We introduced a simple attention based neural network to encode the dialog context and point to the slot values within the conversation. We have also introduced an additional slot carryover model and showed its impact on the model performance. By incorporating the deep contextual word embeddings and combining the traditional fixed vocabulary approach, we significantly improved the joint goal accuracy on MultiWOZ-2.0.

We also did a comprehensive analysis to see to what extent our proposed model can achieve. One interesting and significant finding from the oblation studies suggests the importance of the slot carryover model. We hope this finding can inspire future dialog state tracking research to work towards this direction, i.e., predicting whether a slot of state is none or not.

The field of machine reading comprehension has made significant progress in recent years. We believe human conversation can be viewed as a special type of context and we hope that the developments suggested here can help dialog related tasks benefit from modern reading comprehension models.

Acknowledgements

The authors would like to thank Rahul Goel, Anuj Kumar Goyal, Angeliki Metallinou and other Alexa AI team members for their useful discussions and feedbacks.

References

- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Ultes Stefan, Ramadan Osman, and Milica Gašić. 2018. Multiwoz - a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Danqi Chen. 2018. *Neural Reading Comprehension and Beyond*. Ph.D. thesis, Stanford University.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading Wikipedia to answer open-domain questions. In *Association for Computational Linguistics (ACL)*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Rahul Goel, Shachi Paul, Tagyoung Chung, Jeremie Lecomte, Arindam Mandal, and Dilek Hakkani-Tur. 2018. Flexible and scalable state tracking framework for goal-oriented dialogue systems. *arXiv preprint arXiv:1811.12891*.
- Rahul Goel, Shachi Paul, and Dilek Hakkani-Tür. 2019. Hyst: A hybrid approach for flexible and accurate dialogue state tracking.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Advances in neural information processing systems*, pages 1693–1701.
- Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. 2018. The narrativeqa reading comprehension challenge. *Transactions of the Association of Computational Linguistics*, 6:317–328.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. Race: Large-scale reading comprehension dataset from examinations. *arXiv preprint arXiv:1704.04683*.
- Bing Liu and Ian Lane. 2017. An end-to-end trainable neural network model with belief tracking for task-oriented dialog. *Proc. Interspeech 2017*, pages 2506–2510.
- Nelson F Liu, Matt Gardner, Yonatan Belinkov, Matthew Peters, and Noah A Smith. 2019. Linguistic knowledge and transferability of contextual representations. *arXiv preprint arXiv:1903.08855*.
- Elnaz Nouri and Ehsan Hosseini-Asl. 2018. Toward scalable neural dialogue state tracking model. In *32nd Conference on Neural Information Processing Systems (NeurIPS 2018), 2nd Conversational AI workshop*.
- Julien Perez and Fei Liu. 2017. Dialog state tracking, a machine reading approach using memory network. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 305–314.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392.
- Osman Ramadan, Paweł Budzianowski, and Milica Gasic. 2018. Large-scale multi-domain belief tracking with knowledge sharing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, volume 2, pages 432–437.
- Abhinav Rastogi, Dilek Hakkani-Tür, and Larry Heck. 2017. Scalable multi-domain dialogue state tracking. In *Automatic Speech Recognition and Understanding Workshop (ASRU), 2017 IEEE*, pages 561–568. IEEE.
- Siva Reddy, Danqi Chen, and Christopher D Manning. 2019. CoQA: A conversational question answering challenge. *Transactions of the Association of Computational Linguistics (TACL)*.
- Matthew Richardson, Christopher JC Burges, and Erin Renshaw. 2013. Mctest: A challenge dataset for the open-domain machine comprehension of text. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 193–203.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016. Bidirectional attention flow for machine comprehension. *arXiv preprint arXiv:1611.01603*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Jason D Williams, Pascal Poupart, and Steve Young. 2005. Factored partially observable markov decision processes for dialogue management. In *Proc. IJCAI Workshop on Knowledge and Reasoning in Practical Dialogue Systems*, pages 76–82.
- Jason D Williams and Steve Young. 2007. Partially observable markov decision processes for spoken dialog systems. *Computer Speech & Language*, 21(2):393–422.

- Chien-Sheng Wu, Andrea Madotto, Ehsan Hosseini-Asl, Caiming Xiong, Richard Socher, and Pascale Fung. 2019. Transferable multi-domain state generator for task-oriented dialogue systems. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 808–819.
- Puyang Xu and Qi Hu. 2018. An end-to-end approach for handling unknown slot values in dialogue state tracking. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2018. Global-locally self-attentive dialogue state tracker. *arXiv preprint arXiv:1805.09655*.
- Chenguang Zhu, Michael Zeng, and Xuedong Huang. 2018. Sdnet: Contextualized attention-based deep network for conversational question answering. *arXiv preprint arXiv:1812.03593*.

Cross-Corpus Data Augmentation for Acoustic Addressee Detection

Oleg Akhtiamov^{1,2}, Ingo Siegert³, Alexey Karpov⁴, Wolfgang Minker¹

¹Ulm University, Ulm, Germany

²ITMO University, St. Petersburg, Russia

³Otto-von-Guericke-University, Magdeburg, Germany

⁴SPIIRAS, St. Petersburg, Russia

oakhtiamov@gmail.com, ingo.siegert@ovgu.de,
karpov@iias.spb.su, wolfgang.minker@uni-ulm.de

Abstract

Acoustic addressee detection (AD) is a modern paralinguistic and dialogue challenge that especially arises in voice assistants. In the present study, we distinguish addressees in two settings (a conversation between several people and a spoken dialogue system, and a conversation between several adults and a child) and introduce the first competitive baseline (unweighted average recall equals 0.891) for the Voice Assistant Conversation Corpus that models the first setting. We jointly solve both classification problems, using three models: a linear support vector machine dealing with acoustic functionals and two neural networks utilising raw waveforms alongside with acoustic low-level descriptors. We investigate how different corpora influence each other, applying the mixup approach to data augmentation. We also study the influence of various acoustic context lengths on AD. Two-second speech fragments turn out to be sufficient for reliable AD. Mixup is shown to be beneficial for merging acoustic data (extracted features but not raw waveforms) from different domains that allows us to reach a higher classification performance on human-machine AD and also for training a multipurpose neural network that is capable of solving both human-machine and adult-child AD problems.

1 Introduction

For the past years, the phenomenon of multiparty spoken interaction has drawn many researchers' attention (Busso et al., 2007; Gilmartin et al., 2018; Haider et al., 2018). How do we address other people in such conversations? Normally, we do this either explicitly, directly specifying desirable addressees by their names, or implicitly, using contextual (Ouchi and Tsuboi, 2016; Zhang et al., 2018) and multimodal markers (Tsai et al., 2015; Akhtiamov et al., 2017b; Akhtiamov and

Palkov, 2018; Le Minh et al., 2018). Particularly, we use acoustic markers to emphasise special addressees, such as hard-of-hearing people (Batliner et al., 2008), elderly people, children (Schuller et al., 2017), and automatic spoken dialogue systems (SDSs) (Batliner et al., 2008; Shriberg et al., 2013; Akhtiamov et al., 2017a; Pugachev et al., 2017). We act in this way if we realise that our addressee may have some communicational difficulties, and therefore we modify our normal manner of speech, making it more rhythmical, louder, and generally more understandable as soon as we start talking to such conversational partners (Shriberg et al., 2012; Siegert and Krüger, 2018).

In the present research, we deal with two binary acoustic addressee detection (AD) problems. The first problem of human-machine addressee detection (H-M AD) arises in conversations within a group of users solving a cooperative task by means of an SDS. The users may talk to each other and also contact the system from time to time. The system is supposed to distinguish between machine- and human-directed utterances in order to maintain conversations in a realistic manner. Human-directed utterances do not require a direct system response and should be processed with the system in an implicit way. We use the following two corpora to model the H-M AD problem: the Smart Video Corpus (SVC) (Batliner et al., 2008) and the Voice Assistant Conversation Corpus (VACC) (Siegert et al., 2018). The first competitive VACC baseline is introduced in the present paper. The second problem of adult-child addressee detection (A-C AD) appears in conversations between a group of adults and a child. In this case, our system is supposed to distinguish between child- and adult-directed utterances. A possible application for such a system of adult-child conversation monitoring is the estimation of children's and adults' conversational behaviour that will allow us

to measure Interaction Quality (IQ) (Spirina et al., 2016). According to this complex metric, we will be able to assess the children’s progress in maintaining conversations. We model the A-C AD problem, using the HomeBank Child-Adult Addressee Corpus (HB-CHAAC, mentioned as HB below for simplicity) (Casillas et al., 2017).

We consider both binary classification problems as one: the utterances belonging to the first category are directed to a special addressee that may be an SDS or a child having a lack of communicational skills. The utterances belonging to the second category are directed to ordinary adults without any impairments that may cause miscommunication. In this light, we conduct a series of cross-corpus experiments and merge several corpora with the *mixup* method. This data augmentation technique has already been studied on image classification (Zhang et al., 2017), speech recognition (Medennikov et al., 2018), and acoustic emotion recognition (Fedotov et al., 2018b).

The present paper has the following contributions: the H-M and the A-C AD problem are jointly analysed by means of machine learning; mixup in combination with state-of-the-art classifiers is applied to cross-corpus acoustic AD for the first time; mixup capabilities are investigated over various speech signal representations (including raw data), acoustic context lengths, corpora, domains, languages, and classification problems.

2 Related Work

Several studies have already been conducted on the problem of acoustic H-M AD. The current acoustic SVC baseline was introduced by Akhtiamov et al. (2017a), who applied a feature selection method to a large paralinguistic feature set containing various functionals computed over low-level descriptor (LLD) contours (2013 ComParE feature set described by Eyben (2015)). The ComParE LLDs and their functionals were shown to be a universal solution for a wide range of paralinguistic problems besides AD, e.g. acoustic emotion recognition (Fedotov et al., 2018a), native speech detection, and neurological pathology estimation (Schuller et al., 2015). The same attribute set in combination with a linear support vector machine (SVM) alongside with other models including an end-to-end neural network was applied to the problem of acoustic A-C AD on HB by Schuller et al. (2017). HB was in-

troduced within the Addressee Sub-Challenge of the Interspeech 2017 Computational Paralinguistics Challenge (ComParE) (Schuller et al., 2017) that has already been finished. However, the challenge organisers proposed an extremely competitive baseline (Schuller et al., 2017) that none of the challenge participants managed to surpass, and therefore the HB classification problem remains of great scientific and practical interest.

There also exist speech signal representations designed specially for acoustic H-M AD. Shriberg et al. (2013) suggested modelling speech rhythm and vocal effort with high-abstract attributes: energy contour features, voice quality and spectral tilt features, and delta energy at voicing onsets/offsets. The energy contour and tilt features employed Gaussian mixture models (GMMs) to compute a log likelihood ratio of the two addressee classes. The machine-directed utterances from the corpus used for experiments in the latter study were short predefined commands consisting of three words on average. However, the machine- and child-directed utterances from the data that we have at our disposal were recorded under real-life conditions and usually contain whole sentences of spontaneous speech. Furthermore, it is unclear how these specific attributes perform on A-C AD. Therefore, we would not like to confine to such a narrow attribute set. Instead, we want to use the ComParE features in order to capture all the variety of spontaneous speech. An argument in favour of low-level features, such as LLDs and raw data, is the possibility to use them in combination with deep neural networks capable of performing feature selection and feature transformation implicitly for a specific problem. In the present study, we apply the ComParE functionals jointly with simple linear models, while lower-level features (raw audio and the ComParE LLDs) are used in combination with deep neural networks that learn high-level feature representations for our AD problem. Compared to Mallidi et al. (2018), we do not have that much data for training our networks on acoustic AD. We offset this lack by means of data augmentation.

3 Proposed Approach

3.1 Classifiers

We apply the following three models to audio classification. The first classifier (**func**) is a simple SVM with a linear kernel (Hofmann and Klinken-

berg, 2013). This model deals with the ComParE feature set comprising 6373 functionals (Eyben, 2015) extracted at the utterance level.

The second classifier (**LLD**) consists of two stacked long short-term memory (LSTM) (Hochreiter and Schmidhuber, 1997) layers followed by a global max pooling, a dropout (Srivastava et al., 2014), and a softmax layer. As input, the first layer receives the same LLD sequences used for computing the ComParE functionals. Each sequence element is a vector of 130 LLDs extracted for a sliding time window of 60 ms with an overlap of 50 ms. The sequences are extracted from acoustic context windows of various lengths (from 1/8 to 8 s). The context windows are cut out of audio files with an overlap of 75%. The predictions obtained on several windows belonging to one utterance are averaged to get the final utterance-level prediction.

The third classifier (**e2e**) performing end-to-end speech signal processing differs from the second model in the following way: the sequences of the ComParE LLDs are replaced by the output of a convolutional neural network (CNN). As a result, we obtain a convolutional recurrent neural network (CRNN) that is quite similar to the one suggested by Trigeorgis et al. (2016) for acoustic emotion recognition. However, the initial network architecture specified in the latter study did not provide any reliable results on our AD problem probably due to a lack of perceptive abilities. For this reason, we replaced the initial two-layer CNN by a deeper one. We took the five-layer SoundNet architecture (Aytar et al., 2016) as the reference point for our CNN, cut off its last convolutional layer and scaled the filter sizes and the number of units in each layer in accordance with the input signal resolution and the available amount of our training data. The final shape of the e2e model is depicted in Figure 1.

For the func and LLD models, we use statistical corpus normalisation by bringing the handcrafted features to zero mean and unit variance. For the e2e model, we employ batch normalisation (Ioffe and Szegedy, 2015) between each convolution and activation instead. Training our neural networks, we use the following parameters optimised on a development set: Gaussian noise applied to the input signal if mixup is disabled, 20% dropout applied directly before the softmax layer, rectified linear unit (ReLU) as an activation function for all




Figure 1: E2e classifier. To obtain the LLD model, we replace the middle part of the e2e model by the ComParE LLD sequences. Notation of the layers in the middle part of the e2e model: *layer_name(n_units, filter_size, stride)*, other layers: *layer_name(n_units)*.

convolutional layers, categorical cross-entropy as a loss function, Adam (Kingma and Ba, 2014) as a weight optimisation algorithm, 100 epochs, and a batch size of 32 examples. The initial learning rate is chosen from the set $\{10^{-3}, 10^{-4}, 10^{-5}\}$ and then divided by 10 if there is no performance improvement observed for the past 10 epochs on the development set. We make checkpoints, saving the current model weights at each epoch and using the best checkpoint as the resulting model according to its performance on the development set.

Both neural networks were designed with TensorFlow (Abadi et al., 2016). The func model was implemented with RapidMiner (Hofmann and Klinkenberg, 2013). We used the openSMILE toolkit (Eyben et al., 2013) and its 2013 ComParE feature configuration (Eyben, 2015) to extract acoustic LLDs and their functionals.

3.2 Data Augmentation

We apply a simple yet efficient approach to data augmentation called *mixup* (Zhang et al., 2017). The core idea of this method is to regularise our model by encouraging it to make linear predictions in the vector space between seen data points. The method generates artificial examples as linear combinations of the feature and label vectors taken from two arbitrary real examples and mixed at a proportion λ in the following way:

$$x_{art} = \lambda x_i + (1 - \lambda)x_j, \quad (1)$$

$$y_{art} = \lambda y_i + (1 - \lambda)y_j. \quad (2)$$

λ is randomly generated from a β -distribution for each artificial example. This distribution is defined as follows by a coefficient α that lies within the interval $(0, \infty)$ and determines the probability that our generated example lies close to one of real examples:

$$f(x; \alpha) = x^{\alpha-1}(1-x)^{\alpha-1}. \quad (3)$$

VACC (German)				SVC (German)				HB (English)			
Label	Train	Dev	Test	Label	Train	Dev	Test	Label	Train	Dev	Test
M	1809	501	1493	M	546	90	442	C	1882	420	2182
H	862	218	756	H	557	135	423	A	1160	280	1368
Total	2671	719	2249	Total	1103	225	865	Total	3042	700	3550
	(12)	(3)	(10)		(48)	(10)	(41)		(No speaker info)		
	5639 (25), 2:50:20 s				2193 (99), 3:27:35 s				7292, 3:12:16 s		

Table 1: General characteristics of the considered data sets and their utterance-level labelling. Number of speakers is specified in parentheses. Utterance labels: H - human-, M - machine-, A - adult-, C - child-directed. It is assumed that H = A and M = C.

If y_i and y_j from Equation 2 are different hard targets (one-hot vectors) of a classification problem, y_{art} will be a soft target. This solution provides better model regularisation and generalisation over various classes and partially resolves the problem of imbalanced data.

We declare another mixup parameter k that defines the proportion of the number of artificial examples that should be generated and the number of real examples. When merging n corpora, we generate one batch from each corpus, increasing the amount of training data in n times without using mixup. If we simultaneously apply mixup, artificial batches are generated on the fly from n real batches, increasing the amount of training data in $n(k+1)$ times without any considerable delays in the training process. In most of the mixup applications investigated by Zhang et al. (2017), α lies within the interval $[0.1, 0.5]$, i.e., the algorithm biases toward original examples and thereby generates more realistic artificial ones. We use constant α and k values that equal 0.5 and 2 respectively. For greater α values, mixup leads to underfitting.

4 Corpora

We examine our models on the audio data of the three corpora mentioned above. The VACC data set contains experimental conversations in German between a user, a confederate, and an Echo Dot Amazon Alexa device (Siegert et al., 2018). The SVC data set was collected within large-scale Wizard-of-Oz (WOZ) experiments and consists of realistic conversations in German between a user, a confederate, and a mobile SDS (Batliner et al., 2008). For compatibility with the other corpora, we consider the two-class SVC problem introduced by Batliner et al. (2008). The HB data set contains spoken conversations in English between a child and a group of adults recorded under real-life conditions (Casillas et al., 2017). Each corpus was split into a training, a development, and a test

set at a proportion defined by its developers. There was no development set specified for SVC by Batliner et al. (2008), and therefore we use 20% of the speakers from its initial training set as a development set. The HB test labels are unavailable to us since this corpus was a part of the Interspeech 2017 ComParE Challenge (Schuller et al., 2017) that has already been finished (none of the participants managed to surpass the Addressee Sub-Challenge baseline). Therefore, we use its development set as a new test set and also utilise 20% of the utterances from its initial training set as a new development set. The partitions of the considered corpora are presented in Table 1. A kernel density estimation (KDE) is depicted in Figure 2 for the utterance length distribution of each corpus.




Figure 2: Kernel density estimation (KDE) of the utterance length distributions.

5 Preliminary Experiments with Linear Models

5.1 Feature Selection

Before training the neural network-based classifiers, we conduct preliminary experiments with the func model, aiming to estimate the degree of similarity between the corpora. After feature extraction with the ComParE configuration, we perform recursive feature elimination (RFE), using the coefficients of the normal vector of a linear SVM as attribute weights similarly to Akhtiamov et al. (2017a). Figure 3a demonstrates RFE curves

obtained by applying ten-fold leave-one-speaker-group-out cross-validation (LOGSO) on each corpus without its test set. The resulting performance is calculated as unweighted average recall (UAR) for comparability with the existing studies and averaged over all folds for each reduced feature set. A feature set is considered to be optimal if further RFE leads to a stable performance loss. For each corpus, we choose one optimal feature set obtained on a random fold and analyse their intersection depicted in Figure 3b. The representative acoustic attributes vary essentially: VACC, SVC, and HB have only 450, 2020, and 400 relevant features out of 6373 respectively, while having only 28 features in common: some functionals over *F0final sma*, *audSpec Rfilt sma*, *mfcc sma*, *pcm fftMag spectralRollOff25.0 sma*, *pcm fftMag spectralRollOff50.0 sma*, *voicingFinalUnclipped sma*, and their *deltas* (Eyben, 2015). Besides these attributes, VACC and SVC have only 172 features in common, though these two corpora have the same target classes. The optimal feature set size for SVC is considerably greater than for the other two corpora. This difference was probably caused by the WOZ modelling of SVC dialogues as the WOZ setup did not seem convincing enough to some users, resulting in fuzzy addressee patterns that concerned a greater number of acoustic features.




Figure 3: Preliminary analysis: performance losses during RFE (a), and optimal feature set comparison (b).

5.2 Cross-Corpus and Multitask Classification

We conduct a series of cross-corpus and multitask experiments with the func model, applying a leave-one-corpus-out (LOCO) and an inverse LOCO scheme. In the first scheme, the model is trained on a mixture of all the corpora but one and tested on each of the three corpora. In the second scheme, the model is trained on one corpus and tested on each of the three corpora. In both

cases, the model is trained and tested on the corresponding partitions from Table 1. In these experiments, we do not perform feature selection and do not use mixup. Results of the two experimental series are depicted in Figure 4. Let us denote the matrix from Figure 4a as \bar{A} , its element as $\bar{a}_{i,j}$, the matrix from Figure 4b as \bar{B} , and its element as $\bar{b}_{i,j}$. The resulting UAR ($\bar{a}_{2,2}$) and the optimal feature set size on SVC slightly differ from those obtained by Akhtiamov et al. (2017a) since we apply statistical corpus normalisation in the present study instead of speaker normalisation in order to make our results fairer as the system may face unknown speakers in real applications. Furthermore, there is no information regarding speakers available for HB. $\bar{a}_{1,2}$ and $\bar{a}_{2,1}$ are considerably greater than the other off-diagonal elements of \bar{A} , demonstrating a clear relation between VACC and SVC. This result motivates us to explore the potential of the cross-corpus data augmentation on VACC and SVC by means of mixup and deep learning in our future experiments. \bar{A} does not reveal any relation between HB and the other two corpora, though an interesting trend may be noted in \bar{B} . $\bar{b}_{2,1}$ and $\bar{b}_{3,1}$ are similar to $\bar{a}_{1,1}$, $\bar{b}_{1,2}$ and $\bar{b}_{3,2}$ are close to $\bar{a}_{2,2}$, and $\bar{b}_{1,3}$ and $\bar{b}_{2,3}$ are similar to $\bar{a}_{3,3}$. Altogether, these three results mean that a single func model trained on examples from two arbitrary corpora demonstrates an adequate performance on them both as if the model were trained on each corpora separately or, in other words, that the three classification problems are non-contradictory. However, A-C AD turned out to be essentially more challenging than H-M AD.

		Test		
		(1)	(2)	(3)
Train	(1)	0.788	0.605	0.511
	(2)	0.614	0.770	0.516
	(3)	0.527	0.552	0.602

(a)

		Test		
		(1)	(2)	(3)
Exclude Train	(1)	0.560	0.754	0.604
	(2)	0.758	0.616	0.585
	(3)	0.783	0.756	0.501

(b)

Figure 4: Results of the inverse LOCO (a) and LOCO (b) experiments with the func model. All values are presented in terms of UAR. Corpora: (1) - VACC, (2) - SVC, (3) - HB.

6 Experiments with Neural Networks

6.1 Mixup and Acoustic Context Length

All the experiments below are presented in terms of UAR for comparability with the existing studies. All statistical comparisons are drawn apply-




Figure 5: Classification performance of the LLD model over various context windows and its trends after data augmentation on the considered corpora. In each of the three cases, the training set of the target corpus (on the test set of which UAR is measured) is mixed with itself (*mix(corpus)*) or with itself and with the training set of another corpus (*mix(corpus, another_corpus)*). The points connected with spline interpolation denote exact measurements.

ing a *t*-test with a significance level of 0.05. First, we analyse the sensitivity of our neural networks to acoustic context length variations. This hyperparameter was shown to be critical for paralinguistic problems (Fedotov et al., 2018a). We take a context window length of 1 s as a reference point and then vary it by raising to different powers of two. The context windows are cut out of the audio files with an overlap of 75%. This preprocessing partially resolves the lack of training data. It is possible to align the obtained logarithmic scale with basic acoustic units: given the mean syllable duration estimated by Greenberg (1999) for spontaneous English, we roughly assume that the time intervals between 0, 0.125, 0.5, 1, 2, and 8 s correspond to allophones, syllables, words, collocations/syntagmas, and utterances respectively. In fact, these intervals may significantly overlap since syllable duration is known to be highly speaker-dependent (Greenberg et al., 2003). German words and more complex acoustic units have longer durations compared to their English equivalents.

Performance curves of the LLD classifier tested on LOSGO are depicted in Figure 5. The resulting UAR values are averaged over all folds. The dashed curve is located above the solid one in all three cases, i.e., mixup results in a significant performance improvement already when applied to the same corpus. Adding another corpus to the mixup procedure influences the performance depending on a context window length. Mix(VACC, SVC) significantly surpasses mix(VACC) on VACC for a context window of 2 s. Mix(SVC, VACC) significantly outperforms mix(SVC) on SVC for a context window of 0.5 s. A possible explanation for these two results

is that SVC has generally longer utterances (Figure 2) and probably longer acoustic addressee patterns compared to VACC. Mix(HB) does not benefit from adding another corpus to the mixup procedure.

The curves from Figure 5a flatten beyond 0.5 s, meaning that VACC is less sensitive to context length variations than SVC and HB. The optimal context window length, which provides the highest UAR, is 2 s for VACC and SVC and 1 s for HB. However, the latter corpus demonstrates virtually the same result for a longer window of 2 s. The e2e model shows a similar behaviour on various context windows and reaches the highest UAR for the same context window of 2 s on all three corpora. This fact motivates us to confine to a single context window length of 2 s in our future experiments that corresponds to acoustic patterns at the utterance level. Our results confirm an earlier conclusion drawn by Shriberg et al. (2013) regarding the optimal acoustic context length for H-M AD in English.

Table 2 contains the exact UAR values of the two-second performance slices for both neural networks. Similarly to the results presented in Figure 5, the values from Table 2 are obtained on LOSGO and averaged over all folds. The LLD model demonstrates a higher performance com-

Test Corpus	Model	—	---	-.-.-	mix (all)
VACC	LLD	.879	.890	.901	.873	.886
	e2e	.853	.834	.852	.845	.846
SVC	LLD	.813	.823	.804	.795	.818
	e2e	.764	.756	.758	.749	.761
HB	LLD	.631	.645	.627	.640	.636
	e2e	.647	.632	.633	.616	.631

Table 2: Two-second UAR slices. Each marker corresponds to a curve of the same style in Figure 5.

pared to the e2e model overall, except HB, on which both classifiers behave similarly. In contrast to the LLD model, the e2e classifier does not benefit from mixup. This result contradicts the supposition made by Zhang et al. (2017) to apply mixup to raw speech data and may be naturally explained in the following way: after applying mixup to raw speech signals, our augmented data sounds like crowd noise that confuses the e2e model being unable to handle the cocktail party effect. This is not the case for some handcrafted features, e.g., logarithmic attributes, as applying mixup to them does not necessarily mean a simple overlapping of two waveforms, from which these features were extracted. We conclude that applying mixup makes more sense for acoustic features of a higher abstraction level than raw data, e.g., handcrafted LLDs or features extracted with a CNN. In the present study, we confine to two extreme cases: handcrafted LLDs and raw waveforms.

6.2 Cross-Corpus and Multitask Classification

The experiments below are conducted on the partitions specified in Table 1. Six series of cross-corpus experiments are depicted as performance matrices in Figure 6. Let us denote the matrix from Figure 6a as A and its element as $a_{i,j}$, the matrix from Figure 6b as B and its element as $b_{i,j}$, etc. A and B show inverse LOCO experiments on the LLD model with mixup and on the e2e model without mixup respectively. $a_{1,2}$ and $a_{2,1}$ are considerably greater than the other off-diagonal elements of A . $b_{1,2}$ and $b_{2,1}$ are also significantly greater than the other off-diagonal elements of B . Similarly to the matrix \bar{A} from Figure 4a, these two results demonstrate a clear relation between VACC and SVC that was better captured with the e2e model. The other four matrices from Figure 6 contain results of LOCO experiments: C and D - without mixup, E and F - with mixup. The elements $c_{1,3}$, $c_{2,3}$, $d_{1,3}$, and $d_{2,3}$ are close to a random-choice UAR of 0.5, meaning that both neural networks perceive HB as noise and completely ignore it in favour of another corpus. However, the situation changes if we apply mixup: the elements $e_{1,3}$ and $e_{2,3}$ are similar to $a_{3,3}$ as well as the elements $f_{1,3}$ and $f_{2,3}$ being close to $b_{3,3}$. These two results mean that both neural networks start perceiving both corpora in-

		Test		
		(1)	(2)	(3)
Train	(1)	0.870	0.605	0.520
	(2)	0.609	0.789	0.513
	(3)	0.500	0.530	0.633

(a) LLD+mix, inverse LOCO.

		Test		
		(1)	(2)	(3)
Train	(1)	0.823	0.645	0.556
	(2)	0.665	0.712	0.534
	(3)	0.559	0.559	0.640

(b) e2e, inverse LOCO.

		Test		
		(1)	(2)	(3)
Exclude Train	(1)	0.605	0.684	0.517
	(2)	0.874	0.603	0.521
	(3)	0.885	0.786	0.538

(c) LLD, LOCO.

		Test		
		(1)	(2)	(3)
Exclude Train	(1)	0.663	0.692	0.544
	(2)	0.794	0.588	0.507
	(3)	0.813	0.694	0.504

(d) e2e, LOCO.

		Test		
		(1)	(2)	(3)
Exclude Train	(1)	0.596	0.751	0.622
	(2)	0.885	0.556	0.626
	(3)	0.891	0.785	0.518

(e) LLD+mix, LOCO.

		Test		
		(1)	(2)	(3)
Exclude Train	(1)	0.540	0.681	0.611
	(2)	0.780	0.603	0.625
	(3)	0.816	0.710	0.543

(f) e2e+mix, LOCO.

Figure 6: Results of the inverse LOCO and LOCO experiments with the neural networks. All values are presented in terms of UAR. Corpora: (1) - VACC, (2) - SVC, (3) - HB.

involved in the mixup procedure as efficiently as if the networks were trained on each data set separately. Due to a simpler model architecture, the func classifier did not face such a problem of overfitting to a specific corpus during the experiments with multitask learning presented in Figure 4b.

A similar trend may be noted in Figure 7 that demonstrates experiments on merging all three corpora: if trained on all the corpora without mixup, both LLD and e2e models discriminate SVC and completely ignore HB. Mixup allows us to train a multipurpose neural network that performs equally well on each of the corpora as if there were three networks trained exclusively for single tasks. The classification performance ob-




Figure 7: Results of the experiments on merging all three corpora.

tained on VACC and HB with the func model is generally lower compared to the results of the neural networks, and mixup is unable to improve it. However, the func classifier does not suffer from overfitting to a specific corpus during multitask learning and does not need to be regularised.

7 Experiments with ASR-Based Metafeatures

Some metafeatures obtained from an automatic speech recogniser (ASR) are useful for H-M AD since people speak more clearly than usual when addressing an SDS. Machine-directed speech tends to match the ASR patterns better compared to human-directed speech, resulting in a higher ASR confidence (Tsai et al., 2015). It is interesting to check this approach on A-C AD. Using the Google Cloud ASR for German (on VACC and SVC) and for English (on HB), we extract the following ASR metafeatures at the utterance level: *confidence of the best hypothesis*, *number of hypotheses*, *number of words in the best hypothesis*, and *utterance duration in seconds*. These features except the first one (it is already normalised) are brought to zero mean and unit variance and fed to an SVM with a radial kernel (Hofmann and Klinkenberg, 2013). The UAR values obtained with this classifier on the test partitions from Table 1 are equal to 0.778, 0.657, and 0.515 for VACC, SVC, and HB respectively. The latter value is slightly above a random-choice UAR of 0.5, meaning that ASR confidence is non-representative for A-C AD.

8 Conclusions and Future Work

The H-M and A-C AD problems turned out to be essentially different in certain aspects. The first aspect concerns the previously discussed acoustic patterns of child- and machine-directed speech. On the one hand, none of the considered models managed to reveal any relations between HB and the other two corpora during our inverse LOCO experiments. On the other hand, the LOCO experiments with the linear model demonstrate that the H-M and A-C AD problems are non-contradictory. The second aspect is connected with the degree of how often misunderstanding situations occur in an H-M conversation. People tend to talk to the system in a normal manner in the absence of such situations, and this manner of speech may be acoustically undistinguishable

from human-directed speech. The third aspect concerns what is said during an A-C conversation. Adults' speech often contains separate sounds and intonations and no verbal information when they talk to children, and therefore ASR confidence is non-representative for A-C AD, though it is useful for H-M AD.

Mixup has been shown to be beneficial for neural networks using predefined acoustic features, while not giving any significant performance improvement for e2e models, though Zhang et al. (2017) supposed that it is worth applying the method to raw speech data as well. Linear classifiers do not benefit from mixup neither due to their simple architectures that do not require any regularisation. Another remarkable capability of mixup was revealed in multitask experiments and applies to both handcrafted features and raw data. This method allows us to merge several corpora modelling similar classification tasks in such a way that one neural network trained on this mixture solves all the tasks equally efficiently with single neural networks, each of which was trained on its own corpus. The corpora being utilised for multitask learning may essentially differ, e.g., VACC and SVC were collected in completely different domains, and HB was even collected for another task and uttered in another language. Without mixup, the neural network overfits to the corpus with the strongest correlation between its features and labels (VACC) and starts discriminating the other corpora. Linear models do not suffer from this problem, though they demonstrate a lower classification performance overall.

Two-second speech fragments are optimal for AD and correspond to acoustic patterns at the utterance level. This result confirms an earlier conclusion drawn by Shriberg et al. (2013) regarding H-M AD in English. According to our inverse LOCO experiments, there exists a clear relation between VACC and SVC. Furthermore, applying mixup to these two corpora allows us to improve classification results on VACC significantly. The following UAR values may be taken from Figure 6 as the new baselines: $e_{3,1} = 0.891$ for VACC and $b_{3,3} = 0.640$ for HB. $b_{3,3}$ is the best baseline for standalone classifiers compared to the results introduced by Schuller et al. (2017) on the original HB development set. Our e2e model surpasses the one from (Schuller et al., 2017) that demonstrated a UAR of 0.609. We achieved this performance

improvement due to a more careful choice of the CNN architecture. $a_{2,2} = 0.789$ is similar to the latest SVC baseline of 0.800 established by Akhtiamov et al. (2017a).

In our future work, we plan to extend our experiments, applying mixup to two-dimensional spectrograms and to features extracted with a CNN.

Acknowledgements

The research presented in Section 5 and Section 7 was conducted exclusively within the framework of the Russian Science Foundation project (No. 18-11-00145). The remaining part of the research was supported by DAAD jointly with the Ministry of Science and Higher Education of the Russian Federation within the Michail Lomonosov Program (project No. 8.704.2016/DAAD).

References

- Martín Abadi et al. 2016. Tensorflow: A system for large-scale machine learning. In *Symposium on Operating Systems Design and Implementation (OSDI)*, pages 265–283.
- Oleg Akhtiamov and Vasily Palkov. 2018. Gaze, prosody and semantics: Relevance of various multimodal signals to addressee detection in human-human-computer conversations. In *International Conference on Speech and Computer (SPECOM)*, pages 1–10. Springer.
- Oleg Akhtiamov, Maxim Sidorov, Alexey Karpov, and Wolfgang Minker. 2017a. Speech and text analysis for multimodal addressee detection in human-human-computer interaction. In *Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pages 2521–2525. ISCA.
- Oleg Akhtiamov, Dmitrii Ubskii, Evgeniia Feldina, Aleksei Pugachev, Alexey Karpov, and Wolfgang Minker. 2017b. Are you addressing me? Multimodal addressee detection in human-human-computer conversations. In *International Conference on Speech and Computer (SPECOM)*, pages 152–161. Springer.
- Yusuf Aytar, Carl Vondrick, and Antonio Torralba. 2016. [SoundNet: Learning sound representations from unlabeled video](#). In *Advances in Neural Information Processing Systems 29*, pages 892–900.
- Anton Batliner, Christian Hacker, and Elmar Nöth. 2008. To talk or not to talk with a computer. *Journal on Multimodal User Interfaces*, 2(3):171–186.
- Carlos Busso, Panayiotis Georgiou, and Shrikanth Narayanan. 2007. Real-time monitoring of participants’ interaction in a meeting using audio-visual sensors. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 2, pages 685–688. IEEE.
- Marisa Casillas, Andrei Amatuni, Amanda Seidl, Melanie Soderstrom, Anne Warlaumont, and Elika Bergelson. 2017. What do babies hear? Analyses of child- and adult-directed speech. In *Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pages 2093–2097. ISCA.
- Florian Eyben. 2015. *Real-time speech and music classification by large audio feature space extraction*. Springer.
- Florian Eyben, Felix Weninger, Florian Gross, and Björn Schuller. 2013. Recent developments in opensmile, the munich open-source multimedia feature extractor. In *ACM international conference on Multimedia*, pages 835–838. ACM.
- Dmitrii Fedotov, Denis Ivanko, Maxim Sidorov, and Wolfgang Minker. 2018a. Contextual dependencies in time-continuous multidimensional affect recognition. In *International Conference on Language Resources and Evaluation (LREC)*, pages 1220–1224. ELRA.
- Dmitrii Fedotov, Heysem Kaya, and Alexey Karpov. 2018b. Context modeling for cross-corpus dimensional acoustic emotion recognition: Challenges and mixup. In *International Conference on Speech and Computer (SPECOM)*, pages 155–165. Springer.
- Emer Gilmartin, Benjamin R Cowan, Carl Vogel, and Nick Campbell. 2018. Explorations in multiparty casual social talk and its relevance for social human machine dialogue. *Journal on Multimodal User Interfaces*, 12(4):297–308.
- Steven Greenberg. 1999. Speaking in shorthand – A syllable-centric perspective for understanding pronunciation variation. *Speech Communication*, 29(2-4):159–176.
- Steven Greenberg, Hannah Carvey, Leah Hitchcock, and Shuangyu Chang. 2003. Temporal properties of spontaneous speech – A syllable-centric perspective. *Journal of Phonetics*, 31(3):465–485.
- Fasih Haider, Hayakawa Akira, Saturnino Luz, Carl Vogel, and Nick Campbell. 2018. On-talk and off-talk detection: A discrete wavelet transform analysis of electroencephalogram. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 960–964. IEEE.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Markus Hofmann and Ralf Klinkenberg. 2013. *Rapid-Miner: Data mining use cases and business analytics applications*. CRC Press.

- Sergey Ioffe and Christian Szegedy. 2015. [Batch normalization: Accelerating deep network training by reducing internal covariate shift](#). *CoRR*, abs/1502.03167.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *CoRR*, abs/1412.6980.
- Thao Le Minh, Nobuyuki Shimizu, Takashi Miyazaki, and Koichi Shinoda. 2018. Deep learning based multi-modal addressee recognition in visual scenes with utterances. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1546–1553.
- Sri Harish Mallidi, Roland Maas, Kyle Goehner, Ariya Rastrow, Spyros Matsoukas, and Björn Hoffmeister. 2018. Device-directed utterance detection. In *Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pages 1225–1228. ISCA.
- Ivan Medennikov, Yuri Khokhlov, Aleksei Romanenko, Dmitry Popov, Natalia Tomashenko, Ivan Sorokin, and Alexander Zatvornitskiy. 2018. An investigation of mixup training strategies for acoustic models in ASR. In *Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pages 2903–2907. ISCA.
- Hiroki Ouchi and Yuta Tsuboi. 2016. [Addressee and response selection for multi-party conversation](#). In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2133–2143, Austin, Texas. ACL.
- Aleksei Pugachev, Oleg Akhtiamov, Alexey Karpov, and Wolfgang Minker. 2017. Deep learning for acoustic addressee detection in spoken dialogue systems. In *Conference on Artificial Intelligence and Natural Language (AINL)*, pages 45–53. Springer.
- Björn Schuller, Stefan Steidl, Anton Batliner, Simone Hantke, Florian Hönl, Juan Rafael Orozco-Arroyave, Elmar Nöth, Yue Zhang, and Felix Weninger. 2015. The interspeech 2015 computational paralinguistics challenge: nativeness, parkinson’s & eating condition. In *Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pages 478–482. ISCA.
- Björn Schuller et al. 2017. The interspeech 2017 computational paralinguistics challenge: Addressee, cold & snoring. In *Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pages 3442–3446. ISCA.
- Elizabeth Shriberg, Andreas Stolcke, Dilek Hakkani-Tür, and Larry Heck. 2012. Learning when to listen: Detecting system-addressed speech in human-human-computer dialog. In *Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pages 334–337. ISCA.
- Elizabeth Shriberg, Andreas Stolcke, and Suman V Ravuri. 2013. Addressee detection for dialog systems using temporal and spectral dimensions of speaking style. In *Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pages 2559–2563. ISCA.
- Ingo Siegert and Julia Krüger. 2018. How do we speak with Alexa - Subjective and objective assessments of changes in speaking style between HC and HH conversations. *Kognitive Systeme*, 1.
- Ingo Siegert, Julia Krüger, Olga Egorow, Jannik Nietzold, Ralph Heinemann, and Alicia Lotz. 2018. Voice assistant conversation corpus (VACC): A multi-scenario dataset for addressee detection in human-computer-interaction using Amazon ALEXA. In *LREC 2018 Workshop "LB-ILR2018 and MMC2018 Joint Workshop"*, pages 51–54. ELRA.
- Anastasiia Spirina, Olesia Vaskovskaia, Maxim Sidorov, and Alexander Schmitt. 2016. Interaction quality as a human-human task-oriented conversation performance. In *International Conference on Speech and Computer (SPECOM)*, pages 403–410. Springer.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958.
- George Trigeorgis, Fabien Ringeval, Raymond Brueckner, Erik Marchi, Mihalis Nicolaou, Björn Schuller, and Stefanos Zafeiriou. 2016. Adieu features? End-to-end speech emotion recognition using a deep convolutional recurrent network. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5200–5204. IEEE.
- TJ Tsai, Andreas Stolcke, and Malcolm Slaney. 2015. A study of multimodal addressee detection in human-human-computer interaction. *IEEE Transactions on Multimedia*, 17(9):1550–1561.
- Hongyi Zhang, Moustapha Cissé, Yann N. Dauphin, and David Lopez-Paz. 2017. [Mixup: Beyond empirical risk minimization](#). *CoRR*, abs/1710.09412.
- Rui Zhang, Honglak Lee, Lazaros Polymenakos, and Dragomir Radev. 2018. Addressee and response selection in multi-party conversations with speaker interaction RNNs. In *AAAI Conference on Artificial Intelligence*, pages 5690–5697. AAAI.

A Scalable Method for Quantifying the Role of Pitch in Conversational Turn-Taking

Kornel Laskowski^{1,2} and Marcin Włodarczak¹ and Mattias Heldner¹

¹ Stockholm University, Stockholm, Sweden

² Voci Technologies, Inc., Pittsburgh PA, USA

Abstract

Pitch has long been held as an important signalling channel when planning and deploying speech in conversation, and myriad studies have been undertaken to determine the extent to which it actually plays this role. Unfortunately, these studies have required considerable human investment in data preparation and analysis, and have therefore often been limited to a handful of specific conversational contexts. The current article proposes a framework which addresses these limitations, by enabling a scalable, quantitative characterization of the role of pitch throughout an entire conversation, requiring only the raw signal and speech activity references. The framework is evaluated on the Switchboard dialogue corpus. Experiments indicate that pitch trajectories of both parties are predictive of their incipient speech activity; that pitch should be expressed on a logarithmic scale and Z -normalized, as well as accompanied by a binary voicing variable; and that only the most recent 400 ms of the pitch trajectory are useful in incipient speech activity prediction.

1 Introduction

Intonation is generally held to play an integral role in the phonetic realization of turns and in the prediction of more talk (see e.g. (Bögels and Torreira, 2015) for a review). There is broad consensus that flat pitch segments are associated with turn-holding and that rising or falling pitch segments are associated with turn-yielding (Bögels and Torreira, 2015; Caspers, 2003; Duncan, 1972; Edlund and Heldner, 2005; Ford and Thompson, 1996; Heldner et al., 2009; Heldner and Włodarczak, 2015; Hjalmarsson, 2011; Jefferson, 1984; Kane et al., 2014; Koiso et al., 1998; Laskowski et al., 2009; Local et al., 1986; Selting, 1996; Yanushkevskaya et al., 2014; Zellers, 2013, 2017). Studies considering finer-grained categories of the

pitch contour (Gravano and Hirschberg, 2011; Wennerstrom and Siegel, 2003), additionally including slowly rising and slowly falling pitch, have tended to corroborate those findings. Furthermore, they indicate that the endpoint of a pitch segment is relevant, associating segments reaching the top or bottom of a speaker’s range with turn-yielding and those ending near the middle of the range with turn-holding.

The converging results of so many studies are astonishing given the methodological differences between them, with regard to the speech material (spontaneous vs. task-oriented) and to the pitch-contour categorization method (perceptual judgements, acoustic measurements, or phonologically motivated categories). Perhaps more importantly, the studies in question differ in how the pitch contour is parametrized (e.g. perceptual stylization, functional data analysis, linear or polynomial curve fitting, linear or logarithmic scale), how far back in the speech interval relevant pitch cues are to be found, as well as how cues are evaluated (e.g. perceptually vs. statistically).

It is therefore not very surprising that work which has tried to verify the above claims with acoustic measurements of fundamental frequency (F_0) has also produced some mixed results (see e.g. (Zellers, 2017; Walker, 2017) for reviews). A variety of explanations for these mixed results are believed to exist. First, it has been hypothesized that non-pitch cues may play a more important role than do pitch cues (e.g. (Local and Walker, 2012; Walker, 2017; Zellers, 2017)). Second, it is possible that the role of intonation varies with the communicative situation, and that it is strongly dependent on the number of participants, whether the participants have eye contact, whether the participants know one another, etc. Finally, there may be considerable language-, dialect-, and domain-specific differences in the role of pitch in turn-

taking. At the present time, these explanations continue to be mere hypotheses which — owing to the many methodological differences in published work — cannot be easily evaluated.

The main focus of the current article is to render the evaluation and comparison of such hypotheses tractable, if not outright easy. A key requirement is that the proposed method be *scalable*, i.e. capable of ingesting sufficiently large quantities of conversational material to generate representative results. This in turn requires that it not rely on time-consuming, often-contentious annotation of either turn or pitch phenomena — authors of existing research do not always agree on what constitutes a turn, for example. Furthermore, the method needs to be *quantitative* if it is to permit strict comparison. The method proposed in the current article is both scalable and quantitative; it relies only on the availability of the raw signal and accurate speech activity references, per conversation and per conversation-side. It is presented in Section 3.

To evaluate the method itself, the current article asks three key questions of a large, oft-studied corpus of telephone conversations in English (described in Section 2). These questions are:

Q1. Can attention to pitch reduce the average surprise of incipient speech activity?

Q2. What is the optimal representation of pitch for a speech prediction system?

Q3. How far back into the past should a pitch-sensitive method look?

Experiments described in Section 4 demonstrate that *Q1* can be answered in the affirmative, that binary voicing and Z-normalized log-pitch offer the best results when used together (*Q2*), and that only the most recent 400 ms of pitch history are sufficient (*Q3*). Furthermore, the proposed system is able to answer *Q1* and *Q3* in a nearly fully-automated fashion, for evidently any corpus; the answer to *Q2* may require human-mediated investigations, for which the proposed system provides a suitable and convenient framework.

2 Data

Experiments used the Switchboard-1 Corpus, as re-released in 1997 (Godfrey and Hollimann, 1997). The corpus consists of 2435 dyadic telephone conversations, each approximately 10 minutes in duration. It was iteratively divided into three speaker-disjoint sets as in (Laskowski and

Shriberg, 2012), such that TRAINSET, DEVSET, and TESTSET consist of 762, 227, and 199 conversations, respectively. During the division process, it was not possible to allocate 1247 of the Switchboard-1 conversations, because each of their two speakers had already been placed in different sets. Forced alignments of the manually transcribed words (used as discussed in Subsection 3.2) for both sides of the conversation were provided in (Deshmukh et al., 1998).

3 Methods

This article proposes a means of quantifying the extent to which pitch, represented in a variety of ways, reduces the surprise induced by observing the temporal distribution of speech in unseen conversations. Such a means involves a probabilistic formulation of the problem (Subsection 3.1), a method for obtaining instantaneous binary speech activity (Subsection 3.2), a method for measuring pitch (Subsection 3.3), and a model for approximating the probabilities given those features, together with a metric for quantifying model performance (Subsection 3.4).

3.1 Stochastic Turn Taking

As in (Laskowski, 2012), the methodology employed here relies on forming a probability distribution over the side-attributed speech activity in entire dyadic conversations. This eliminates a dependency on the specific definition of a turn; the resulting probability models attempt to account for all speech, effectively marginalizing out alternative definitions of what turns are and where they start and end.

The most direct means of modeling conversations for this purpose is to discretize their temporal extent; here, a frame-step and size of 100 ms is used, representing approximately half of a normative syllable. Such discretization results in a $K \times N$ *chronogram* for each conversation, ie.

$$\mathbf{Q} = \begin{bmatrix} \dots & \blacksquare\blacksquare\blacksquare\blacksquare\blacksquare\blacksquare\blacksquare & \dots \\ \dots & \square\square\square\blacksquare\blacksquare\blacksquare\blacksquare & \dots \end{bmatrix}, \quad (1)$$

where the k th row, $1 \leq k \leq K$, represents the speech activity of one of the $K = 2$ sides to the conversation, and each column \mathbf{q}_n , $1 \leq n \leq N$, represents one 100-ms interval. Each $\mathbf{q}_n[k] \in \{\square, \blacksquare\} \equiv \{0, 1\}$, indicating that the k th party is either not-speaking or speaking in frame n , respectively.

The probability \mathcal{P} of a given \mathbf{Q} is then given by

$$\mathcal{P}(\mathbf{Q}) = \prod_{n=1}^N \mathcal{P}(\mathbf{q}_n | \mathbf{q}_1^{n-1}) \quad (2)$$

$$\approx \prod_{n=1}^N \mathcal{P}(\mathbf{q}_n | \mathbf{q}_{n-\tau}^{n-1}) \quad (3)$$

$$\approx \prod_{n=1}^N \prod_{k=1}^K \mathcal{P}(\mathbf{q}_n[k] | \mathbf{q}_{n-\tau}^{n-1}) \quad (4)$$

where Equation 3 represents a Markovian truncation of the history to the most recent τ frames, and Equation 4 assumes that participants are conditionally independent of one another in any given frame, but dependent on their joint past $\mathbf{q}_{n-\tau}^{n-1}$. The term *target participant* is used to refer to that side of the conversation for which the interior factor on the right-hand-side of the equation is being evaluated; when evaluating the left-hand-side over all cells in \mathbf{Q} , each of the $K = 2$ participants becomes the target participant half the time.

In this framework, quantifying the impact of pitch — or any other side information available in $K \times N$ matrix form as \mathbf{X} — entails comparing the probability in Equation 4 to

$$\mathcal{P}(\mathbf{Q}|\mathbf{X}) \approx \prod_{n=1}^N \prod_{k=1}^K \mathcal{P}(\mathbf{q}_n[k] | \mathbf{q}_{n-\tau}^{n-1}, \mathbf{x}_{n-\tau}^{n-1}) .$$

By excluding the current and future \mathbf{x}_n^N from the conditioning context, the factor $\mathcal{P}(\mathbf{q}_n[k] | \mathbf{q}_{n-\tau}^{n-1}, \mathbf{x}_{n-\tau}^{n-1})$ is observed to be a causal prediction.

3.2 Speech Activity

The above equation forms a probability density over speech activity \mathbf{Q} that *actually happened*, rather than speech activity that can be measured. The most accurate means currently available for producing \mathbf{Q} is to perform forced time-alignment of the k th participant’s audio channel to the words spoken by that participant. The resulting word boundaries are then aligned to the 100-ms frame boundaries which define \mathbf{Q} , and each $\mathbf{q}_n[k]$, $1 \leq n \leq N$ and $1 \leq k \leq K$, is assigned to 1 if the k th participant was speaking for 50% or more of the temporal support of the n th frame.

3.3 Pitch

Pitch was extracted using the `get_f0` implementation available in the Snack Sound Toolkit

(Sjölander, 2001). In order to avoid contagion from the future (pitch tracking uses context to smooth candidate per-frame fundamental frequency estimates), a separate pitch track was extracted for the τ -duration conditioning context of every frame n in every channel k of every conversation¹. Snack’s default frame step is 10 ms; the resulting sequence of 10-ms pitch estimates was then aligned to the 100-ms frames in \mathbf{Q} , yielding side-information \mathbf{P} . Each cell $\mathbf{p}_n[k]$ of \mathbf{P} was assigned to the mean of those voiced 10-ms pitch estimates of the k th participant’s speech which fell entirely within the temporal support of frame n ; it therefore sufficed for only one 10-ms-frame to be deemed as voiced by Snack in order for the 100-ms frame in \mathbf{P} to be considered voiced²; unvoiced frames in \mathbf{P} were assigned to NaN.

Note that pitch computed as described above may exhibit doubling and halving errors; the exploration of the impact of (manually) corrected pitch is beyond the scope of the current article. Similarly, phenomena such as diplophonia and creakiness are not explicitly treated.

3.4 Models and Metrics

The prediction probabilities described in Subsection 3.1 were approximated using a feed-forward neural network

$$\mathcal{P}(\mathbf{q}_n[k] | \mathbf{q}_{n-\tau}^{n-1}, \mathbf{x}_{n-\tau}^{n-1}) \approx f(\mathbf{q}_{n-\tau}^{n-1}, \mathbf{x}_{n-\tau}^{n-1})$$

with one hidden layer of H tanh units³, and one sigmoid output unit — representing the probability that the k th participant is speaking at frame n . For most experiments in the current article, $H = 8$. Note that the network has no recurrence since determining the exact extent of the usefully conditioning history is of primary interest. Network weights were trained on TRAINSET, using

¹This brute-force and seemingly inefficient approach proved to have considerable impact on the numerical results presented in Section 4, indicating that basing incremental predictions on non-incremental pitch extraction would have been a form of cheating.

²Other policies were explored, notably that in which at least half of the 10-ms frames need to be voiced; the results exhibited the same trends as those reported here, although numerically the cross entropy rates were slightly larger. It appears that better predictions are possible when more of the 100-ms frames in \mathbf{P} are deemed voiced, even when some of those cells are more sensitive to outliers in the underlying 10-ms pitch trajectory.

³Note that tanh activation units in the network implicitly map NaN features to zero. This approach is likely sub-optimal, but provides a well-understood and simple-to-train baseline for improvements like those described in (Laskowski, 2015).

1000 iterations of scaled conjugate gradient (SCG; (Møller, 1993)) descent — a second-order, deterministic rather than stochastic procedure.⁴

The appropriate objective function given a single sigmoid output unit is the cross entropy error (Bishop, 1995); it was used during SCG training as well as in the subsequent evaluation of trained models. Since, for any given conversation and participant, the evaluation of the model for a sequence of frames can be thought of as a causal prediction, during testing the error is henceforth referred to as the *cross-entropy rate*, and is expressed in bits per 100-ms frame.

4 Results

4.1 Representation

The first suite of experiments attempts to identify an optimal representation of pitch for the analysis task at hand. To put the ensuing results into perspective, the baseline is a system which excludes all pitch information; Figure 1 depicts as “ Q^τ ” the achieved cross entropy rate as a function of the number τ of past speech activity frames which comprise the conditioning context. As can be seen, the cross entropy rate exhibits a nearly linear decline over the range $\tau \in [1, 10]$ for all three of TRAINSET, DEVSET, and TESTSET. Q^{10} achieves 0.274371 bits/frame on DEVSET, which is 0.014200 bits/frame lower than the 0.288571 bits/frame achievable when only that target participant’s speech activity is considered (not shown in the figure, but henceforth lowercase q^{10}).

In all subsequent experiments in this subsection, the conditioning context consists of Q_1^{10} — all 10 most recent frames of speech activity from both participants to the conversation — plus the τ most recent frames of one of several representations of pitch for the target participant. The first of these is just P , as computed in Subsection 3.3. As can be seen in Figure 1 (where for notational convenience the lowercase “ p ” indicates target participant only), the most recent frame of pitch P_1^1 by itself already provides an improvement over Q_1^{10} for TRAINSET. It appears that reductions in TRAINSET cross entropy rates begin to asymptote at $\tau = 3$ frames⁵. This indicates that the

⁴For each experimental setting, a single randomly seeded model was trained.

⁵It should be noted that each model at τ , visually connected by a line to the point at $\tau - 1$, contains all of the features of that point. As a result, the curves can reasonably be expected to be monotonically decreasing or asymptotically

proposed model learns to exploit pitch for speech activity prediction, and that therefore recent pitch must be correlated with incipient speech activity in TRAINSET. The fact that the same trends are observed for DEVSET indicates that the correlations which the model learns on TRAINSET generalize to data unseen during model training. The model achieves a cross entropy rate minimum on DEVSET at $\tau = 3$ of 0.270831 bits/frame, which is 0.0035400 bits/frame lower than the best value for Q_1^{10} alone.

Absolute pitch, as represented by P , is patently speaker-dependent; for the model to have successfully leveraged absolute pitch, it must be ignoring a significant portion of the variability observed in P . To quantify this, an experiment was conducted which uses binary voicing V (instead of P), whose elements $v_k[n]$ are unity if the corresponding $p_k[n]$ is non-NaN and zero otherwise. Denoted as “ $Q^{10} \cup v^\tau$ ” in Figure 1, the curve exhibits a minimum on DEVSET at $\tau = 8$ of 0.271698, which 0.0026730 lower than for Q_1^{10} alone and represents 76% of the reduction observed for P . This is surprisingly high and implies that the actual value of absolute pitch is not as relevant for prediction as is its (non-NaN) existence. Exposing the model to both V and P for the target participant (in addition to Q_1^{10}), denoted “ $Q^{10} \cup v^\tau \cup p^\tau$ ” in Figure 1, is seen to lower the cross entropy rate to 0.270304 bits/frame at $\tau = 9$, by 0.000527 bits/frame. It is possible that the availability of V allows the model to focus on extracting information from frames in which absolute pitch is known to exist, and not waste its finite capacity on inferring this by itself.

Since, as expected, variability in absolute pitch P appears to present a problem for the model, an experiment was conducted which Z-normalizes P by each speaker’s mean and standard deviation. These two quantities must be known a priori; assuming that they do not deviate from a speaker’s conversation-specific statistics permits their estimation from each conversation separately. This leads to a new representation, Z , whose elements $z_k[n]$ are equal to $(p_k[n] - \mu_P) / \sigma_P$ where $p_k[n]$ is non-NaN, and NaN otherwise. The curve in Figure 1, denoted “ $Q^{10} \cup z^\tau$ ”, exhibits a DEVSET minimum of 0.270877 bits/frame at $\tau = 8$.

flat. That they are not reflects the effect of random seeding and the fact that each point represents one model rather than an average over multiple, differently-seeded but otherwise-same, models.




Figure 1: Cross entropy rate, along the y -axis in bits per 100-ms frame, for several representations of pitch on top of 10 frames of speech activity from both participants, as a function of the duration of the pitch history, along the x -axis in number of 100-ms frames. Rates are shown from left to right for TRAINSET, DEVSET, and TESTSET.

This is only negligibly different from the minimum of 0.270831 bits/frame achieved for absolute pitch (cf. the previously-discussed curve denoted “ $Q^{10} \cup p^\tau$ ”) at $\tau = 3$, and at first glance suggests the infelicity of Z-normalization. Closer inspection reveals that while Z-normalization usefully removes inter-speaker variability, it also brings values close to the speaker’s mean close to zero, which makes them — from the model’s point of view — indistinguishable from unvoiced frames. Exposing the model to both \mathbf{V} and \mathbf{Z} corrects this, and yields a cross entropy rate of 0.268366 at $\tau = 4$, as can be seen in Figure 1 for the curve denoted $Q^{10} \cup v^\tau \cup z^\tau$. This is lower than the rate achieved by the $Q^{10} \cup v^\tau \cup p^\tau$ curve by 0.0019380 bits/frame, almost 4 times more than the reduction observed when including \mathbf{V} with \mathbf{P} .

Pitch is claimed to be perceived on a logarithmic scale; to explore whether log-pitch outperforms pitch on the speech activity prediction task, $\mathbf{L} \equiv \log_2 \mathbf{P}$ was formed. Its elements $\mathbf{l}_k[n]$ are equal to $\log_2 \mathbf{p}_k[n]$ when $\mathbf{p}_k[n]$ is non-NaN, and NaN otherwise. Z-normalizing \mathbf{L} instead of \mathbf{P} yields a new representation \mathbf{Y} , whose elements $\mathbf{y}_k[n]$ are equal to $(\mathbf{l}_k[n] - \mu_L) / \sigma_L$ if $\mathbf{l}_k[n]$ is non-NaN, and NaN otherwise. Denoted by the curve “ $Q^{10} \cup v^\tau \cup y^\tau$ ” in Figure 1, this representation yields a DEVSET cross entropy rate minimum of 0.268441 at $\tau = 4$. This is actually higher than the DEVSET minimum of the “ $Q^{10} \cup v^\tau \cup z^\tau$ ” curve, but it is lower for all values $\tau \neq 4$, and also

smoother over the entire $\tau \in [1, 10]$ range.

The last experiment of this subsection builds on the logarithmic version, including voicing and z-normalized log-pitch not just for the target participant but also for their interlocutor. This is denoted in Figure 1 by “ $Q^{10} \cup V^\tau \cup Y^\tau$ ”, and its minimum is reached at $\tau = 4$ with a value of 0.267864 bits/frame. It can be tentatively concluded that model sensitivity to the non-target participant’s recent pitch history reduces average surprise, by the small amount of 0.000577 bits/frame.

4.2 History Duration

Experiments in the previous subsection show that recent pitch appears to be correlated with incipient speech activity, and that a predictor exposed to 10 frames of most-recent speech activity should also be exposed to at least 4 most-recent frames of voicing (\mathbf{V}_1^4) and Z-normalized log-pitch (\mathbf{Y}_1^4). Although it cannot be concluded that this particular representation is optimal, it is the most optimal representation from amongst those investigated for the Switchboard corpus. The experiments shown in Figure 2 aim to establish whether this is true even when much longer histories of speech activity are considered; (Laskowski and Shriberg, 2012) had shown that speech activity histories as long as 8 s (80 100-ms frames, compressed quasi-logarithmically) continue to improve predictions.

Figure 2 depicts the same “ Q^τ ” curve shown in Figure 1, but extends this to $\tau = 20$ 100-




Figure 2: Cross entropy rate, along the y -axis in bits per 100-ms frame, for voicing (V) and speaker-dependent Z-normalized log-pitch (Y) on top of either 10 or 20 frames of speech activity from both participants (shown for reference with enlarged markers on the curve for Q alone), as a function of the duration of the pitch history, along the x -axis in number of 100-ms frames. Rates are shown from left to right for TRAINSET, DEVSET, and TESTSET. Lines connecting points are drawn for the purposes of visualization.

ms frames of speech activity history. It can be seen, for both TRAINSET and DEVSET (as well as TESTSET), that the nearly-linear decrease in cross entropy rate as τ increases continues, albeit less steeply. Also shown in the figure is the same curve as “ $Q^{10} \cup V^\tau \cup Y^\tau$ ”, for which the DEVSET minimum can be found at $\tau = 4$. What is new in the figure is the curve denoted as “ $Q^{20} \cup V^\tau \cup Y^\tau$ ”, which depicts the impact of pitch when the speech activity history is 2 seconds rather than 1 second long. As can be seen, this third curve exhibits its DEVSET minimum also at $\tau = 4$. A system trained on $Q_1^{10} \cup V_1^4 \cup Y_1^4$ reduces the cross entropy rate of a system trained on Q_1^{10} alone by $0.274515 - 0.267864 = 0.0066510$ bits/frame; one that is trained on $Q_1^{20} \cup V_1^4 \cup Y_1^4$ exhibits a reduction over a system trained on Q_1^{20} alone by $0.272448 - 0.265396 = 0.0070520$ bits/frame. This is not only a larger reduction in absolute terms, it appears even larger relative to the speech-only baseline. It suggests that the usefulness of the most recent 400 ms of pitch grows as the duration of speech activity history increases.

4.3 Model Complexity and Training

A final suite of experiments was conducted in order to shed light on potential under-training or over-fitting of the model, given the fixed size of TRAINSET. The representation identified at the end of Subsection 4.1 was used, namely $Q_1^{10} \cup$

$V_1^4 \cup Y_1^4$; there, the model consisted of 8 units in its hidden layer and its training consisted of 1000 iterations of SCG descent. Figure 3 compares cross-entropy rates when the number of training iterations and the number of hidden units are varied in $\{1000, 2000, 3000, 4000\}$ and $\{8, 16, 32, 64\}$, respectively. Note that these numbers of hidden units correspond to 305, 609, 1217, and 2433 free parameters, given an input representation dimensionality of 36.

As can be seen in the figure, extending the training regimen to 2000 iterations is clearly beneficial; extending it further to 3000 iterations yields only negligibly lower DEVSET cross entropy rates. Increasing the model complexity from 8 to 64 hidden units is also beneficial, but on DEVSET the improvement from 32 to 64 units is much smaller than on TRAINSET, indicating not-yet overfitting but getting close. The DEVSET cross entropy rate for 64 units and 4000 iterations is already higher than that for 64 units and 3000 iterations. Note that there is no evidence that more than 400 ms of pitch might benefit any of these larger systems.

5 Discussion

5.1 Generalization

The models presented in this article have all been trained using TRAINSET alone; model selection has been conducted using cross entropy rate min-




Figure 3: Cross entropy rate, along the y -axis in bits per 100-ms frame, for four models differing in the number H of hidden units and using 4 100-ms frames of voicing (V) and speaker-dependent Z-normalized log-pitch (Y) on top of 10 100-ms frames of speech activity from both participants, as a function of the number of iterations of SCG training, along the x -axis in thousands. Rates are shown from left to right for TRAINSET, DEVSET, and TESTSET. Lines connecting points are drawn for the purposes of visualization.

Feature Set	H	I	\mathcal{X}
Q_1	8	1000	0.285379
Q_1^{10}	8	1000	0.275052
$Q_1^{10} \cup V_{tar,1}^{\tau=8}$	8	1000	0.272502
$Q_1^{10} \cup P_{tar,1}^{\tau=3}$	8	1000	0.271450
$Q_1^{10} \cup V_{tar,1}^{\tau=9} \cup P_{tar,1}^{\tau=9}$	8	1000	0.271006
$Q_1^{10} \cup V_{tar,1}^{\tau=9} \cup Z_{tar,1}^{\tau=4}$	8	1000	0.269953
$Q_1^{10} \cup V_{tar,1}^{\tau=9} \cup Y_{tar,1}^{\tau=4}$	8	1000	0.269690
$Q_1^{10} \cup V_1^{\tau=9} \cup Y_1^{\tau=4}$	8	1000	0.269568
$Q_1^{10} \cup V_1^{\tau=9} \cup Y_1^{\tau=4}$	64	1000	0.267630
$Q_1^{10} \cup V_1^{\tau=9} \cup Y_1^{\tau=4}$	64	3000	0.267358

Table 1: Cross entropy rates \mathcal{X} in bits per 100-ms frame, obtained for TESTSET using several representations of pitch, numbers H of hidden units, and numbers I of training iterations. All models trained on TRAINSET, and model selection (over τ , H , and/or I as applicable) performed using DEVSET.

imization on DEVSET. TESTSET has been left untouched, and therefore presents a suitable candidate set for characterizing how the proposed framework generalizes to *completely* unseen data. Table 1 summarizes these achievements, from the right-hand-side of Figures 1 and 2.

As can be seen, the absolute reduction in cross entropy rate due to the inclusion of pitch information (in the form of voicing and Z-normalized log-pitch) is $0.275052 - 0.267358 = 0.0076940$

bits/frame. This magnitude represents approximately 75% of the reduction observed when pitch information is excluded and the speech activity context is increased from 1 frame to 10 frames ($0.285379 - 0.275052 = 0.010327$ bits/frame, ie. rows 1 and 2 in the table). All trends observed for TESTSET in Figures 1, 2, and 3 are nearly identical to those observed for DEVSET.

5.2 Normalization

That the prediction of speech activity can successfully make use of approximately 8 s of most-recent speech activity history (Laskowski and Shriberg, 2012), but of only 400 ms of most-recent pitch history, is surprising and somewhat deflating. However, it is important to note that the optimal representation of pitch was determined to involve Z-normalization, for which the conversation-side mean and standard deviation were assumed to be known a priori. In reality, these statistics would need to be accumulated from the start of each conversation, up to and including the $(n - 1)$ th frame. It is also possible that estimation of these statistics should favor the recent past, yielding local Z-normalization statistics which themselves evolve over time. This is currently under investigation.

5.3 Reproducibility

The experiments presented in this article number just shy of 150; each experiment took approx-

imately 6 hours to run on a hyper-threaded 6-core Intel Xeon E5645 2.40GHz machine, running Debian Linux 3.16. The complete experiment suite, including all source and intermediate Switchboard Corpus data, are available at www.cs.cmu.edu/~kornel/software/stt.html.

5.4 Potential Impact

For Switchboard conversations, the proposed framework has demonstrated that attentiveness to the pitch trajectories of both conversation sides reduces the average surprise of incipient side-attributed speech activity. It appears that it suffices for the considered pitch trajectories to be quite short (400 ms). The Switchboard corpus thereby provides sufficient proof that the proposed framework is capable of yielding findings such as these, in cases in which only the actual speech activity is available and for which pitch can be automatically measured. The framework is agnostic to the much more contentious attempts to define and annotate what a turn is, and not reliant on additional turn-landmark or pitch-trajectory annotation.

The direct impact of this work is that it enables the automated analysis — with regard to the role of pitch in turn-taking — of large corpora which would otherwise be intractable to analyze in their entirety. Due to its quantitative nature, the framework enables direct comparisons between corpora which differ in arguably important ways, such as language, dialect, or domain.

Furthermore, an indirect impact of the findings of which the proposed framework is capable is that such findings may inform automated speech processing systems operating under specific language, dialect, or domain conditions, for example mixed-initiative dialog systems. Knowledge of how such conditions affect the interplay between pitch and turn-taking would enhance the naturalness and flexibility of those systems.

6 Conclusions

Pitch has long been held as an important signalling channel when planning and deploying speech in conversation, and myriad studies have been undertaken to determine the extent to which it actually plays this role. Unfortunately, these studies have required considerable human investment in data preparation and analysis, and have therefore often been limited to a handful of specific conversational

contexts. This has made it difficult to compare and contrast, in a quantitative way, the role played by pitch in turn-taking as a function of language, dialect, domain, channel, other-party familiarity, etc.

The framework proposed in this article addresses these limitations, by enabling a nearly-automatic quantitative characterization of the role of pitch throughout an entire conversation, requiring only the raw signal and speech activity references. Although the latter may require prior manual transcription of the lexical content (followed by forced alignment), this is far easier than manually annotating turn landmarks or pitch trajectories, and is often already available for a corpus under study. The framework is adaptable to the role-in-turn-taking analysis of any feature which can be measured from the raw signal.

This article has evaluated the proposed framework by answering three specific questions regarding the role of pitch in turn-taking, in the Switchboard corpus. First, the presented evidence suggests that pitch can be leveraged to reduce the average surprise of incipient speech. Its inclusion, on top of a conditioning context containing 1 second of speech activity from both dialogue parties, yields a cross entropy reduction of 0.014200 bits per 100 ms; this is approximately half as much as is gained by including the non-target participant's 1-second of speech activity, over just the target participant's, in the first place. Second, the optimal representation of pitch appears to be Z -normalized log-pitch, together with the binary indicator variable of voicing; at least in part, the role of the latter is to differentiate between unvoiced frames and voiced mean-log-pitch frames. Finally, experiments indicate that the dynamic pitch trajectory information which is useful for speech activity prediction is limited to the most recent 400 ms; pitch trajectory information less recent than that is necessary only to provide static Z -normalization statistics. Furthermore, the reduction in average surprise appears to be a function of the duration of the considered speech activity history; the longer the speech activity history, the more valuable do those most recent 400 ms of pitch seem to be.

Acknowledgments

This work was funded in part by the Stiftelsen Marcus och Amalia Wallenbergs Minnesfond project MAW 2017.0034, *Hidden events in turn-taking*.

References

- C. Bishop. 1995. *Neural Networks for Pattern Recognition*. Oxford University Press, New York NY, USA.
- S. Bögels and F. Torreira. 2015. Listeners use intonational phrase boundaries to project turn ends in spoken interaction. *Journal of Phonetics*, 52:46–57.
- J. Caspers. 2003. Local speech melody as a limiting factor in the turn-taking system in Dutch. *Journal of Phonetics*, 31(2):251–276.
- N. Deshmukh, A. Ganapathiraju, A. Gleeson, J. Hamaker, and J. Picone. 1998. Resegmentation of SWITCHBOARD. In *Proceedings of the 5th International Conference on Spoken Language Processing (ICSLP1998)*, pages unnumbered, paper 0685, Sydney, Australia.
- S. Duncan. 1972. Some signals and rules for taking speaking turns in conversations. *Journal of Personality and Social Psychology*, 23(2):283–292.
- J. Edlund and M. Heldner. 2005. Exploring prosody in interaction control. *Phonetica*, 62(2-4):215–226.
- C. Ford and S. Thompson. 1996. *Interactional units in conversation: Syntactic, intonational, and pragmatic resources for the management of turns*, chapter 3. Cambridge University Press, Cambridge MA, USA.
- J. Godfrey and E. Hollimann. 1997. *Switchboard-1 Release 2*. Catalog Number LDC97S62, Linguistic Data Consortium, Philadelphia PA, USA.
- A. Gravano and J. Hirschberg. 2011. Turn-taking cues in task-oriented dialogue. *Computer Speech and Language*, 25(3):601–634.
- M. Heldner, J. Edlund, K. Laskowski, and A. Pelcé. 2009. *Prosodic features in the vicinity of silences and overlaps*, pages 95–105. Peter Lang, Frankfurt am Main, Germany.
- M. Heldner and M. Włodarczak. 2015. *Pitch slope and end point as turn-taking cues in Swedish*, pages 10–15. Glasgow, Scotland.
- A. Hjalmarsson. 2011. The additive effect of turn-taking cues in human and synthetic voice. *Speech Communication*, 53(1):23–35.
- G. Jefferson. 1984. *Transcript notation*, pages ix–xvi. Cambridge University Press, Cambridge MA, USA.
- J. Kane, I. Yanushevskaya, C. de Looze, B. Vaughan, and A. Ní Chasaide. 2014. *Analysing the prosodic characteristics of speech-chunks preceding silences in task-based interactions*, pages 333–337. Singapore.
- H. Koiso, Y. Horiuchi, S. Tutiya, A. Ichikawa, and Y. Den. 1998. An analysis of turn-taking and backchannels based on prosodic and syntactic features in Japanese map task dialogs. *Language and Speech*, 41(3–4):295–321.
- K. Laskowski. 2012. Exploiting loudness dynamics in stochastic model of turn-taking. In *Proceedings of the 4th IEEE Workshop on Spoken Language Technology (SLT2012)*, pages 79–84, Miami FL, USA.
- K. Laskowski. 2015. Auto-imputing radial basis functions for neural network turn-taking models. In *Proceedings of the 15th Annual Conference of the International Speech Communications Association (INTERSPEECH2015)*, pages 1820–1824, Dresden, Germany.
- K. Laskowski, M. Heldner, and J. Edlund. 2009. *Exploring the prosody of floor mechanisms in English using the fundamental frequency variation spectrum*, pages 2539–2543. Glasgow, Scotland.
- K. Laskowski and E. Shriberg. 2012. Corpus-independent history compression for stochastic turn-taking models. In *Proceedings of the 37th IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP2012)*, pages 4937–4940, Kyoto, Japan.
- J. Local, J. Kelly, and W. Wells. 1986. Towards a phonology for conversation: Turn-taking in Tyne-side English. *Journal of Linguistics*, 22(2):411–437.
- J. Local and G. Walker. 2012. How phonetic features project more talk. *Journal of the International Phonetic Association*, 42(3):255–280.
- M. Møller. 1993. A scaled conjugate gradient algorithm for fast supervised learning. *Neural Networks*, 6(4):525–533.
- M. Selting. 1996. On the interplay of syntax and prosody in the constitution of turn-constructive units and turns in conversation. *Pragmatics*, 6:357–388.
- K. Sjölander. 2001. Snack sound toolkit 2.2.10. <http://www.speech.kth.se/snack/>.
- G. Walker. 2017. Pitch and the projection of more talk. *Research on Language and Social Interaction*, 50(2):206–225.
- A. Wennerstrom and A. F. Siegel. 2003. Keeping the floor in multiparty conversations: Intonation, syntax, and pause. *Discourse Processes*, 36(2):77–107.
- I. Yanushevskaya, J. Kane, C. de Looze, and A. Ní Chasaide. 2014. pages 959–963. Dublin, Ireland.
- M. Zellers. 2013. *Pitch and lengthening as cues to turn transition in Swedish*, pages 248–252. Lyon, France.
- M. Zellers. 2017. Prosodic variation and segmental reduction and their roles in cuing turn transition in Swedish. *Language and Speech*, 60(3):454–478.

A Large-Scale User Study of an Alexa Prize Chatbot: Effect of TTS Dynamism on Perceived Quality of Social Dialog

Michelle Cohn¹, Chun-Yen Chen², Zhou Yu²

¹Department of Linguistics, ²Department of Computer Science

University of California, Davis

{mdcohn, abtchen, joyu}@ucdavis.edu

Abstract

This study tests the effect of cognitive-emotional expression in an Alexa text-to-speech (TTS) voice on users' experience with a social dialog system. We systematically introduced emotionally expressive interjections (e.g., "Wow!") and filler words (e.g., "um", "mhm") in an Amazon Alexa Prize socialbot, Gunrock. We tested whether these TTS manipulations improved users' ratings of their conversation across thousands of real user interactions ($n=5,527$). Results showed that interjections and fillers each improved users' holistic ratings, an improvement that further increased if the system used both manipulations. A separate perception experiment corroborated the findings from the user study, with improved social ratings for conversations including interjections; however, no positive effect was observed for fillers, suggesting that the role of the rater in the conversation—as active participant or external listener—is an important factor in assessing social dialogs.

1 Introduction

Dialog systems, despite recent improvements, still face a fundamental issue of how to convey interest and emotion via text to speech (TTS) synthesis. Many TTS voices have been described as "robotic" or "monotonous" by human listeners (Baker, 2015), an issue further exacerbated for generation of longer utterances (Németh et al., 2007). This is particularly relevant for non-task-oriented dialog systems, such as those that aim to engage users in social chitchat (Akasaki & Kaji, 2017; Liu et al., 2017); for example, Tokuhisa & Terashima (2009) found that affective (i.e., emotion conveying) productions relate to

perceptions of speaker enthusiasm in non-task-oriented human-human conversation. In another study, adjustment of the prosodic features of computer TTS affects listeners' perceptions of the system's type of clarification request (Skantze et al., 2006), signaling its "cognitive state". Still, the ability to design a computer or robot system to convey cognitive-emotional expressiveness remains an area of rich study in the field of Affective Computing (AC) (cf. Tao & Tan, 2005). While prior approaches to model human-like expressiveness in various systems have involved manipulation of the overall TTS prosody, including pitch, rate, and volume (e.g., Gálvez et al., 2017; Henning & Chellali, 2012; Montero et al., 1998; Mustafa et al., 2010; Nass & Lee, 2001; Schröder, 2007), the present paper tests whether adding minimal and discrete emotional-cognitive expressions in a TTS voice impacts user experience with a social dialog system. More specifically, we examine whether a full "overhaul" of prosody is necessary to meaningfully improve a dialog system, or whether we can inject units of cognitive-emotional expression in carefully specified locations to produce a similar effect.

Yet, our understanding of what types of TTS modifications will result in believable and sincere expressions of emotion and cognitive states in a dialog system remains an open question; there have been mixed findings as to whether "human-like" TTS adjustments, such as adding filler words, result in improved user metrics (e.g., Syrdal et al., 2010; Pfeifer & Bickmore, 2009).

Critically, the vast majority of human-computer dialog studies have been run on a limited number of participants and conversations (e.g., $n=96$ in in Brave et al., 2005) and in a lab setting where users are recruited to interact with the systems (e.g., Brave et al., 2005; Cowan et al., 2015; Qvarfordt et al., 2005; Yu et al.,

2016); that is, users may not be interacting with real intents. For one, the presence of an experimenter could impact the way users interact with the system (cf. Orne, 1962). This is also true for dialog systems; users may be less comfortable to engage in more naturalistic conversation, or may be more willing to accept errors or incongruencies by a computer system while in the lab. Additionally, having fewer observations, as well as a participant pool largely consisting of college age students (e.g., Cowan et al., 2015) may impact researchers' ability to generalize findings to other user demographic groups (cf. Henrich & Heine, 2010).

In this paper, we describe an experiment where we systematically manipulated the Amazon Alexa TTS generation in Gunrock, the 2018 Alexa Prize winner socialbot (Chen et al., 2018). Our participants included over 5,000 real users who engaged with the system from their own homes and devices. We targeted two types of TTS manipulations: interjections (e.g., "Awesome!") and filler words. We selected these two elements as they are ways humans communicate their cognitive-emotional states, but vary in their intensity: while interjections express enthusiasm and strong emotion, filler words communicate the speaker's cognitive states (e.g., "Um... let me think") in a more tempered fashion. Both interjections and fillers have also been proposed to serve as socio-affective "glue" between interlocutors, expressing emotional and cognitive states that serve to strengthen relational bonds between humans and computers (Auberge et al., 2013; Sasa & Auberge, 2014; 2017).

In addition to its scope, this study is novel in several regards. First, no prior work, to our knowledge, has explored how individuals respond to emotion generated by a voice-activated digital assistant (e.g., Amazon's Alexa, Apple's Siri); users may have a more personal connection with and may even show greater personification of these increasingly prevalent household devices (Lopatovska, & Williams, 2018). Additionally, this paper introduces a methodology for designing and inserting interjections and filler words, both in terms of their context as well as their acoustic adjustments using Speech Synthesis Markup Language (SSML). Furthermore, no prior experiments have parametrically tested the presence of these two

elements in controlled studies; doing so allows us to test whether there is a cumulative effect of these cognitive-emotional insertions. Finally, conducting an experiment directly through the Alexa system is an innovative approach that builds on past work that has largely relied on naturalness ratings of synthetic voices with no interactive component for the rater themselves (e.g., Marge et al., 2010; Gálvez et al., 2017; Hennig & Chellali, 2012; Schmitz et al., 2007).

This study can serve as a test to the 'Computers are Social Actors' theoretical framework (CASA: Nass et al., 1994; Nass & Moon, 2000) that proposes that humans apply social norms from human-human interaction to *computers* when they detect a cue of humanity in the system. One empirical question for the CASA framework is what cues can trigger computer personification and to what extent this personification graded; that is, do we see cumulative effects of introducing multiple human-like features in a dialog system, or do listeners display a more categorical response to human-likeness? In particular, we ask whether individuals' ratings of social dialog quality vary according to the type and combination of addition for interjections and filler words.

In the following section, we will review the literature for related work on cognitive-emotional expression via interjections and filler words in human-human and human-computer interaction (HCI). Then, we will introduce our overall chatbot dialog system design and our interjection/filler insertion methodology in Section 3, our user study experiment in Section 4, and a perception experiment in Section 5.

2 Related Work

2.1 Limited Prior Work on Interjections and Exclamations in HCI

Despite the prevalence of interjections in human speech patterns, few groups have explored inserting interjections in TTS systems. In human speech, interjections constitute words or phrases that can display emotion (e.g., emotive interjections such as "Yuck!"; cf. Wierzbicka, 1999) or reveal the speaker's "information state" (e.g., "Aha!"). Some interjections are based on existing words (e.g., "Neat!"), while others are based on non-lexical vocal productions (e.g., "Ooh!"; cf. Yang, 2010). Interjections can also

signal that the information is newsworthy (e.g., “Really?” in Pammi, 2012). Still, the addition of interjections in TTS voices remains a largely understudied area, while much greater attention has been given to overall prosodic adjustments over the scope of a phrase or utterance (e.g., pitch, duration, etc.) (e.g., Németh et al., 2007) or the introduction of non-linguistic affective bursts in robots (e.g., beeps, buzzes in Read & Belpaeme, 2012). While not introducing interjections per se, but rather modeling new TTS productions based on positive or negative interjections (e.g., “Great!” vs. “Oh dear!”), Syrdal and colleagues (2010) found that speech trained on positive exclamations resulted in higher listener ratings in a 7-utterance simulated dialog; they observed no such effect for TTS adjustments for negative exclamations (e.g., “Oh dear!”, “Oops!”). One novel line of research we explore in the present study is whether the presence of an interjection – and the degree of prosodic dynamism in the interjection, such as exaggerating the pitch contour and increasing duration – contributes to a user’s perception of the system as being more cognitive-emotionally expressive.

2.2 Mixed Results for Fillers in HCI

Another element signaling cognitive-emotional expression in human conversations is filler words. In certain instances, filler words, or filled pauses (e.g., “um”), can be considered to be a type of disfluency or hesitation in a speaker’s production (Clark & Tree, 2002), demonstrating more time for the speaker to “collect” their thoughts (cf. Brennan & Williams, 1995). At the same time, filler words can signal information about the speaker’s cognitive state; for example, longer filler words have been shown to signal greater uncertainty or degree of thought on the conversational subject, while the pitch contour on the filler word communicates the speaker’s level of understanding (Ward, 2004). In some studies, introduction of filler words in dialog systems has a facilitatory effect on perceived naturalness and expressiveness of the voice (Gallé, et al., 2017; Goble & Edwards, 2018; Marge et al., 2010; Wigdor et al., 2016). For instance, a user’s “sensation of engagement” in a conversation with a robot improves with the addition of filler words (Gallé, et al., 2017). Filler words additionally have been shown to impact perceived likeability and engagement with a computer, even for individuals not

directly talking to the computer/robot; independent raters gave higher naturalness ratings for “overheard” human-computer conversations when the computer voice included filler words (e.g., using the Talkie dialog system in Marge et al., 2010).

Yet, at the same time, other studies have reported no effect of introducing filler words (e.g., “Hmmm”, “uh huh” in Syrdal et al., 2010), or a negative effect for some listeners (e.g., Pfeifer & Bickmore, 2009). This negative response might be expected given their association with as markers of anxiety and unpreparedness for some subjects. However, Christenfeld (1995) additionally observed that listeners’ evaluations varied based on their task: when asked to focus on the speech *style*, subjects reported more negative ratings of the filler “um”, but subjects had no such negative judgments when they were asked to focus on the content. This raises an important question: how might the experimental task impact the way users perceive these more human-like, but in some cases more “marked”, displays of cognitive-emotional expressiveness? Addressing a limitation of prior work having subjects rate stimuli presented in isolation (e.g., Syrdal et al., 2010), our study tests both actual user’s responses as well as external raters in assessing the introduction of fillers.

3 Dialog System Design Amazon Alexa Prize Chatbot

For the past two years, Amazon has launched the Alexa Prize Socialbot Challenge to support universities in building conversational bots to advance human-computer interaction. General public users with an Alexa-enabled device or free Alexa application can access the system and talk to the system about various topics (e.g., music, sports, animals, movies, food, weather, etc.) in a conversational manner. When a user engaged the social mode by saying “Let’s chat”, one of the socialbots in the competition was randomly invoked. After talking to the system, the Alexa Skill system automatically solicited user feedback (“How likely are you to talk to this bot again, on a scale from one to five?”), providing a measure of user engagement.

Competing in the 2018 Alexa Prize competition, our chatbot, Gunrock (Chen et al., 2018), aims to produce engaging and coherent conversations with real human users. During the competition, our bot achieved an average rating

of 3.62 (on a 1-to-5 scale) in over 40,000 conversations; conversations had an average of 18.9 turns, averaging 4.35 minutes in duration. Our bot uses automatic speech recognition and text-to-speech models are provided by Amazon. It has a three-stage natural language understanding pipeline including ASR correction, sentence segmentation, constituency parsing, and dialog act prediction to aid user intent detection. Our system has a hierarchical agenda-based dialog manager that covers different topics, such as movies, music, etc., and a templated-based natural language generation module that allows the system to fill slots with data retrieved from various knowledge sources. Please refer to Chen et al. (2018) for system implementation details.

3.1 Methods of Inserting Interjections (Speechcons)

We designed a framework to introduce 52 distinct interjections pre-recorded by the US English Alexa voice actor. These interjections, known as Speechcons (Amazon, 2018), are “special words and phrases that Alexa pronounces more expressively”. For a listening sample, refer to the Speechcon website (Amazon, 2018). We inserted these interjections using Speech Synthesis Markup Language (SSML) tags in the Alexa Skills Kit. These interjections were longer in duration and showed wider pitch variations and exaggerated pitch contours, relative to their unmodified counterparts (see Figure 1).




Figure 1: Pitch and duration differences for Speechcon and unmodified production of “Cool!” generated in Praat (Boersma & Weenik, 2018).

Of the 52 interjections (see Table 1 for a breakdown), we inserted 39 phrase-initially using a rule-based system, for the following 5 contextual scenarios, defined by conversational template: when the bot wanted to signal interest about the user’s response to encourage the user to elaborate, to resolve an error, to accept a request, to change the topic, and to express

agreement of opinion. In each context, we randomly inserted an interjection appropriate for that context (from the subset of pre-categorized interjections) to increase variation and retain user interest. Note that insertion of interjections did not result in any pauses or other incongruencies in the Alexa TTS generation.

Context	Purpose	# Interjections
1	Signal interest	12
2	Error resolution	14
3	Accept request	4
4	Change topic	4
5	Agree	2
<i>Utterance-specific</i>		13

TOTAL: 52

Table 1: Total number of possible interjections added to defined slots in conversational templates.

Interjections were selected for each context by a native English speaker (Author 1) based on the acoustic production of the interjection and its semantic/pragmatic fit in the utterance. First, we selected positive interjections (e.g., “Wow!”) that could be used to signal interest (Context 1) and negative interjections (e.g., “Darn!”) in error resolution (Context 2); we used the widest variety of interjections for these two contexts as these situations arose most frequently in conversation. We denote the interjection version of words with an exclamation (e.g., “Awesome!”).

- **Context 1: To signal interest about the user’s response and elicit user’s expansion.**

We added 12 interjections phrase-initially to show Alexa’s interest in the user’s answer (after Alexa asks a question and the user provides a response); these interjections included “Awesome!”, “Cool!”, “Fantastic!”, “Super!”, “Wow!”, “Ooh la la!”, “No way!”, “Fancy that!”, “Interesting!”, and more (for a full list, see Appendix A). For example: “[Wow!... | Interesting!... | Ooh la la!...]. Tell me more about it.”

- **Context 2: Error resolution.**

We also introduced 14 interjections in error resolution templates in order to show Alexa’s “feelings” about her misunderstanding. Possible interjections included “Whoops a daisy!”, “Darn”, “Oh brother”. For example: “[Whoops-a-daisy!... | Baa!... | Darn!...] I think you said probably. Can you say that one more time?”

- **Context 3: To accept a request.** We inserted 4 interjections phrase-initially to reflect Alexa’s acceptance of the user’s request (e.g., such as to change topic), including: “Okey dokey!”, “Righto!”, “As you wish!” and “You bet!”. For example: “[*Okey dokey!... | Righto!... | As you wish!...*] Here’s some more info.”
- **Context 4: To change the topic.** We used 4 interjections to transition to a new topic, simulating a scenario where Alexa “just remembered” something she wanted to share with the user. We generated 2 interjection versions of “Ooh!” and “Ah!” to use in this context. For example: “[*Ooh!... | Ah!... | All righty!...*] tell me more about you! What else are you interested in? Do you like [music | movies | animals]?”
- **Context 5: To express agreement of opinion.** We inserted 2 interjections phrase-initially to show Alexa’s emphatic agreement to the user’s opinion: “Yes!” “High Five!”. For example: “[*High Five!... | Yes!...*] We share the same thoughts!”

Overall, our rule-based system resulted to the insertion of interjections in 12-18% of turns in each conversation. We implemented these interjections with a following pause (ranging from 150-300ms), using SSML. Note that 13 unique interjections, of the total 52, were added to very specific utterances (e.g., using “Moo!” with cow jokes) without using this rule-based system (see [Appendix B](#) for stimuli and descriptions). All the interjections were rated on two axes by a native English speaker (see [Appendix A](#) for full word list and classifications; see Table 5 for an example conversation log from in-lab user tests). **Axis 1** is valence: Positive, neutral, or negative. For example, the interjection “Awesome!” was rated as having a positive valence, while “Darn!” was rated as having a more negative valence. **Axis 2** is the interjection emotional orientation: self- or other-oriented (cf. [Brave et al., 2005](#)).

3.2 Methods of Inserting Fillers

We added 9 fillers used in American English ([Barbieri, 2008](#)) in the conversational templates: “um”, “hmm”, “huh”, “ah”, “uh”, “oh”, “ooh”, “uh huh”, “mhm” (see Table 5 for an example conversation log from in-lab user tests). In all cases, we used SSML to add a pause (ranging from 150-200ms) following the filler word and slow the production of the word “so” (80% of

original rate), if it occurred before or after the filler to improve naturalness. We added certain subsets of filler words in three specific contexts: to change topics, when retrieving Alexa’s backstory, and as an acknowledgment to the user’s utterance. Overall, this resulted in fillers added to a total of 7.8-7.9% of total turns.

- **Context 1: To change topic.** We added two fillers, “um” and “uh”, either before or after “so” to introduce a new topic. We additionally reduced the rate of “so” (indicated by underlining in the following examples). For example: “[*Um...sooo, | Sooo, um... | Uh... sooo | Sooo... uh,*] I’ve been meaning to ask you: do you like to play videogames?”
- **Context 2: When retrieving Alexa’s backstory.** We added six fillers (“mhm”, “hmm”, “um”, “uh”, “oh”, and “ooh”) at the beginning of the utterance when the user had asked Alexa a question, simulating that Alexa needed time to consider her own experience and/or opinions. For example: “[*Hmm..., | Uh... | Oh... | Ooh... | Mhmm...*] I love all animals, but I think my favorite is probably the elephant”.
- **Context 3: As an acknowledgment to the user’s answer to Alexa’s question.** We added the fillers to act as feedback response tokens. Specifically, we added “ah”, “oh”, “uh huh”, “mhm”, “huh”, and “ooh” at the beginning of the utterance to show Alexa’s acknowledgment of the content provided by the user (e.g., “*Oh... legos? Interesting choice!*”). Note that while these utterances are often used for backchanneling, where one speaker provides verbal feedback while the other continues to hold the floor (e.g., “uh huh” in [Pammi, 2012](#)), we do not classify them as such they did not occur during the user’s turn. Given the limitations of the text transcripts of the conversations—in the absence of acoustic-phonetic data—we could not implement a real-time backchanneling mechanism

4 Experiment 1: Chatbot User Study

In the current study, we systematically tested the impact of adding interjections and fillers in the Alexa TTS voice in our chatbot ([Chen et al., 2018](#)). We hypothesize that in a social dialog system, adding interjections (e.g., “Awesome!”) and filler words (e.g., “um”) in appropriate locations, with emotional valence consistencies,

will improve overall user ratings. This prediction stems from related work conducted in laboratory settings with other types of interlocutors (e.g., robot in Gallé et al., 2017; Marge et al., 2010), with greater expressiveness of the voice relating to positive ratings by users (e.g. Hennig & Chellali, 2012).

4.1 Experimental Conditions

From November 20, 2018 to December 3, 2018 we conducted an ablation study with four possible conditions, varying according to the presence of interjections and fillers (see Table 2). Condition A was filtered to include interjections (and exclude filler words). Condition B was filtered to include filler words and exclude interjections. Condition C included both interjections and fillers, while Condition D excluded both elements. Condition was randomly invoked for each user. During this timeframe, no other code updates were implemented. A total of 5,527 users participated in the study for a total of 5,582 conversations, with 62,130 conversational turns.

Condition	Interjection	Filler	Users	Conversations
A	✓	-	n=1511	n=1523
B	-	✓	n=1183	n=1196
C	✓	✓	n=1423	n=1443
D	-	-	n=1410	n=1420
Total			n=5527	n=5582

Table 2: Experimental conditions & summary statistics

4.2 Statistical Analysis & Results

We modeled user rating (produced at the end of the interaction on a scale from 1-to-5) with a mixed effects linear regression with the *lme4* R package (Bates et al., 2015), with the fixed effect of Condition (A: Interjection only, B: Filler only, C: Interjection and Filler, or D: Neither) and by-user random intercepts. Effects were contrast coded relative to Condition D (baseline condition).

The linear regression model revealed a main effect of Condition on users’ ratings, with significantly higher ratings for the three conditions with manipulations (A: Interjection, B: Filler, and C: Interjection & Filler) relative to baseline (see Table 3 and Figure 2 below). The highest rating improvement was observed for

Condition C (Interjection & Filler) with an average increase of 0.749.

User Study: Holistic Ratings

	Coef	SE	df	t	p
Intercept	2.88	0.025	8.09E+03	114.37	<0.001 ***
A- Interjection only	0.34	0.019	2.99E+04	17.50	<0.001 ***
B- Filler only	0.34	0.019	2.92E+04	18.02	<0.001 ***
C- Interjection + Filler	0.75	0.019	2.98E+04	38.56	<0.001 ***

Num. observations = 31,065 *REML criterion at convergence: 18039.8*

Table 3: Hierarchical linear regression model output: User ratings based on Condition, relative to the baseline condition (“D”).




Figure 2: Mean user rating by Condition (error bars represent standard error; asterisks depict significance ($p<0.001$) relative to the baseline condition, “D”)

The revealed linear regression model, with Condition C as the reference, tested whether the combined condition (Interjections & Fillers) showed higher ratings relative to the addition of interjections or fillers alone. Results revealed that Condition C indeed showed higher user ratings than Conditions A (Interjections only: $\beta=-0.561$, $t=-26.16$, $p<0.001$) or B (Filler only: $\beta=-0.326$, $t=-15.33$, $p<0.001$).

4.3 Interjections Subset Analysis & Results:

We conducted a more fine-grained analysis on the subset of conversations that included the interjections (i.e., Condition A: Interjection, and Condition C: Interjection and filler). In this section, we test whether valence (positive, neutral, negative), emotion orientation (self- versus other), and interjection function (error resolution, change

topic, signal interest, etc.) differentially affect user ratings. We predict that more positive interjections, interjections that communicate more other-oriented displays of emotion, and interjections that are used to signal interest (relative to other functions, such as changing topic) will show higher user ratings, in line with prior work (e.g., Bono & Ilies, 2006; Brave et al., 2005; Gibbs & Mueller, 1988).

A mixed effects linear regression model tested the interjection classifications on user’s ratings. Fixed effects included Interjection Valence (positive, negative, neutral), Emotion Orientation (self-oriented, other-oriented), and Context (Error resolution, change topic, play, etc). Given the overlap between Emotional Valence and Function (with positive interjections exclusively used to Signal Interest and negative interjections almost always used in Error Resolution, see Appendix A), we tested these two variables in separate models. Random effects included by-user random intercepts.

Model comparisons based on the corrected AIC (Burnham et al., 2011) were conducted with the *MuMIN* R package (Barton, 2017) to test the inclusion of Emotion Valence or Function as main effects, given their colinearity. Model comparisons revealed that the model with the fixed effects of Valence and Emotion Orientation best fit the data ($AIC_c=1689.9$), relative to the model including Function and Emotion Orientation ($AIC_c=1694.78$). The retained model output (see Table 4) revealed a main effect of Emotion Orientation, with “other” oriented emotional displays (e.g., “Wow!”) associated with higher rating than more self-oriented productions (e.g., “ah”). No differences were observed on the basis of interjection Valence.

User Ratings: By Interjection characteristics

	Coef	SE	df	t	p
Intercept	3.25	0.079	451.0	40.96	<0.001 ***
Emotion-Other	0.06	0.022	245.1	2.74	0.007 **
Valence-Positive	0.02	0.032	248.50	0.76	0.448
Valence-Negative	0.03	0.035	243.60	0.85	0.398

REML criterion at convergence: 1670.8

Table 4: Mixed effect linear regression output for interjection classification effects on user ratings.

4.4 Qualitative User Study

As part of the Alexa Prize Competition, we additionally recruited users to interact with the system for feedback and bug testing for earlier versions of the dialog system. In September and October 2018, we recorded the interactions of twenty volunteers (12 undergraduates, 8 graduate students). After talking to the socialbot, subjects were asked about their interaction. Several subjects mentioned that they liked the filler words in Alexa’s speech as it “sounded like she was actually thinking” or “seemed more realistic”. Additionally, we noted that subjects often laughed or smiled when they heard the hyper-expressive interjections while they were part of the conversation (e.g., “Wowza!”).

USR:	Hey Alexa, let’s chat.
SYS:	Hi, this is an Alexa Prize Socialbot. I’m glad to meet you. How are you doing today?
USR:	I’m okay. How are you?
SYS:	I’m good! So, tell me more about yourself! Do you like movies?
USR:	I do like movies.
SYS:	I’m excited to talk about movies. <i>Sooo...umm...</i> (Filler), have you seen any movies lately?
USR:	A Star is Born.
SYS:	<i>Ah...</i> (Filler) I’ve seen A Star is Born too! What would you rate this movie on a scale from 1 to 10?
USR:	A ten.
SYS:	<i>Wow!...</i> (Interjection) A perfect 10! It really is an exceptional movie! Here is something I just learned. Some of the scenes were filmed at the Coachella Festival in 2017...

Table 5: Conversation transcription from an in-lab user test. Interjections and filler words are denoted by *italics* and labeled (original, annotations).

5 Experiment 2: Perception Study

While our user study suggests an improvement on the basis of interjections and fillers, it is possible that other factors played a role in the final ratings (e.g., specific phrasing), as well as the co-occurrence of certain interjections, with particular dialog acts (e.g., Alexa using “Darn!” to resolve errors). To disentangle these factors, we conducted a psycholinguistic experiment using a

Qualtrics survey administered through Amazon’s Mechanical Turk¹.

5.1 Participants, Stimuli, and Procedure

A total of 85 Amazon Mechanical Turk workers (i.e., “Turkers”) participated in the rating task (note that all Turkers had to have an approval rating of 97% or higher and at least 1000 prior HITs). Stimuli consisted of four 3-utterance dialogs between Alexa and a human male talker (a native English speaker, age 29). The conversation topics were based on those discussed in the main social bot (animals and movies), though were novel utterances. The dialogs systematically varied as to whether the expression of emotion in the interjection (if expressed) was self- or other-oriented and had positive or negative valence.

Using the rules for inserting interjections and fillers (see Sections 3.2 and 3.3) and mirroring the Condition structure from Experiment 1, we systematically generated four conditions for each dialog: A) Interjection addition, B) Filler addition, C) Interjection and Filler addition, and D) Baseline. In each of these conditions, we held the human’s response exactly the same, as well as all of the wording (for an example, see Table 6). Using a between-subjects design, we additionally tested whether the conversational context for filler words in the first utterance affects their ratings (e.g., following: “So” versus “Yeah, movies can be really fun....So”).

CONDITION 1A: Interjection	CONDITION 1B: Filler
<p>Alexa: So, I’ve been meaning to ask you. What else are you interested in? Do you like animals?</p> <p>Human: I love animals!</p> <p>Alexa: <i>Awesome!</i> I think my favorite animal is the elephant.</p>	<p>Alexa: <i>Sooo, um...</i> I’ve been meaning to ask you. What else are you interested in? Do you like animals?</p> <p>Human: I love animals!</p> <p>Alexa: <i>Awesome.</i> I think my favorite animal is the elephant.</p>

Table 6: Example dialog (Conditions A and B) excerpt used in the perceptual ratings study. Interjections and fillers are annotated in italics.

In the experiment, subjects heard each utterance (randomly presented) and were asked to rate Alexa on several dimensions using a sliding bar (on a scale of 0-to-100): likeability, naturalness, expressiveness, and engagement

(e.g., “How engaged does Alexa sound in the conversation?”). Two listening comprehension questions were included to ensure that Turkers were attending to the stimuli and task at hand (e.g., “What was Alexa’s favorite animal?” Correct response: An elephant).

5.2 Analysis and Results

Subjects’ ratings for each variable were analyzed with separate linear mixed effects models, with a fixed effect of Condition and by-Subject random intercepts. Results showed a main effect of Condition, where introducing interjections significantly increased ratings of engagement ($\beta=6.1, t=3.1, p<0.01$), naturalness ($\beta=3.7, t=3.5, p<0.001$), expressiveness ($\beta=9.0, t=7.7, p<0.001$), and likeability ($\beta=3.4, t=3.1, p<0.001$) of Alexa. Furthermore, we observed a significant improvement of introducing both interjections and fillers on perceived expressiveness ($\beta=8.1, t=7.0, p<0.001$). When introducing fillers only, we observed a negative effect on ratings of likeability ($\beta=-2.8, t=-2.5, p<0.05$) and engagement ($\beta=-2.4, t=-2.1, p<0.05$) (see Figure 3).




Figure 3: Perceptual Ratings of Alexa for each Condition.

Subset analyses on interjections (Conditions B and C) relative to the baseline were conducted to test for an interaction of Condition*Orientation (self- versus other- oriented emotion) and Condition*Valence (positive, negative, neutral). The models showed significant interactions for both: interjections that were other-oriented ($p<0.001$) and positive in valence ($p<0.001$) showed higher ratings for likeability, engagement,

¹ www.MTurk.com

and expressiveness. The subset analysis testing an interaction between the filler condition (relative to baseline) and Conversational Context revealed no effect on ratings.

6 Discussion

This paper combines a large-scale user study with a targeted perceptual ratings experiment to test the effect of adding hyper-expressive interjections (e.g., “Awesome!”) and filler words (e.g., “um”, “um”) in a 2018 Amazon Alexa Prize chatbot. Overall, our user study provides evidence that introducing these discrete expressions of cognitive-emotional expression improves users’ experience talking to a social dialog system; this was evidenced by a higher holistic rating that they provided at the end of the interaction on a scale from 1-to-5. Using both a large sample size and in-situ experiment of an Amazon Alexa Skill, such that users directly engaged with their own devices, is a novel methodology for assessing TTS expressiveness that extends prior in-lab studies on users recruited to engage with the system (e.g., Brave et al., 2005; Cowan et al., 2015; Qvarfordt et al., 2005; Yu et al., 2016).

The cumulative effect of adding interjections and fillers (e.g., in Condition C) suggests that individuals might respond better to dialog systems that use greater TTS dynamism, or *variation*, in the ways in which cognitive-emotional expressiveness is conveyed. These findings can inform theoretical frameworks of computer personification (Nass, 1994; Nass & Moon, 2000); while in a conversation with the system, users appear to be reading the minimal and discrete “human” cognitive-emotional cues generated by the TTS voice – and these effects are additive. Additionally, our results support the classification of fillers and interjections as “socio-affective glue” in developing rapport in human-computer interaction (cf., Sasa & Auberger, 2014).

The facilitatory effect of interjections in the user study was additionally replicated in our perceptual ratings study: we found higher ratings of naturalness, expressiveness, and engagement when Alexa used interjections (e.g., `<speechcon>"Awesome!"</spcon>`) versus unmodified productions of the same words (e.g., “Awesome.”). At the same time, we find that introducing filler words improves ratings when the user is directly engaging with the socialbot, but independent raters, who are not directly part

of the conversation, give lower ratings for filler words. This suggests that the role of the user in the conversation, as well as the conversational context (as being more socially oriented) may be important considerations in evaluating TTS manipulations to improve cognitive-emotional expressiveness.

Finally, this work has practical applications for other dialog system designers, with the Alexa system (e.g., using Speechcons), but also more broadly. That we see an improvement across thousands of users and unique conversations suggests that inserting interjections and fillers plays a key role in perceptions of social dialog quality. We see the potential to use this expressiveness in other types of interactions, including task-oriented dialog (e.g., in tutoring, counselling sessions, etc.).

7 Conclusion

Overall, we present a methodology for inserting interjections and filler words in a socialbot dialog system and empirical validation of their use in a large-scale user study. In comparison to utterance- or phrase- level prosodic manipulations, these word-level “infusions” of cognitive-emotional expression are easier to implement and appear to improve users’ experience. For one, that we see an improvement in ratings across a large-scale pool of users, each with a unique conversation, suggests that introducing these minimal TTS manipulations in other types of dialog systems may be beneficial. Future work testing the implementation of interjections and/or fillers in task versus non-task-oriented systems can further tease apart their generalizability.

Acknowledgments

We would like to acknowledge the help from Amazon in terms of financial and technical support and Dr. Georgia Zellou for feedback on the project.

References

- Akasaki, S., & Kaji, N. (2017). Chat Detection in an Intelligent Assistant: Combining Task-oriented and Non-task-oriented Spoken Dialogue Systems. *ArXiv:1705.00746 [Cs]*. Retrieved from <http://arxiv.org/abs/1705.00746>.
- Amazon. (2018). Speechcon Reference (Interjections): English (US) | Custom Skills. Retrieved from

- <https://developer.amazon.com/docs/custom-skills/speechcon-reference-interjections-english-us.html>.
- Auberge, V., Sasa, Y., Robert, T., Bonnefond, N., & Meillon, B. (2013). Emoz: a wizard of Oz for emerging the socio-affective glue with a non humanoid companion robot. In *WASSS 2013*. HAL ID: [hal-00953780](https://hal.archives-ouvertes.fr/hal-00953780).
- Baker, F. S. (2015). Emerging realities of text-to-speech software for nonnative-English-speaking community college Students in the freshman year. *Community College Journal of Research and Practice*, 39(5), 423-441. doi.org/10.1080/10668926.2013.835290.
- Barbieri, F. (2008). Patterns of age-based linguistic variation in American English I. *Journal of Sociolinguistics*, 12(1), 58-88. doi.org/10.1111/j.1467-9841.2008.00353.x.
- Barton, K. (2017). MuMIn : multi-model inference. R package (Version 1.7.2). Retrieved from ci.nii.ac.jp/naid/10030918982/.
- Bates, D., Bolker, B., & Walker, S. (2015). *Fitting Linear Mixed-Effects Models Using lme4*. Retrieved from: doi.org/10.18637/jss.v067.i01.
- Boersma, P., & Weenink, D. (2018). Praat: doing phonetics by computer (Version 6.0.37). Retrieved from <http://www.praat.org/>.
- Bono, J. E., & Ilies, R. (2006). Charisma, positive emotions and mood contagion. *The Leadership Quarterly*, 17(4), 317-334. doi.org/10.1016/j.leaqua.2006.04.008.
- Brave, S., Nass, C., & Hutchinson, K. (2005). Computers that care: investigating the effects of orientation of emotion exhibited by an embodied computer agent. *International Journal of Human-Computer Studies*, 62(2), 161-178. dx.doi.org/10.1016/j.ijhcs.2004.11.002.
- Brennan, S. E., & Williams, M. (1995). The feeling of another's knowing: Prosody and filled pauses as cues to listeners about the metacognitive states of speakers. *Journal of memory and language*, 34(3), 383-398. doi.org/10.1006/jmla.1995.1017.
- Burnham, K. P., Anderson, D. R., & Huyvaert, K. P. (2011). AIC model selection and multimodel inference in behavioral ecology: some background, observations, and comparisons. *Behavioral Ecology and Sociobiology*, 65(1), 23-35. doi.org/10.1007/s00265-010-1029-6.
- Chen, C.-Y., Yu, D., Wen, W., Yang, Y. M., Zhang, J., Zhou, M., ... Iyer, S. (2018). Gunrock: Building A Human-Like Social Bot By Leveraging Large Scale Real User Data. *2nd Proceedings of Alexa Prize*. m.media-amazon.com/images/G/01/mobile-apps/dex/alex/alexaprize/assets/pdf/2018/Gunrock.pdf.
- Christenfeld, N. (1995). Does it hurt to say um?. *Journal of Nonverbal Behavior*, 19(3), 171-186. doi.org/10.1007/BF02175503.
- Clark, H. H., & Tree, J. E. F. (2002). Using uh and um in spontaneous speaking. *Cognition*, 84(1), 73-111. [doi.org/10.1016/S0010-0277\(02\)00017-3](https://doi.org/10.1016/S0010-0277(02)00017-3).
- Cowan, B. R., Branigan, H. P., Obregón, M., Bugis, E., & Beale, R. (2015). Voice anthropomorphism, interlocutor modelling and alignment effects on syntactic choices in human-computer dialog. *International Journal of Human-Computer Studies*, 83, 27-42. doi.org/10.1016/j.ijhcs.2015.05.008.
- Gallé, M., Kynev, E., Monet, N., & Legras, C. (2017). Context-aware selection of multi-modal conversational fillers in human-robot dialogs. In *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)* (pp. 317-322). doi.org/10.1109/ROMAN.2017.8172320.
- Gálvez, R. H., Benuš, Š., Gravano, A., & Trnka, M. (2017). Prosodic Facilitation and Interference while Judging on the Veracity of Synthesized Statements. *Proc. Interspeech 2017*, 2331-2335. doi.org/10.21437/Interspeech.2017-453.
- Gibbs Jr, R. W., & Mueller, R. A. (1988). Conversational sequences and preference for indirect speech acts. *Discourse Processes*, 11(1), 101-116. doi.org/10.1080/01638538809544693.
- Goble, H., & Edwards, C. (2018). A robot that communicates with vocal fillers has... Uhhh... greater social presence. *Communication Research Reports*, 35(3), 256-260. doi.org/10.1080/08824096.2018.1447454.
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). Most people are not WEIRD. *Nature*, 466(7302), 29. doi.org/10.1038/466029a.
- Hennig, S., & Chellali, R. (2012). Expressive synthetic voices: Considerations for human robot interaction. In *RO-MAN, 2012 IEEE* (pp. 589-595). IEEE. doi.org/10.1109/ROMAN.2012.6343815.
- Liu, H., Lin, T., Sun, H., Lin, W., Chang, C.-W., Zhong, T., & Rudnicky, A. (2017). RubyStar: A Non-Task-Oriented Mixture Model Dialog System. *ArXiv:1711.02781 [Cs]*. Retrieved from <http://arxiv.org/abs/1711.02781>.
- Lopatovska, I., & Williams, H. (2018, March). Personification of the Amazon Alexa: BFF or a mindless companion. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval* (pp. 265-268). ACM. doi.org/10.1145/3176349.3176868.
- Marge, M., Miranda, J., Black, A. W., & Rudnicky, A. I. (2010). Towards improving the naturalness of social conversations with dialog systems. In *Proceedings of the 11th Annual Meeting of the Special Interest Group on Discourse and Dialog* (pp. 91-94). Association for Computational Linguistics. <http://aclweb.org/anthology/W10-4318>.
- Montero, J. M., Gutierrez-Arriola, J. M., Palazuelos, S., Enriquez, E., Aguilera, S., & Pardo, J. M.

- (1998). Emotional speech synthesis: From speech database to TTS. In *Fifth International Conference on Spoken Language Processing*. https://www.isca-speech.org/archive/archive_papers/icslp_1998/i98_1037.pdf.
- Mustafa, M. B., Ainon, R. N., Zainuddin, R., Don, Z. M., Knowles, G., & Mokhtar, S. (2010). Prosodic Analysis And Modelling For Malay Emotional Speech Synthesis. *Malaysian Journal of Computer Science*, 23(2), 102-110. <https://ejournal.um.edu.my/index.php/MJCS/article/view/6399>.
- Nass, C., & Lee, K. M. (2001). Does computer-synthesized speech manifest personality? Experimental tests of recognition, similarity-attraction, and consistency-attraction. *Journal of Experimental Psychology: Applied*, 7(3), 171. doi.org/10.1037/1076-898X.7.3.171.
- Nass, C., & Moon, Y. (2000). Machines and mindlessness: Social responses to computers. *Journal of Social Issues*, 56(1), 81–103. doi.org/10.1111/0022-4537.00153.
- Nass, C., Steuer, J., & Tauber, E. R. (1994). Computers are social actors. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 72–78). ACM. doi.org/10.1145/191666.191703.
- Németh, G., Fék, M., & Csapó, T. G. (2007). Increasing prosodic variability of text-to-speech synthesizers. In *Eighth Annual Conference of the International Speech Communication Association*. https://www.isca-speech.org/archive/interspeech_2007/i07_0474.html.
- Orne, M. T. (1962). On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American psychologist*, 17(11), 776. doi.org/10.1037/h0043424.
- Pammi, S. C. (2012). Synthesis of listener vocalizations: towards interactive speech synthesis. PhD thesis, Naturwissenschaftlich-Technische Fakultät I, Universität des Saarlandes, Saarbrücken, Germany. doi.org/10.22028/D291-26277.
- Pfeifer, L. M., & Bickmore, T. (2009, September). Should agents speak like, um, humans? The use of conversational fillers by virtual agents. In *International Workshop on Intelligent Virtual Agents* (pp. 460-466). Springer, Berlin, Heidelberg. http://doi.org/10.1007/978-3-642-04380-2_50.
- Qvarfordt, P., Beymer, D., & Zhai, S. (2005, September). Realtourist—a study of augmenting human-human and human-computer dialog with eye-gaze overlay. In *IFIP Conference on Human-Computer Interaction* (pp. 767-780). Springer, Berlin, Heidelberg. doi.org/10.1007/11555261_61.
- Read, R., & Belpaeme, T. (2012, March). How to use non-linguistic utterances to convey emotion in child-robot interaction. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction* (pp. 219-220). ACM. doi.org/10.1145/2157689.2157764.
- Sasa, Y., & Auberge, V. (2014). Socio-affective interactions between a companion robot and elderly in a Smart Home context: prosody as the main vector of the " socio-affective glue". In *SpeechProsody 2014*. hal.inria.fr/hal-00953723.
- Sasa, Y., & Aubergé, V. (2017, October). SASI: perspectives for a socio-affectively intelligent HRI dialog system. In *1st Workshop on "Behavior, Emotion and Representation: Building Blocks of Interaction"*. hal.inria.fr/hal-01615470/.
- Schmitz, M., Krüger, A., & Schmidt, S. (2007). Modelling personality in voices of talking products through prosodic parameters. In *Proceedings of the 12th international conference on Intelligent user interfaces* (pp. 313–316). ACM. doi.org/10.1145/1216295.1216355.
- Schröder, M. (2007, September). Interpolating expressions in unit selection. In *International Conference on Affective Computing and Intelligent Interaction* (pp. 718-720). Springer, Berlin, Heidelberg. doi.org/10.1007/978-3-540-74889-2_66.
- Skantze, G., House, D., & Edlund, J. (2006). User responses to prosodic variation in fragmentary grounding utterances in dialog. In *Ninth International Conference on Spoken Language Processing*. [isca-speech.org/archive/interspeech_2006/i06_1229.html](https://www.isca-speech.org/archive/interspeech_2006/i06_1229.html).
- Syrdal, A. K., Conkie, A., Kim, Y. J., & Beutnagel, M. C. (2010). Speech acts and dialog TTS. In *Seventh ISCA Workshop on Speech Synthesis*.
- Tao, J., & Tan, T. (2005). Affective computing: A review. In *International Conference on Affective computing and intelligent interaction* (pp. 981–995). Springer. doi.org/10.1007/11573548_125.
- Tokuhisa, R., & Terashima, R. (2009, July). Relationship between utterances and enthusiasm in non-task-oriented conversational dialogue. In *Proceedings of the 7th SIGdial Workshop on Discourse and Dialogue* (pp. 161-167). <https://dl.acm.org/citation.cfm?id=1654628>.
- Ward, N. (2004). Pragmatic functions of prosodic features in non-lexical utterances. In *Speech Prosody 2004, International Conference*.
- Ward, N., & Tsukahara, W. (2000). Prosodic features which cue back-channel responses in English and Japanese. *Journal of Pragmatics*, 32(8), 1177-1207. [doi.org/10.1016/S0378-2166\(99\)00109-5](https://doi.org/10.1016/S0378-2166(99)00109-5).
- Wierzbicka, A. (1999). *Emotions across languages and cultures: Diversity and universals*.

- Cambridge University Press.
doi.org/10.1017/CBO9780511521256.
- Wigdor, N., de Greeff, J., Looije, R., & Neerinx, M. A. (2016, August). How to improve human-robot interaction with Conversational Fillers. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)* (pp. 219-224). IEEE.
doi.org/10.1109/ROMAN.2016.7745134.
- Yang, L. C. (2010). Meaning and Context: Prosodic Variation of Interjections. In *Conversational Speech*. In *Speech Prosody 2010-Fifth International Conference*. https://www.isca-speech.org/archive/sp2010/papers/sp10_380.pdf
- Yu, Z., Nicolich-Henkin, L., Black, A. W., & Rudnicky, A. (2016). A wizard-of-oz study on a non-task-oriented dialog systems that reacts to user engagement. In *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialog* (pp. 55-63). <http://www.aclweb.org/anthology/W16-3608>.

Appendix A. Interjection (Speechcon) Classifications

Function	Valence			Emotional Orientation
	<i>Positive</i>	<i>Neutral</i>	<i>Negative</i>	
<i>Signal Interest</i>	Great! Awesome! Fantastic! Super! Ooh la la! Wowza!	Wow! Cool! Interesting! Fancy that! No way!		<i>Other</i>
	Aha!			<i>Self</i>
<i>Resolve error</i>		Jiminy cricket! Whoops a daisy!	Darn! Shoot! Yikes! Oh boy! Oh dear! Oh brother! Ouch! Tsk tsk!	<i>Other</i>
		Ruh roh!	Baa! Oof! Uh oh!	<i>Self</i>
<i>Accept request</i>		Okey dokey! Righto! As you wish! You bet!		<i>Other</i>
<i>Change topic</i>	Spoiler alert* (only with disclosure)	Ahem! All righty!		<i>Other</i>
		Ooh! Ah!		<i>Self</i>
<i>Express agreement of opinion</i>	High five! Yes!			<i>Other</i>
<i>Joke (phrase-finally)</i>	Just kidding!*		Wah wah* Neener neener!*	<i>Other</i>
			D'oh!*	<i>Self</i>
<i>Joke (specific context)</i>		Woof!^ Moo!^ Meow!^ Kerplop!^ Honk!^		<i>Other</i>
<i>Other context and module-specific interjection</i>	Yum!^ Aww!^			
<i>Response to user after telling a joke</i>	Tee hee!^			

Table A1. Interjections that are only used in very constrained contexts are annotated with an asterisk (*); those that are only used in one, specifically specified sentence are annotated with a carat (^).

Appendix B. Methods of Inserting Sentiment-Specific Interjections

We additionally added 10 interjections in sentiment-specific utterances. These were not interchangeable (unlike Contexts 1-4 described in Section 3.3). We used the interjection, “Spoiler Alert!” to change the topic by leading in to a disclosure by Alexa (see example A below). We introduced 2 interjections as a response to humor, that occurred after a response to a joke. “Tee hee!” and “Woohoo!” (see examples B and C). We implemented “Yum!” specifically in the food module, in response to the user’s favorite food (see example D). Similarly, we added the interjection, “Aww!” as a response to the user disclosing information about their pet in the animal module and “Woof!” and “Meow!” to respond if they indicated they liked dogs or cats, respectively (see examples E-G).

- | |
|--|
| <ul style="list-style-type: none">a) <i>Spoiler alert!</i>... Did you know? I am definitely more of a dog person than a cat person. How about you? Do you like animals?b) <i>Woohoo!</i>... I’m glad you get my awesome humor.c) <i>Tee hee!</i>... I LOL’d at that as well If I could giggle I would.d) <i>Yum!</i>... That sounds really delicious.e) <i>Woof!</i>... I love dogs.f) <i>Meow!</i>... I love cats.g) <i>Aww!</i>... That’s so cute. |
|--|

Table B1: Examples of sentiment-specific interjections (denoted in italics).

We added several interjections (e.g., “Moo!”, “Honk!”, “Woof!”, “Just Kidding!”) at the end of utterances to complement jokes and express playfulness (see examples G-K in Table B-2).

- | |
|---|
| <ul style="list-style-type: none">h) What do you call a cow during an earthquake? ... A milkshake. ... <i>Moo!</i>i) What do you call blueberries playing the guitar?... A jam session. ... <i>Wah wah!</i>j) What did the traffic light say to the car? ...Don't look! I'm about to change... <i>Honk!</i>k) Why wouldn't the shrimp share his treasure?... Because he was a little shellfish... <i>Neener neener!</i>l) Yeah, wouldn't it be (interesting weird) if I could poop? ... <i>Kerplop!</i> |
|---|

Table B2: Examples of sentiment-specific interjections (denoted in italics) added phrase-finally

Additionally, we added “Kerplop!” in our response if a user asked Alexa if she “poops” (a frequent question in the user studies) (see Table B2 above).

Influence of Time and Risk on Response Acceptability in a Simple Spoken Dialogue System

Andisheh Partovi and Ingrid Zukerman

Clayton School of Information Technology, Monash University
Clayton, Victoria 3800, Australia

Andi.Partovi@monash.edu, Ingrid.Zukerman@monash.edu

Abstract

We describe a longitudinal user study conducted in the context of a Spoken Dialogue System for a household robot, where we examined the influence of time displacement and situational risk on users' preferred responses. To this effect, we employed a corpus of spoken requests that asked a robot to fetch or move objects in a room. In the first stage of our study, participants selected among four response types to these requests under two risk conditions: low and high. After some time, the same participants rated several responses to the previous requests — these responses were instantiated from the four response types. Our results show that participants did not rate highly their own response types; moreover, they rated their own response types similarly to different ones. This suggests that, at least in this context, people's preferences at a particular point in time may not reflect their general attitudes, and that various reasonable response types may be equally acceptable. Our study also reveals that situational risk influences the acceptability of some response types.

1 Introduction

Spoken Dialogue Systems (SDSs) must often engage in follow-up interactions to deal with Automatic Speech Recognizer (ASR) errors or elucidate ambiguous or inaccurate requests (which are exacerbated by ASR errors):

- ASR errors, although significantly reduced in recent times,¹ may produce wrong entities or actions, or ungrammatical utterances that cannot be processed by a Spoken Language Understanding (SLU) system (e.g., “the plate inside the microwave” being misheard as “of plating sight the microwave”).²

¹9to5google.com/2017/06/01/google-speech-recognition-humans/.

²All the sample ASR outputs in this paper are real.

- People often express themselves ambiguously or inaccurately (Trafton et al., 2005; Moratz and Tenbrink, 2006; Funakoshi et al., 2012; Zukerman et al., 2015). An ambiguous reference to an object matches several objects well, while an inaccurate reference matches one or more objects partially. For instance, a reference to a “big blue mug” is ambiguous if there is more than one big blue mug, and inaccurate if there are two mugs – one big and red, and one small and blue.

In the last two decades, research in response generation has focused on techniques that generate response policies that optimize dialogue completion, using *Markov Decision Processes (MDPs)*, e.g., (Singh et al., 2002; Lemon, 2011), and *Partially Observable MDPs (POMDPs)*, e.g., (Williams and Young, 2007; Gašić and Young, 2014). Recently, deep-learning algorithms have been used to generate dialogue responses on the basis of request-response pairs, e.g., (Li et al., 2016; Prakash et al., 2016; Serban et al., 2017). Human and simulation-based evaluations of MDP and POMDP systems focus on dialogue completion, while evaluations of deep-learning algorithms focus on individual responses.

In this paper, we draw inspiration from research in Recommender Systems, where Amatriain et al. (2009) and Said and Bellogín (2018) showed that over time, users gave inconsistent ratings to items, leading to the “magic barrier” to prediction accuracy in Recommender Systems (Said and Bellogín, 2018). This prompted us to posit that people may also be inconsistent when assessing responses in a dialogue at different times, which may affect the results of human evaluations.

To investigate this claim, we conducted a longitudinal study in the context of an SDS for a household robot. We first collected a corpus of spoken requests that asked a robot to fetch or move

objects in a room. Our participants were shown the top ASR outputs for these requests (the intention was to replicate the information available to an SDS, without the extra information people can glean from what they hear). They were also told that these requests had to be executed under two risk conditions: *low risk*, where the consequences of performing the wrong action are trivial, and *high risk*, where performing the wrong action could significantly inconvenience the speaker. The participants had to choose among four response types: Do the request without further interaction, CONFIRM the intended object, ask the requester to CHOOSE between a few candidate objects, or ask the requester to REPHRASE all or part of the request. After 1.5-2 years, the same participants were shown the original requests and ASR outputs, and were asked to rate responses generated from their previously selected response types and from other sources, in particular response types selected by one of the authors and by a classifier trained on the author’s chosen response types.

Our findings show that (1) participants down-rated responses sourced from their previously chosen response types; and (2) these responses were liked as much as *different* responses sourced from the response types selected by one of the authors or by the above-mentioned classifier. The first result indicates that, at least in the context of one-shot dialogues with an SDS for a household robot, people’s preferred response types at a particular point in time may not reflect their general attitudes. The second result suggests that, instead of one best response type, several reasonable response types may be acceptable, including those selected by a classifier trained on a non-target but relevant corpus.

We also investigated the influence of situational risk on the acceptability of response types. We found that (3) as expected, under the high-risk condition, the preferred response types were generally more conservative than under the low-risk condition; but (4) surprisingly, participants’ attitudes toward certain response types, e.g., CONFIRM, were not affected by risk.

The rest of this paper is organized as follows. In the next section, we discuss related work. Our experimental setup is described in Section 3. In Section 4, we present our classifier and the features used to train it. The results of our experiment are described in Section 5, and concluding remarks appear in Section 6.

2 Related Work

Decision-theoretic approaches have been the accepted standard for response generation in dialogue systems for some time (Carlson, 1983). These approaches were initially implemented in SDSs as Bayesian reasoning processes that optimize a system’s confidence when making myopic (one-shot) decisions regarding dialogue acts (Paek and Horvitz, 2000; Sugiura et al., 2009), and as *Dynamic Decision Networks* that make decisions about dialogue acts over time (Horvitz et al., 2003; Liao et al., 2006).

MDPs (Singh et al., 2002; Lemon, 2011), *POMDPs* (Williams and Young, 2007; Gašić and Young, 2014), and their extensions *Hidden Information State Model* (Young et al., 2010, 2013) and *Conversational Entity Dialogue Model* (Ultes et al., 2018) were used, often in combination with *Reinforcement Learning (RL)*, to learn policies that optimize dialogue completion on the basis of feedback given by real or simulated users.

Recently, deep learning has been applied to various aspects of SDSs (Wen et al., 2015; Li et al., 2016; Mrkšić et al., 2017; Prakash et al., 2016; Serban et al., 2017; Tseng et al., 2018; Yang et al., 2017). Wen et al. (2015) and Tseng et al. (2018) considered the generation of linguistically varied responses; Li et al. (2016) and Prakash et al. (2016) produced dialogue contributions of chatbots; and Serban et al. (2017) generated helpdesk responses and Twitter follow-up statements. Mrkšić et al. (2017) proposed a dialogue-state tracking framework, and Yang et al. (2017) a mechanism for slot tagging and user-intent and system-action prediction in slot-filling applications. A combination of deep learning and RL has been used in end-to-end dialogue systems that query a knowledge-base, where user utterances are mapped to a clarification question or a knowledge-base query (Williams and Zweig, 2016; Zhao and Eskenazi, 2016; Dhingra et al., 2017). All these systems harness large corpora comprising request-response pairs to learn responses that are assumed to be better than alternative options.

Like evaluations based on simulated users, human evaluations of (PO)MDP/RL systems focus on successful dialogue completion (Singh et al., 2002; Thomson et al., 2008; Young et al., 2010), while human evaluations of deep-learning systems assess individual responses (Wen et al., 2015; Li et al., 2016; Prakash et al., 2016; Serban et al., 2017; Dhingra et al., 2017).

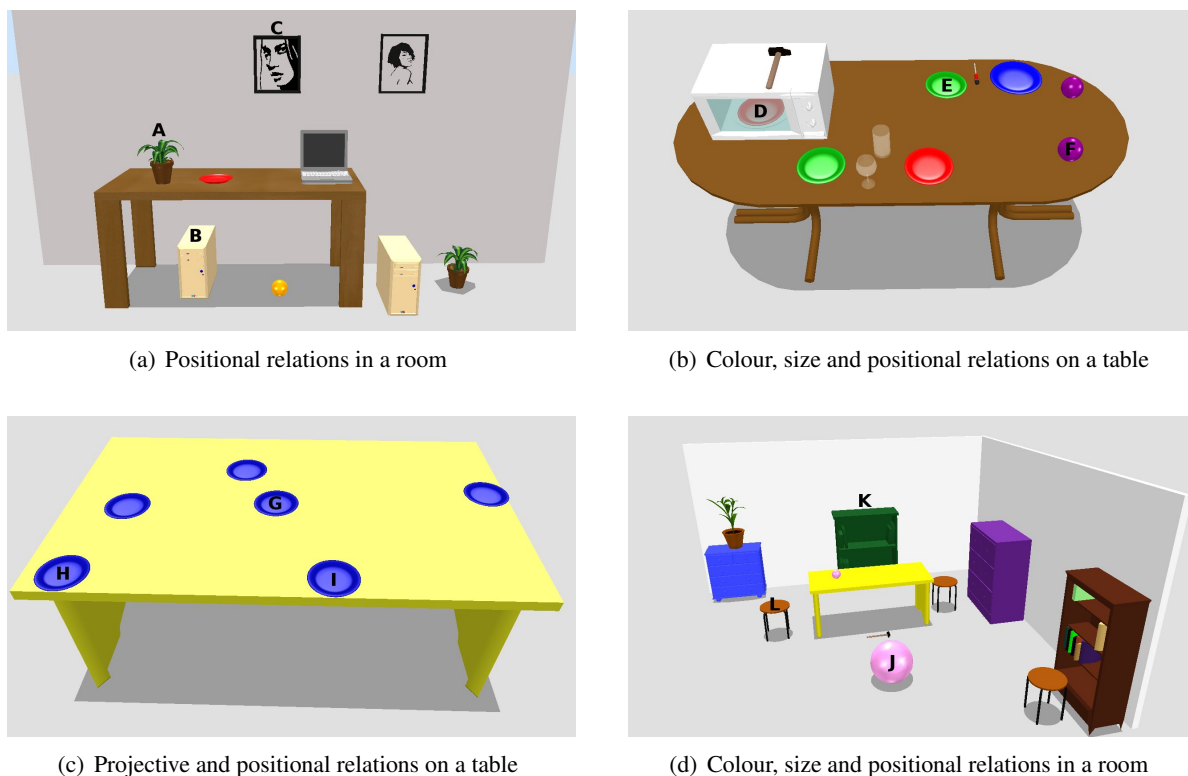


Figure 1: Household scenes used in our study

The findings reported in this paper contribute to (PO)MDP/RL research by determining whether there are factors other than dialogue completion that affect the suitability of responses, and to deep-learning research by ascertaining whether indeed there is a single best response to each request.

The research described in (Jurčiček et al., 2011) and (Liu et al., 2016) shed light on ancillary aspects of human evaluations of system responses. The former compared evaluations by Amazon Mechanical Turk workers with evaluations by participants recruited for a lab experiment; and the latter conducted user studies to determine the validity of word-based evaluation metrics.

This paper also addresses ancillary aspects of human response evaluations, viz the influence of temporal displacement and situational risk on users’ attitudes toward response types, and users’ opinions of response types obtained from different sources (including a classifier trained on a corpus that differs from the target corpus).

3 Experimental Setup

Our experiment comprises two main stages: (1) responding to requests, and (2) rating responses to the same requests.

Creating a corpus of requests

We created a corpus of requests by collecting

a corpus of spoken descriptions, and converting them to requests.

To collect the spoken descriptions, we replicated the experiment described in (Zukerman et al., 2015), but we used the Google ASR, instead of the Microsoft Speech API. In our experiment, the top-ranked outputs produced by this ASR had a 13% word error rate, which resulted in 53% of the descriptions having imperfect top-ranked ASR outputs. In addition, 33% of the descriptions had errors in all top four ASR outputs.

Following the protocol in (Zukerman et al., 2015), 35 participants were asked to describe 12 designated objects (labeled A to L) in four scenes (Figure 1); speakers were allowed to restate the description of an object up to two times. In total, we recorded 478 descriptions such as the following: “the flower on the table” (object A in Figure 1(a)), “the plate inside the microwave” (object D in Figure 1(b)), “the plate at the center of the table” (object G in Figure 1(c)), and “the large pink ball in the middle of the room” (object J in Figure 1(d)). 20% of the descriptions had an unintelligible object in all ASR outputs, e.g., “the *Heartist* under the table”, 17.9% were ambiguous (several objects matched the description), and only 3.8% were inaccurate (no object matched the description perfectly).

We retained 292 descriptions,³ and for each description, we used the top four ASR outputs. The corpus of requests, denoted *RequestCorpus*, was created by prefixing the verb “get” (for small objects) or “move” (for large objects) to each ASR output (which remained unchanged), e.g., “get the flower on the table”. This corpus was divided into sets of at most 12 requests (one request per object, mostly from one speaker).

Demographic and risk-propensity information

We gathered information about the participants’ gender, English nativeness, age, education and risk propensity. For the last item, we showed the participants twelve statements obtained from (Rohmann, 2005): six risk-proneness statements, e.g., “I follow the motto ‘nothing ventured, nothing gained’”, and six risk-aversion statements, e.g., “My decisions are always made carefully and accurately”; (dis)agreement was indicated on a 1-5 Likert scale. The hope was that these information items would assist in predicting participants’ responses.

Stage 1 – Responding to requests

This corpus was collected through an online survey where participants had to indicate how they would respond to potentially misheard requests. Each participant was shown at most 12 requests from *RequestCorpus* (spoken by other people). Each request consisted of four verb-prefixed ASR outputs, and was accompanied by a version of the appropriate image in Figure 1 where the objects were numbered (to enable participants to identify any object as the referent). Each participant was then asked to select one of four response types for each request: DO, CONFIRM, CHOOSE OR REPHRASE. Figure 3 in Appendix A displays a screenshot containing a numbered version of Figure 1(a), four ASR outputs for a request for object #5 (labeled **B** in Figure 1(a)), and the four response types.

Prior to presenting the survey questions, participants were given a training example containing the descriptions shown below in italics:

Do: *Fetch object # [This response is suitable if you are sure which object you should get].*

Here participants were asked to enter the number of the object they would get or move.

³186 descriptions were removed as follows: 20 and 45 descriptions that were not tagged by Stage 1 and Stage 2 participants respectively, 59 descriptions that could not be processed by the SLU system, and 62 descriptions that had more than one prepositional phrase (to simplify the dataset used to train our classifier, Section 4).

CONFIRM: *Ask: Did you mean object #? [This response is suitable if you feel the need to confirm the requested object before taking action].*

Here too participants were asked to enter the number of the object they were confirming.

CHOOSE: *Ask: Which object did you mean? [This response is suitable when you are hesitating between several objects].* In this case, participants were asked to enter the numbers corresponding to their candidate objects.

REPHRASE: *Ask: Please rephrase your request. [This response is suitable when a request is so garbled you can’t understand it].*⁴

These choices were made under two risk conditions: *low risk* – where participants were told that the requested object must be delivered to someone in the same room; and *high risk* – where they were told that the object must be delivered to a remote location (Figure 3). These settings were designed to discriminate between situations where mistakes are fairly inconsequential and situations where mistakes are costly.

40 people took part in this stage of the experiment, but six dropped out after this stage. Half of the remaining participants were male, and 18 were native English speakers. 4 participants were between 18-24 years of age, 16 between 25-34 years of age, 7 between 35-44, and 7 over 45. In terms of education, 5 participants had a secondary education, 16 had a Bachelor, 8 a Masters, and 5 a PhD. To assess the participants’ risk propensity, we subtracted their total risk-aversion score from their total risk-proneness score (the total risk-aversion/proneness score was calculated by adding up the Likert score of the six risk-aversion/proneness statements): 16 participants were risk prone, 8 were risk averse, and 10 were fairly neutral (the difference between the scores was less than 3).

In total, this corpus, denoted *ResponseCorpus*, contains 584 response types (= 292 requests × 2 conditions), which are distributed as shown in Columns 2 and 3 of Table 1.

To determine the influence of speaker diversity on classifier performance (Section 4), we created a second corpus, denoted *AuthorCorpus*, where one of the authors selected response types for all the

⁴As seen in Figure 3, this response type comprised three options: REPHRASE OBJECT, REPHRASE POSITION and REPHRASE ALL. But we merged them into just REPHRASE owing to their low frequency in the dataset (Table 1).

Response type	<i>ResponseCorpus</i>		<i>AuthorCorpus</i>	
	Low risk	High risk	Low risk	High risk
DO	61.3%	45.5%	56.2%	50.3%
CONFIRM	8.9%	17.8%	14.4%	20.2%
CHOOSE	20.2%	23.3%	22.9%	22.9%
REPHRASE	9.6%	13.4%	6.5%	6.5%

Table 1: Response type distribution under high- and low-risk conditions

requests. The distribution of their response types appears in Columns 4 and 5 of Table 1.

Stage 2 – Rating responses to the same requests

After 1.5-2 years, we were able to reach 34 participants from Stage 1, and we built *RatingsCorpus* as follows. Each participant was shown the requests they had seen before (without alerting them to this fact) together with several candidate responses. They were then asked to rate the suitability of each response on a 1-5 Likert scale under the low- and high-risk conditions.

The candidate responses were sourced from the response types chosen by the participant (*ResponseCorpus*) and the author (*AuthorCorpus*) in Stage 1, and the response types returned by a classifier trained on *AuthorCorpus* (Section 4).⁵ In addition, for every DO response from Stage 1, we also presented a CONFIRM response in Stage 2, and vice versa. Clearly, if more than one source had the same response type for a request, this response type was presented only once in Stage 2. Figure 4 in Appendix A displays a screenshot of Stage 2 survey questions regarding the same request as that in Figure 3, presented to the same participant.

Two Stage 2 responses, viz DO and REPHRASE, are direct renditions of the corresponding Stage 1 response types. However, to enable participants to rate CONFIRM and CHOOSE response types, we needed to refer to specific objects. We decided to use images to mimic pointing in CONFIRM responses (e.g., “Do you want this [PICTURE]?”) and in CHOOSE responses with two or three candidate objects (e.g., “There are two things on the table, do you want this [PICTURE 1] or that [PICTURE 2]?”). We restricted the number of CHOOSE responses with images because we deemed it unnatural to

⁵We chose this classifier as it posts high accuracy when trained with limited data, while at the same time, representing a “worst case” for *ResponseCorpus*, as it was trained on a different corpus (the difference between the corpora is statistically significant, χ^2 with p -value < 0.05).

1 Is there an ASR output with all correct words?
2 % of wrong words in the top ASR output
3 % of wrong words in all ASR outputs
4 % of ASR outputs with all correct words

Table 2: Features that reflect the ASR’s confidence

point to more than three things.⁶ In addition, all CHOOSE responses were realized as text only, e.g., “There are two things on the table, which one do you want?”. That is, there were two CHOOSE responses with two or three candidate objects, and one CHOOSE response with more candidate objects. Figure 4 illustrates two CHOOSE responses, a CONFIRM response and a DO response.

4 Using a Classifier to Select Responses

One of the aims of this project is to determine whether we can generate acceptable responses using a classifier trained on a small non-target but relevant corpus. As noted in Section 3, in order to simplify the classifier, we removed descriptions with more than one prepositional phrase. Hence, most descriptions have semantic segments corresponding to an OBJECT, a POSITION SPECIFIER and a LANDMARK (only 22 (7.5%) descriptions have no prepositional phrase, e.g., “the big pink ball”).

4.1 Classification features

To extract features of interest, we assume an SLU system that returns several ranked interpretations, and can represent (a) the ASR’s confidence in the correctness of its candidate outputs, and (b) how well an interpretation (in the context of the room) matches a given description.

We employed the output of the SLU system described in (Zukerman et al., 2015), and for each description, we automatically extracted features that represent the above two types of information. We also included information about situational risk (high or low); and for *ResponseCorpus*, we added the participants’ demographic characteristics gender, English nativeness, age and education, and the difference between their risk-proneness and risk-aversion scores (Section 3).

Features that reflect the ASR’s confidence. These features are shown in Table 2. They reflect the ASR’s “opinion” of the correctness of its output, rather than the ground truth. The last feature is noteworthy because the ASR may have high confidence in a few ASR outputs, e.g., “the flower on

⁶Only 10 (7.9%) CHOOSE responses under both risk conditions had more than three candidate objects.

1 # of interpretations with similar total match score to that of the top-ranked interpretation	($\times 1$)
2 How well the relative position of OBJECT and LANDMARK in an interpretation matches the position specified in the description	($\times 10$)
3 Lexical-match score of the OBJECT, LANDMARK and POSITION SPECIFIER in an interpretation with the corresponding semantic segment in the description	($\times 30$)
4-6 Other match scores of each OBJECT and LANDMARK in an interpretation with the corresponding semantic segment in the description	
4 Colour match score	($\times 20$)
5 Size match score	($\times 20$)
6 # of <i>Unknown</i> modifiers	($\times 20$)

Table 3: Features extracted from top-10 SLU system interpretations

the table” and “the *flour* on the table”, even if only one is intended by the speaker.

Features that represent how well an interpretation matches a description. These features are summarized in Table 3. They are calculated for the top- N interpretations returned by the SLU system, where $N = 10$ (in this system, the correct interpretation is among the top ten in about 90% of the cases). The scores calculated by the SLU system for these features are combined into a total match score for each interpretation, which determines its ranking. For instance, given the description “the brown stool near the table”, two stools in Figure 1(d) have a high total match score, as both are brown and near the table: the stool to the right of the table and stool **L**, which is to the left of the table. However, since the former stool is closer to the table, it has a slightly higher total score, and is ranked first, while stool **L** is ranked second.

The first feature in Table 3 represents the ambiguity of a description through the similarity between the total match score of the top-ranked interpretation and that of subsequent interpretations. We encode this similarity as the ratio between the total score of the i -th interpretation ($i = 1, \dots, N$) and the total score of the top-ranked interpretation. All the interpretations whose ratio is above an empirically-derived threshold are deemed similar to the top-ranked interpretation.

The second feature, computed for each of the top- N interpretations, represents the goodness of the match between the position of the OBJECT in the interpretation (i.e., in the room) and its requested position in the description. For example, both stools in Figure 1(d) are *near* the table, but the position match score of the stool to the right of the table is higher than that of stool **L**.

The rest of Table 3 contains features that represent the quality of the match between individ-

ual elements in an interpretation and their corresponding semantic segments in the given description. Feature #3 represents how well the canonical name of each element in an interpretation matches the corresponding lexical item in the description. For instance, the terms “stool” and “table” respectively match perfectly the terms that designate stool **L** and the yellow table in Figure 1(d). However, if the speaker had said “ottoman”, the lexical match with the canonical term for stool **L** would have been poorer.

Features #4-6 pertain to intrinsic attributes of things, which are normally stated as noun modifiers in a description. They are computed for the OBJECT and LANDMARK of each of the top- N interpretations. Following Zukerman et al. (2015), we have focused on colour and size modifiers, designating other modifiers, e.g., composition or shape, as *Unknown*. Features #4 and #5 respectively reflect the goodness of a match between the color and size of an OBJECT or LANDMARK in an interpretation and the colour and size specifications in the corresponding semantic segment in the given description. For example, a request for a “brown stool” in the context of Figure 1(d) returns a high colour match with stool **L**, while a request for a “blue stool” would return a low colour match. Finally, the match score for Feature #6, which pertains to *Unknowns*, e.g., “the *plastic* stool”, reflects the badness of a match.

4.2 Classifying responses

We considered several classification algorithms to learn response types from the corpora collected in Stage 1 of our experiment (Section 3):⁷ Naïve Bayes, Support Vector Machines, Decision Trees, Random Forest (RF) and Recurrent Neural Nets

⁷We tried over- and under-sampling to deal with the large majority class (DO, Table 1), and applied Principal Components Analysis to reduce the number of features, but these measures did not affect classifier performance.

Response type	<i>ResponseCorpus</i> + <i>Gender & English</i> + <i>RiskPronenessDiff</i>		<i>AuthorCorpus</i>	
	Precision	Recall	Precision	Recall
DO	0.77	0.83	0.945	0.945
CONFIRM	0.44	0.41	0.842	0.842
CHOOSE	0.77	0.72	0.985	0.985
REPHRASE	0.70	0.55	1.00	1.00
Accuracy	0.72		0.94	

Table 4: Per-class and overall classifier performance

(RNNs). RF yielded the best performance for both *ResponseCorpus* and *AuthorCorpus* (RNNs under-performed, as there were not enough data).

Table 4 displays the per-class and overall performance of the RF classifier with 10-fold cross validation for both corpora. As seen in Table 4, RF performed much better for *AuthorCorpus* than for *ResponseCorpus*. This is attributable to the consistency of the 584 ratings provided by one person in *AuthorCorpus*, compared to the variability among participants in *ResponseCorpus* (different participants selected different responses for requests that had the same features).

The demographic features gender and English nativeness and the difference between risk-proneness and risk-aversion scores mitigated the impact of speaker diversity in *ResponseCorpus* (age and education had no effect). In addition, situational risk had some influence on classification results in *ResponseCorpus*. This is consistent with the observation that the vast majority of the differences between the low- and high-risk condition were due to changes from DO to more conservative response types, in particular CONFIRM (represented in Columns 2 and 3 in Table 1). Despite this, most of the misclassifications were also between DO and CONFIRM.

Although the performance of the RF classifier on *ResponseCorpus* is disappointing, this result is tangential to the main thrust of this paper. In Section 5, we examine participants’ attitudes toward responses obtained from the RF classifier trained on *AuthorCorpus* (which is significantly different from *ResponseCorpus*, Section 3).

5 Results

The main objective of our experiment is to determine whether participants’ attitudes toward responses remain consistent over time. That is, how well do participants like their own previous responses? And do they prefer them to other re-

sponses? As mentioned in Section 3, these other responses were sourced from the response types in *AuthorCorpus* and the response types chosen by the RF classifier trained on *AuthorCorpus*.

In addition, we sought to gain insights about the feasibility of using a classifier trained on the responses of one person, and to determine the influence of situational risk on people’s attitudes toward response types.

Hypotheses pertaining to fewer than 200 samples were tested using Wilcoxon matched-pairs signed-rank test, and for more than 200 samples, we used the Normal approximation of this test (Siegel and Castellan, 1988).

How well do people like their previously selected response types? In order to answer this question, we had to address the following issues:

1. In Stage 1, participants selected a response type for each request, while in Stage 2, they rated responses. To compare Stage 1 selections to Stage 2 ratings, we ascribed ratings to the response types selected in Stage 1. In order to account for participants’ rating bias, we assigned to each response type selected by a participant in Stage 1 the highest rating this participant gave to any response in Stage 2 (87% of these highest ratings were 5 – the maximum on the Likert Scale, Section 3).
2. In Stage 2, we offered two options for CHOOSE response types with two or three candidate objects: CHOOSE+pictures and CHOOSE+text (Section 3). For each description, we assigned to a Stage 2 CHOOSE response type the maximum of the ratings of the two options.

We tested the hypothesis that participants’ Stage 1 response types yield highly rated responses in Stage 2 under both risk conditions. The result of this test was that *participants’ Stage 2 ratings of responses sourced from their own Stage 1 response types were significantly lower than the ratings ascribed to these Stage 1 response types under the low- and high-risk conditions* ($p\text{-value} \ll 0.01$).

Figure 2 displays a histogram of the differences between the ratings ascribed to Stage 1 response types and the ratings given to the corresponding responses in Stage 2 under both risk conditions. For example, the leftmost bars indicate that the ratings of 159 response types under the low-risk condition and 123 response types under the high-risk condition did not change between Stage 1 and




Figure 2: Differences between ratings ascribed to Stage 1 response types and ratings of the corresponding Stage 2 responses under low- and high-risk conditions

Stage 2 (the difference is 0). In other words, participants lowered their ratings of 133 response types under the low-risk condition and 169 response types under the high-risk condition. Do (majority class) accounts for 71% of these down-rated response types under the low-risk condition, and 60% under the high-risk condition.

Do users prefer their previously selected response types to other response types? To answer this question, for each risk condition, we collected the participants’ Stage 1 response types that *differ* from those in *AuthorCorpus* for the same request, and their response types that differ from those chosen by the RF classifier trained on *AuthorCorpus*.

Table 5 compares participants’ ratings of responses (R_{S1}) sourced from their Stage 1 response types ($S1$) with their ratings of responses (R_d) sourced from *different* response types (d) selected by the RF classifier for the same requests under the low- and high-risk conditions. In total, 107 response types chosen by the classifier differ from the participants’ selected response types under the low-risk condition, and 126 under the high-risk condition. In 47 of the low-risk cases and 46 of the high-risk cases, the responses sourced from the classifier’s response types received a higher rating than the responses sourced from the participants’ own response types (the results are similar for *AuthorCorpus*). Table 6 illustrates two of these low-risk cases, and two of these high-risk cases. For instance, in the high-risk example pertaining to Figure 1(a), the participant chose REPHRASE in Stage 1, but gave it a rating of 1 in Stage 2, while CONFIRM received a rating of 5.

As seen in Table 5, under the low-risk condition, participants generally preferred the responses sourced from the classifier response types, while the opposite effect was observed under the high-

Users’ Stage 1 response type ($S1$) versus a <i>different</i> response type (d)	Low risk	High risk
$Rating(R_{S1}) > Rating(R_d)$	32	55
$Rating(R_{S1}) = Rating(R_d)$	28	25
$Rating(R_{S1}) < Rating(R_d)$	47	46
# of requests where $S1 \neq d$	107	126

Table 5: Comparison between participants’ ratings of responses sourced from their Stage 1 response types and responses sourced from different classifier-selected response types

risk condition (these findings are corroborated by the results in Table 7). Nonetheless, when we tested the hypothesis that participants liked responses sourced from their own previous response types as much as responses sourced from different response types in *AuthorCorpus* and different response types chosen by the classifier, both tests returned the same result: *there were no statistically significant differences between users’ ratings of responses sourced from their own Stage 1 response types and their ratings of responses sourced from different response types under the low- and high-risk conditions* (p -value > 0.15).

How does situational risk affect participants’ attitudes toward different response types? As seen in Table 1, the proportion of Dos in *ResponseCorpus* decreased under the high-risk condition, while the proportion of the other response types increased (the difference between the low- and high-risk response types is statistically significant, χ^2 with p -value $\ll 0.01$). This indicates that participants preferred more conservative (risk-averse) response types under the high-risk condition.

Figure 2 suggests that participants were also more critical of their own previous response types under the high-risk condition than under the low-risk condition (they reduced the ratings of 169 response types under the high-risk condition compared to only 133 under the low-risk condition). This observation is confirmed by the mean ratings of the Stage 2 responses in our corpora under the low- and high-risk conditions, which are shown in Table 7 for the responses sourced from *ResponseCorpus* and the responses obtained from the RF classifier (the *AuthorCorpus* results are similar).

In addition, the ratings of Do and of both versions of CHOOSE were significantly lower under the high-risk condition than under the low-risk condition (p -value $\ll 0.01$ for Do and CHOOSE+text, and p -value < 0.05 for CHOOSE+pictures). In con-

Top four ASR outputs	a. get the paint on the wall b. get the paint on the walls c. get the paint on the world d. <i>get the painting on the wall</i>	a. get the green light next to the blue plate b. get the green light next to the Blue Plate c. get the green light next to the blue planet d. get the green light next to the blue plates
Figure, requested object	I(a), C	I(b), E
Situational risk	High	High
Stage 1 response type	REPHRASE (rating: 1)	CHOOSE (rating: 1)
Stage 2 preferred response type	CONFIRM (rating: 5)	CONFIRM (rating: 5)
Top four ASR outputs	a. <i>move the green book rack</i> b. move the Greene book rack c. move the Green Book rack d. move the green book RAC	a. get the blue light on the left corner of the table b. <i>get the blue plate on the left corner of the table</i> c. get the bloop light on the left corner of the table d. get the Blue Planet on the left corner of the table
Figure, requested object	I(d), K	I(c), H
Situational risk	Low	Low
Stage 1 response type	DO (rating: 1)	CHOOSE (rating: 3)
Stage 2 preferred response type	CONFIRM (rating: 4)	DO (rating: 5)

Table 6: Examples where users gave lower ratings in Stage 2 to responses sourced from their selected Stage 1 response types than to responses sourced from different response types chosen by the RF classifier; the correct ASR output is italicized

<i>ResponseCorpus</i>		RF Classifier	
Low risk	High risk	Low risk	High risk
3.99 (1.31)	3.59 (1.49)	4.09 (1.29)	3.54 (1.49)

Table 7: Mean (Stdev) of response ratings under low- and high-risk conditions

trast, no statistically significant differences were found with respect to CONFIRM and REPHRASE under the two risk conditions. Also, participants preferred CONFIRM to DO and CHOOSE+pictures to CHOOSE+text under both risk conditions ($p\text{-value} \ll 0.01$).

These findings suggest that situational risk influences the acceptability of certain response types, but further research is required to identify these response types in a broader context.

6 Conclusion

We have offered a longitudinal study where participants initially selected response types for ASR outputs of spoken requests; and after some time, they rated responses sourced from their own response types, as well as responses sourced from other response types. Our results show that the participants did not think that their original choices were the best, and that overall, they had the same opinion of responses sourced from their own response types, the response types chosen by one of the authors and those selected by a classifier trained on the response types of the author. These findings suggest that, at least in the context of one-shot dialogues with a household robot, people’s response preferences at a particular point in time may not reflect their general attitudes, and that var-

ious reasonable responses may be equally acceptable. Our results also indicate that, at least in this context, a classifier trained on a small non-target but relevant corpus may yield adequate responses.

Our experiment also distinguished between two types of situational risk: low and high. We found that risk influences people’s general attitudes toward responses — they were more risk averse and critical under high-risk conditions than under low-risk conditions. However, this attitude was directed toward some response types (DO and CHOOSE) and not others (CONFIRM and REPHRASE). This finding, if generalized, may influence response type selection.

The implications of our findings for deep-learning systems are that training on a single best response may be unjustified, as several responses are equally acceptable. Further studies are required to determine whether our findings generalize to longer dialogues in more complex domains. If this is the case, (PO)MDP/RL systems do not need to take into account people’s preferences when generating a response. However, if extra-linguistic factors such as risk come into play, they should be incorporated into policy-learning algorithms to bias response selection in favour of risk-sensitive responses preferred by people. Finally, our findings regarding rating inconsistency over time may affect the results of comparative studies, such as that of Liu et al. (2016).

7 Acknowledgments


This research was supported in part by grants DP1201001030 and LP150100060 from the Australian Research Council.

References

- X. Amatriain, J.M. Pujol, and N. Oliver. 2009. I like it ... I like it not: Evaluating user ratings noise in recommender systems. In *UMAP'2009 – Proceedings of the 2009 Conference on User Modeling Adaptation and Personalization*, pages 247–258, Trento, Italy.
- L. Carlson. 1983. *Dialogue Games: An Approach to Discourse Analysis*. D. Reidel Publishing Company, Dordrecht, Holland, Boston.
- B. Dhingra, L. Li, X. Li, J. Gao, Y.N. Chen, F. Ahmed, and L. Deng. 2017. Towards end-to-end reinforcement learning of dialogue agents for information access. In *ACL'17 – Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 484–495, Vancouver, Canada.
- K. Funakoshi, M. Nakano, T. Tokunaga, and R. Iida. 2012. A unified probabilistic approach to referring expressions. In *SIGDIAL'2012 – Proceedings of the 13th SIGdial Meeting on Discourse and Dialogue*, pages 237–246, Seoul, South Korea.
- M. Gašić and S.J. Young. 2014. Gaussian processes for POMDP-based dialogue manager optimization. *IEEE/ACM Transactions on Audio, Speech & Language Processing*, 22(1):28–40.
- E. Horvitz, C. Kadie, T. Paek, and D. Hovel. 2003. Models of attention in computing and communication: From principles to applications. *Communications of the ACM*, 46(3):52–57.
- F. Jurčićek, S. Keizer, M. Gašić, F. Mairesse, B. Thomson, K. Yu, and S.J. Young. 2011. Real user evaluation of spoken dialogue systems using Amazon Mechanical Turk. In *Proceedings of Interspeech 2011*, pages 3061–3064, Florence, Italy.
- O. Lemon. 2011. Learning what to say and how to say it: Joint optimisation of spoken dialogue management and natural language generation. *Computer Speech and Language*, 25(2):210–221.
- J. Li, W. Monroe, A. Ritter, M. Galley, J. Gao, and D. Jurafsky. 2016. Deep reinforcement learning for dialogue generation. In *EMNLP2016 – Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Austin, Texas.
- W. Liao, W. Zhang, Z. Zhu, Q. Ji, and W.D. Gray. 2006. Toward a decision-theoretic framework for affect recognition and user assistance. *International Journal of Human-Computer Studies*, 64:847–873.
- C-W. Liu, R. Lowe, I.V. Serban, M. Noseworthy, L. Charlin, and J. Pineau. 2016. How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *EMNLP2016 – Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132, Austin, Texas.
- R. Moratz and T. Tenbrink. 2006. Spatial reference in linguistic human-robot interaction: Iterative, empirically supported development of a model of projective relations. *Spatial Cognition & Computation: An Interdisciplinary Journal*, 6(1):63–107.
- N. Mrkšić, Ó.S. Diarmuid, T.H. Wen, B. Thomson, and S.J. Young. 2017. Neural belief tracker: Data-driven dialogue state tracking. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 1777–1788, Vancouver, Canada.
- T. Paek and E. Horvitz. 2000. Conversation as action under uncertainty. In *UAI-2000 – Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence*, pages 455–464, Stanford, California.
- A. Prakash, C. Brockett, and P. Agrawal. 2016. Emulating human conversations using convolutional neural network-based IR. In *Proceedings of the NeurIR'16 SIGIR Workshop on Neural Information Retrieval*, Pisa, Italy.
- B. Rohrmann. 2005. Risk attitude scales: Concepts, questionnaires, utilizations. Technical report, University of Melbourne.
- A. Said and A. Bellogín. 2018. Coherence and inconsistencies in rating behavior: Estimating the magic barrier of recommender systems. *User Modeling and User-Adapted Interaction*, 28:97–125.
- I.V. Serban, T. Klinger, G. Tesauro, K. Talamadupula, B. Zhou, Y. Bengio, and A. Courville. 2017. Multiresolution recurrent neural networks: An application to dialogue response generation. In *AAAI-17 – Proceedings of the 31st National Conference on Artificial Intelligence*, pages 3288–3294, San Francisco, California.
- S. Siegel and N.J. Castellan. 1988. *Non-Parametric Statistics for the Behavioral Sciences*, second edition. McGraw–Hill, Inc.
- S. Singh, D. Litman, M. Kearns, and M. Walker. 2002. Optimizing dialogue management with reinforcement learning: Experiments with the NJFun system. *Artificial Intelligence Research*, 16:105–133.
- K. Sugiura, N. Iwahashi, H. Kashioka, and S. Nakamura. 2009. Bayesian learning of confidence measure function for generation of utterances and motions in object manipulation dialogue task. In *Proceedings of Interspeech 2009*, pages 2483–2486, Brighton, United Kingdom.
- B. Thomson, M. Gašić, S. Keizer, F. Mairesse, J. Schatzmann, K. Yu, and S. Young. 2008. User study of the Bayesian update of dialogue state approach to dialogue management. In *Proceedings of Interspeech 2008*, pages 483–486, Brisbane, Australia.

- J.G. Trafton, N.L. Cassimatis, M.D. Bugajska, D.P. Brock, F.E. Mintz, and A.C. Schultz. 2005. Enabling effective human-robot interaction using perspective-taking in robots. *IEEE Transactions on Systems, Man and Cybernetics – Part A: Systems and Humans*, 35(4):460–470.
- B-H. Tseng, F. Kreyssig, P. Budzianowski, I. Casanueva, Y-C. Wu, S. Ultes, and M. Gašić. 2018. Variational cross-domain natural language generation for spoken dialogue systems. In *SIGDIAL'2018 – Proceedings of the 19th SIGdial Meeting on Discourse and Dialogue*, pages 338–343, Melbourne, Australia.
- S. Ultes, P. Budzianowski, I. Casanueva, L.M. Rojas Barahona, B-H. Tseng, Y-C. Wu, S. Young, and M. Gašić. 2018. Addressing objects and their relations: The Conversational Entity Dialogue Model. In *SIGDIAL'2018 – Proceedings of the 19th SIGdial Meeting on Discourse and Dialogue*, pages 273–283, Melbourne, Australia.
- T.H. Wen, M. Gašić, N. Mrkšić, P. Hao Su, D. Vandyke, and S.J. Young. 2015. Semantically conditioned LSTM-based natural language generation for spoken dialogue systems. In *EMNLP2015 – Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1711–1721, Lisbon, Portugal.
- J.D. Williams and S. Young. 2007. Partially observable Markov decision processes for spoken dialog systems. *Computer Speech and Language*, 21(2):393–422.
- J.D. Williams and G. Zweig. 2016. End-to-end LSTM-based dialog control optimized with supervised and reinforcement learning. *arXiv preprint arXiv:1606.01269*.
- X. Yang, Y.N. Chen, D. Hakkani-Tür, P. Gao, and L. Deng. 2017. End-to-end joint learning of natural language understanding and dialogue manager. In *ICASSP'2017 – IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 5690–5694, New Orleans, Louisiana.
- S.J. Young, M. Gašić, S. Keizer, F. Mairesse, J. Schatzmann, B. Thomson, and K. Yu. 2010. The Hidden Information State Model: A practical framework for POMDP-based spoken dialogue management. *Computer Speech and Language*, 24(2):150–174.
- S.J. Young, M. Gašić, B. Thomson, and J. Williams. 2013. POMDP-based statistical spoken dialogue systems: A review. *Proceedings of the IEEE*, 101(5):1160–1179.
- T. Zhao and M. Eskenazi. 2016. Towards end-to-end learning for dialog state tracking and management using deep reinforcement learning. In *SIGDIAL'2016 – Proceedings of the 17th SIGdial Meeting on Discourse and Dialogue*, pages 1–10, Los Angeles, California.
- I. Zukerman, S.N. Kim, Th. Kleinbauer, and M. Mosh-taghi. 2015. Employing distance-based semantics to interpret spoken referring expressions. *Computer Speech and Language*, 34:154–185.

A Screenshots for Stage 1 and Stage 2



2.1. The ASR has returned the following alternatives for a particular spoken request in the context of the above image:

- a. get the CPU under the table
- b. get ICP you under the table
- c. get icpu under the table
- d. get SCP you under the table

Assuming that you are in the same room as the speaker, we would like you to choose how would you respond to this request, given all the alternative texts returned by the ASR and the image above. You may choose one of the following responses:

- Get object #
- Did you mean object #
- Which of these objects did you mean? (enter the numbers separated by blanks)

Ask the speaker to rephrase one of the following:

- The part about the intended object
- The part about the position of the intended object
- The whole sentence


2.2. Now assume the speaker is in a remote location, and the requested action would be significantly more time consuming than simply handing over the intended item. Would your response be different? Please tick the appropriate button.

- Same response
- Get object #
- Did you mean object #
- Which of these objects did you mean? (enter the numbers separated by blanks)

Ask the speaker to rephrase one of the following:

- The part about the intended object
- The part about the position of the intended object
- The whole sentence

Figure 3: Screenshot for Stage 1



1. The ASR has returned the following alternatives for a spoken request in the context of the above image:

- get the CPU under the table
- get ICP you under the table
- get icpu under the table
- get SCP you under the table

How would you rate the following responses?


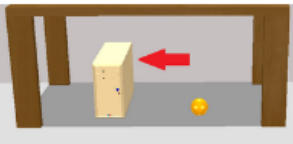
	Low-risk condition (requester is in the same room as you)					High-risk condition (requester is in a far-away location)				
	Very Unsuitable	Somewhat unsuitable	Neutral	Somewhat suitable	Very suitable	Very unsuitable	Somewhat unsuitable	Neutral	Somewhat suitable	Very Suitable
1. I didn't hear what you want from under the table. There are two things under the table. Which one do you want?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
2. I didn't hear what you want from under the table. There are two things under the table. Do you want this or that? 	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
3. Is this what you want? 	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
4. The robot just gets an object without asking any questions.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Figure 4: Screenshot for Stage 2

Characterizing the Response Space of Questions: a Corpus Study for English and Polish

Jonathan Ginzburg Zulipiye Yusupujiang Chuyuan Li Kexin Ren
Université de Paris, CNRS, Laboratoire de Linguistique Formelle
yonatan.ginzburg@univ-paris-diderot.fr

Paweł Łupkowski
Adam Mickiewicz University, Poznań
Pawel.Lupkowski@amu.edu.pl

Abstract

The main aim of this paper is to provide a characterization of the response space for questions using a taxonomy grounded in a dialogical formal semantics. As a starting point we take the typology for responses in the form of questions provided in (Łupkowski and Ginzburg, 2016). This work develops a wide coverage taxonomy for question/question sequences observable in corpora including the BNC, CHILDES, and BEE, as well as formal modelling of all the postulated classes. Our aim is to extend this work to cover *all* responses to questions. We present the extended typology of responses to questions based on a corpus studies of BNC, BEE and Map-task with include 506, 262, and 467 question/response pairs respectively. We compare the data for English with data from Polish using the Spokes corpus (205 question/response pairs). We discuss annotation reliability and disagreement analysis. We sketch how each class can be formalized using a dialogical semantics appropriate for dialogue management.

1 Introduction

There are various theories of what questions are (Groenendijk and Stokhof, 1997; Wiśniewski, 2015), and several computational theories of dialogue (Poesio and Rieser, 2010; Asher and Lascarides, 2003; Ginzburg, 2012), but no attempt yet at a comprehensive characterization of the response space of queries.

This task, nonetheless, is of considerable theoretical and practical importance: it is an important ingredient in the design of dialogue systems, spoken or text-based; it provides benchmarks for dialogue/question theories, and of course is a component in explicating intelligence to pass the Turing test (Turing, 1950).

(Łupkowski and Ginzburg, 2013, 2016) tackled one part of this problem, offering an empirical and theoretical characterization of the range of *query*

responses to a query. Based on a detailed analysis of the British National Corpus and three other corpora, two task-oriented (BEE (Rosé et al., 1999) and AmEx (Kowtko and Price, 1989)) and a sample from CHILDES (MacWhinney, 2000), they identified 7 classes of questions that a given query gives rise to; we refer to these classes as the L(upkowski)G(inzburg) classes of question responses.¹ We take their work as a starting point and make the following hypothesis:

- (1) Main hypothesis: responses drawn from or concerning the LG classes plus direct and indirect answerhood exhaust the response space of a query.

Specifically this amounts to the following general types of responses (we present the detailed taxonomy in section 3).

1. Question-Specific:
 - (a) Answerhood;
 - (b) Dependent queries (A: Who should we invite? B: Who is in town?);
2. Clarification Requests.
3. Evasion responses:
 - (a) Ignore (address the situation, but not the question);
 - (b) Change the topic ('Answer my question');
 - (c) Motive ('Why do you ask?');
 - (d) IDK ('I don't know');

¹The study sample consisted of 1,466 query/query response pairs. As an outcome the following query responses (q-responses) taxonomy was obtained: (1) CR: clarification requests; (2) DP: dependent questions, i.e. cases where the answer to the initial question depends on the answer to a q-response; (3) MOTIV: questions about an underlying motivation behind asking the initial question; (4) NO ANSW: questions aimed at avoiding answering the initial question; (5) FORM: questions considering the way of answering the initial question; (6) QA: questions with a presupposed answer, (7) IGNORE: responses ignoring the initial question—for more details see (Łupkowski and Ginzburg, 2016, p. 355).

(e) Difficult to provide a response.

The hypothesis has to be understood *relationally*—one is not really interested in the extension of the semantic entities (primarily propositions and questions) that can be given as responses. Rather, as exemplified in (2), one is interested in the class each such entity is classified as since that is what determines the subsequent contextual evolution.

- (2) I do not want to talk about that question.
 (Direct answer to *what do you not want to do?* Evasion answer to *Where were you last night?*).

We provide a brief discussion of the existing literature in section 2. Following this, we provide a description of the proposed taxonomy, in section 3. We then set out to test our main hypothesis in an initial study, using three corpora in English (BNC, BEE, MapTask) and one corpus in Polish (Spokes (Pezik, 2015)). By and large, the hypothesis achieves wide coverage, as we discuss in section 5. We sketch an account of how the different classes can be characterized, taking a fairly general perspective and building on the initial characterization of (Łupkowski and Ginzburg, 2016) while drawing some metatheoretical conclusions. Finally, section 8 offers a variety of extensions we plan to undertake.

2 Related work

Berninger and Garvey (1981) introduce their rich taxonomy of possible replies for children conversation in a nursery school. The taxonomy covers six categories, categories that are co-extensive with the ones mentioned in the introduction to this paper, though no semantic explication or interannotator study is offered: (i) Indirect answers. (ii) Confessions of ignorance. (iii) Clarification questions. (iv) Evasive replies. (v) Miscellaneous.

An extensive 10-language comparative project on question/response sequences in ordinary conversation was carried out from 2007 as the part of the Multimodal Interaction Project at the Max Planck Institute for Psycholinguistics (Stivers et al., 2010). The coding scheme for the response types covered categories of **Non-response**, **Non-answer response**, **Answer**, and **Can't determine** (Stivers and Enfield, 2010, p. 2624).

The results were 76% answer responses, 19% non-answers, and 5% non-responses. (Stivers,

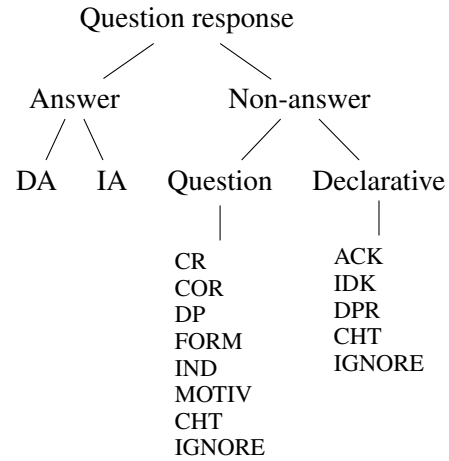


Figure 1: Response space of questions

2010, p. 2778) Interestingly, (Yoon, 2010) reports results for Korean which though indicative of a similar pattern (Answer > Non-Answer > Non-response) indicate a markedly different distribution: of the sample of 326 questions-responses, 52% were answers, 33% non-answers and 15% non-responses (Yoon, 2010, p. 2790). It is worth stressing that the question sample was limited to questions that functionally sought information, confirmation or agreement see (Yoon, 2010, p. 2783).

The work discussed in this section indicate the need for a wider corpus study of the whole spectrum of answers to questions.² The studies discussed are limited in terms of analyzed examples. They also imposed certain limitations in terms of numbers of response categories to be identified—they were mainly aimed at understanding the answer/non-answer difference. An extensive corpus study is needed for a fine grained characterization of the response space of questions. Moreover, we aim at providing an explicit dialogical semantics for each category of our corpus-based typology.

3 A taxonomy of responses to queries

We start with the most general division of question responses to answers and non-answers as discussed in the previous section. In the answer class we distinguish direct and indirect answers—see figure 1.

²For a detailed review of the literature on query responses, see (Łupkowski and Ginzburg, 2016), pp. 245–49, which discusses work from the question generation literature, in particular (Graesser et al., 1992).

Direct answers (DA) are (i) either sentential and denote propositions that are answers or (ii) are non-sentential and convey an answer as their content.³ This is clearly visible in the following example—B is providing information required by A:

- (3) A: Who is going to check that?
B: *Well I can check it.*

Indirect answers (IA) involve an inference of an answer from the utterance, as in (4):⁴

- (4) A: What is it?
A: What's he done?
B: *Ehm, you know what I've said before.*

Here A has to infer the answer to his/her questions from B's suggestion that this issue has been addressed before.

For the non-answer group the taxonomy (mostly) reuses the classes proposed in (Łupkowski and Ginzburg, 2013, 2016) with some minor renaming.

Clarification questions (CR) address something that was not completely understood in initial question (q1)⁵, like:

- (5) A: Why are you in?
B: *What?*

Corrections (COR) are declarative counterparts of CRs in that they assert rather than query about the original speaker's intended meaning. This is exemplified in B's answer in (6):

- (6) A: what is it?
A: Something forty <unclear>.
A: UB forty?
B: *WD forty.*

³ For the direct answers category we allow for additional sub-categories, which we return to discuss briefly in section 7. These include: (1) no/yes answer to polar questions; (2) simple answer to wh-questions; (3) partial polar answer; (4) partial wh-question answer.

⁴ As with the direct answers category, we have also used the following sub-categories of indirect answers, but do not elaborate on this here for reasons of space: (i) indirect answer addressing wh-question; (2) q-widening IAs (over-informative answer to a polar question, addressing a more general wh-question).

⁵ This class contains intended content queries, repetition requests and relevance clarifications—for detailed discussion see e.g. (Purver, 2006) or (Ginzburg, 2012).

A: WD.

Dependent questions (DP) constitute the case where the answer to the initial question (q1) depends on the answer to the query-response (q2), as in:

- (7) A: Do you want me to <pause> push it round?
B: *Is it really disturbing you?*
[cf. *Whether I want you to push it around depends on whether it really disturbs you.*]

See more in section 7.1.

Question responses may also address that the way the answer to q1 will be given depends on the answer to q2 (**FORM**). This type of question response differs from DP as the response concerns only the form in which the answer to q1 will be given (how it will be formulated). This may be noticed in (8), where the way B answers A's question will be dictated by A's answer to q2—whether or not A wants to know details point by point.

- (8) A: Okay then, Hannah, what, what happened in your group?
B: *Right, do you want me to go through every point?*

One also encounters q2, which is rhetorical and in this sense does not need to be answered and **indirectly provides an answer** to q1 (IND).

- (9) A: Are you Gemini?
B: *Well if I'm two days away from your, what do you think?*

As for evasive question-responses we have one type which addresses the **motivation underlying asking q1 (MOTIV)**. Whether an answer to q1 will be provided depends on a satisfactory answer to q2, as in the following example:

- (10) A: What's the matter?
B: *Why?*

Another type of evasive question-response is **change-the-topic (CHT)**. These are cases wherein q2 enables the speaker to avoid answering q1 while attempting to force the other speaker to answer q2 first. Instead of answering q1, the agent provides q2 and attempts to “turn the table” on the original querier. The original querier is pressured to answer q2 and put q1 aside.

- (11) A: Why is it recording me?
B: *Well why not?*

An **IGNORE** type of query-response appears when q2 relates to the situation described by q1 but not directly to the initial question:

- (12) A: I've got Mayfair <pause> Piccadilly, Fleet Street and Regent Street, but I never got a set did I?
B: *Mum, how much, how much do you want for Fleet Street?*

A and B are playing Monopoly. A asks a question, which is ignored by B. It is not that B does not wish to answer A's question and therefore asks q2. Rather, B ignores q1 and asks a question related to the situation (in this case, the board game). See also the following example:

- (13) A: Just one car is it there?
B: *Why is there no parking there?*

If a question response is not an answer and it is a declarative we consider the following cases. For a start declarative responses can serve the same purpose as ignoring query-response:

- (14) a. A: So does that mean that the ammeter is not part of the series, just hooked up after to the tabs?
B: *Let's take a step back.*
b. A: What have you been doing Melvin? <laugh>
B: *I ain't talking cos you've got that bloody thing on.*

Acknowledgement (ACK)—a speaker acknowledges that s(he) has heard the question, e.g. *mhm, aha* etc.

- (15) A: that's about it innit?
B: *Mm mm.*

The speaker states that s(he) **does not know the answer (IDK)**.

- (16) A: When's the first consignment of Scottish tapes?
B: *Erm <pause> don't know.*

The speaker states that it is **hard to provide an answer (DPR)**, points at a different information source, etc.

- (17) A: Why?
B: *I'm not exactly sure.*

An utterance signals that speaker does not want to answer, s(he) **changes the topic**, gives an evasive answer (CHT).⁶

- (18) A: What's dolly's name?
B: *It's raining.*

4 Corpus data used for the study

In order to test our main hypothesis, we used corpora from two languages, English and Polish.

4.1 English: BNC, BEE, MapTask

The data for English comes from the BNC, BEE, and the MapTask corpora (Burnard, 2000; Rosé et al., 1999; Anderson et al., 1991). 506 Q-R turns were taken from the BNC, 256 Q-R turns from BEE, and 467 Q-R turns from the MapTask. In each case starting points where questions occur were chosen by randomly selecting turn numbers, and coding the subsequent questions in that extract. Questions were turn units ending with a '?'; however, tag questions and turns with missing text (the BNC's 'unclear') were eliminated from considerations. The BNC data covers mainly topically unrestricted conversations. As for BEE and MapTask dialogues are more task oriented—BEE contains tutorial dialogues from electronics courses and MapTask consists of dialogues recorded for a direction-providing task.

4.2 Polish: the Spokes Corpus

The data used for this study was drawn from the Spokes corpus (Pezik, 2015). The corpus currently contains 247,580 utterances (2,319,291 words) in

⁶These can occur in text as well:

- (i) So, in answer to the question: Is Jeremy Corbyn an anti-Semite? My response would be that that's the wrong question. The right questions to ask are: Has he facilitated and amplified expressions of anti-Semitism? Has he been consistently reluctant to acknowledge expressions of anti-Semitism unless they come from white supremacists and neo-Nazis? Will his actions facilitate the institutionalisation of anti-Semitism among other progressives? Sadly, my answer to all of these is an unequivocal yes. (D Lipstadt, *Antisemitism: Here and Now*)

transcriptions of spontaneous conversations. For the study four files were selected from the corpus (10,244 words, 1,424 turns)⁷. Within each file the question-response pairs (Q-R) were selected manually. In total we obtained 205 Q-R pairs for the study.

5 Results

For the annotation all the question-response pairs were supplemented with a full context. The guideline for annotators contained explanations of all the classes and examples for each category. Also the OTHER category was included. The tagset used to annotate gathered data is presented in Table 1. The detailed results of the annotation are presented in figure 2. We discuss the annotation reliability in section 6.

5.1 English

In all three cases, the OTHER class is less than 3%, hence coverage is above 97%. The most frequent classes of responses in all three corpora are direct answers (DA); in the BNC the next biggest are clarification requests, for BEE these are indirect answers, whereas for the MapTask the second biggest are IGNORE.

5.2 Polish

The two most frequent classes of responses for Spokes are answers: direct ones (DA=51.71%) and—much smaller—indirect ones (IA=13.66%). The next two most frequent classes are IDK (stating that a person does not know the answer to the question, IDK=10.24%) and utterances ignoring the question asked (questions and declaratives, IGNORE=9.76%).

5.3 Discussion

As might be expected from the results presented in (Łupkowski and Ginzburg, 2016), the most frequent question-response for English and Polish data is the clarification request. What is more surprising is that by adding declaratives into the picture a relatively high number of ignoring responses is observed for both English and Polish. Łupkowski and Ginzburg (2016) analyzed only question-responses and this type was observed rarely (0.57% for n=1,051 for BNC). Other evasive responses (relatively) frequent in both lan-

⁷Files 016O, 019w, 01AO, 01dL cover casual conversation concerning youth, wine and travelling plans.




Figure 2: Frequency of responses to questions for the BNC (n=506), BEE (n=256), MapTask (n=467) and Spokes (n=205) studies

guages are CHT and IDK. For the latter, we observe that it was more frequent in Polish than in the English data. This may be a consequence of the lower number of examples analyzed for Polish—Spokes is smaller and less varied than the BNC.

As regards cross-corpus differences, BNC and Spokes data cover mainly topically unrestricted conversations, while BEE and MapTask contain task-oriented dialogues. Correspondingly, MapTask has the highest number of direct answers (79.0%), and BEE almost the same (77.5%). However, for BNC and Spokes these numbers are lower (respectively 61.26% and 51.71%). For both clarification requests and evasive response types frequencies are lower for task-oriented corpora than for BNC and Spokes (this is in line with results for BNC and BEE reported in (Łupkowski and Ginzburg, 2016, p. 256–257)).

6 Annotation reliability

6.1 Inter-annotator studies

Table 1: Tagset used for annotation of the data

Category	TAG
1. Direct answer	DA
2. Indirect answer	IA
3. Clarification response	CR
4. Dependent question	DP
5. The utterance does not relate to the question, but to the situation	IGNORE
6. Question being an indirect answer	IND
7. Question addressing the form of answer to be given	FORM
8. Question about the motivation for the initial question	MOTIV
9. I do not know	IDK
10. Difficult to provide an answer	DPR
11. Correction	COR
12. Acknowledgement	ACK
13. Utterance signals that speaker does not want to answer, s(he) changes the topic, gives evasive answer	CHT
14. Utterance that does not fit in any of the above	OTHER

For English: For the inter-annotator study a sample of nearly 800 Q-Rs from the BNC were annotated by two advanced graduate students in computational linguistics, L2 speakers of English, who underwent several training sessions with one of the authors, a native speaker of English with significant experience in dialogue annotation. The first annotator coded 622 Q-Rs and the second annotator annotated 730 Q-Rs. Then we chose the initial 515 Q-Rs, which were commonly annotated

by both annotators, deleting 9 Q-Rs which were incomplete or unclear utterances to yield the 506 commonly annotated QR pairs from the BNC. For these we calculated the κ (Carletta, 1996) and α (Krippendorff, 2011) measures. We used the data mining and data analysis tool (Pedregosa et al., 2011) in Python with its *sklearn.metrics* package for calculating Cohen’s kappa, and also used the Python implementation *Krippendorff*⁸ for the calculation of Krippendorff’s alpha. In this case, Cohen’s Kappa for two annotators is 0.65 (substantial), and Krippendorff’s alpha is 0.66. All disagreements were then discussed in detail by one of the annotators and the afore-mentioned author and resolved (though some ambiguous cases remain, as discussed below.).

For Polish: The entire sample of 205 Q-Rs was annotated by the main annotator and two other annotators (one of whom has previous experience in corpus data annotation, all annotators were Polish native speakers). Fleiss’ Kappa for all three annotators was 0.53 (i.e. moderate). For the first and the second annotator—Cohen’s Kappa 0.66 (substantial). For the first and the third annotator—Cohen’s Kappa 0.49 (moderate).⁹ Krippendorff’s alpha for all three annotators is 0.742. For the first and second annotator the score is 0.617, while for the first and the third annotator it is 0.379. All measures were calculated using the *irr* package (Gamer et al., 2012) from R (R Core Team, 2013), version 3.3.1.

Disagreement analysis For reasons of space, we restrict attention to English here. Among the valid commonly annotated 506 BNC Q-Rs, there are 94 cases where the annotation disagreements between two annotators occurred. The main disagreements concerned DA versus IA (34), IGNORE versus CHT/ACK/DP/DA (16), and ACK versus OTHER (5), as exemplified in (19). Invariably, the direct/indirect disagreements occurred with ‘why’, ‘how’ and ‘what is X doing’ questions, where answers are by and large sentential and for which there has been significant controversy in the theoretical literature on how to characterize answerhood (Kuipers and Wiśniewski, 1994; Asher and Lascarides, 1998).

⁸<https://pypi.org/project/krippendorff/>

⁹Whereas the first and second annotators have much experience in dialogue annotation, the third annotator is a logician with less annotation experience.

(19) a. ANON5: Why do they pretend not to know?

ANON5: <pause> I mean they should be fully aware of of of our <unclear>

ANON2: **Val, well this is a new guy.**
[DA v. IA, resolved to IA.]

b. ANN: That's not very nice.

STUART: It is.

ANN: No It isn't.

STUART: Well it is. Why isn't it?

ANN: **Cos it isn't.** [DA v. IGNORE, resolved to IA since indirectly indicates that there is no reason.]

c. JOHN: Can you spell box?

SIMON: **Mhm.** [ACK v. OTHER, resolved to DA, after consideration of surrounding context.]

After carefully discussing all disagreements, we concluded that there are (at least) 10 cases which are truly ambiguous and should not be resolved; this is in line with a recent trend in dialogue annotation (e.g., Passonneau and Carpenter, 2014); though we have not implemented the more complex approach this inevitably requires in the current work. We exemplify two such cases. (20a,b) involve an ambiguity between CR and IND, and DA and IA, respectively; both are hard to resolve conclusively.

(20) a. FRANCIS: What is five?

FRANCIS: Tell me <unclear>.

UNKNOWN: <pause> **is there five people?**

b. HUG: What's he working on Rog?

ROG: **Oh he's off work <unclear> and you see he has all the time off for councils and you know it isn't as if he's there fulltime.**

7 Formal Analysis

In this section, we discuss briefly the requirements on a computational semantic theory to be able to characterize the response space of a query in terms of the notions discussed in previous sections. Łupkowski and Ginzburg (2016) assume such a characterization should be formulated in dialogical terms, for instance as dynamics of agent information states, since this makes the analysis usable for dialogue analysis. Indeed, to the extent that the empirical work here verifies our main hypothesis (1), the formal rules provided in (Łupkowski and Ginzburg, 2016) yield a complete characterization of the response space for questions in implementable form (for a sketch see (Maraev et al., 2018)). However, using a proof theoretic approach along the lines of erotetic logics like IEL (Wiśniewski, 2013) is conceivable, assuming it can be extended in certain respects, as we will explain.

7.1 Question-specificity

Any speaker of a given language can recognize, independently of domain knowledge and of the goals underlying an interaction, that certain propositions are *about* or *directly concern* a given question. This is the answerhood relation needed for characterizing direct answerhood.

The most basic notion of answerhood—*simple answerhood* (Ginzburg and Sag, 2000)—is the range of the propositional abstract, plus their negations.

(21) a. $\text{SimpleAns}(\lambda\{ \}p) = \{p, \neg p\};$

b. $\text{SimpleAns}(\lambda x.P(x)) = \{P(a), P(b), \dots, \neg P(a), \neg P(b) \dots\}$

In fact, *simple answerhood*, though it has good coverage, is not sufficient. *Aboutness* must be sufficiently inclusive to accommodate conditional, weakly modalized, and quantificational answers, all of which are pervasive in actual linguistic use (Ginzburg and Sag, 2000).

How to formally and empirically characterize aboutness is an interesting topic researched within work on the semantics of interrogatives (see e.g. Ginzburg and Sag, 2000; Groenendijk, 2006), though a comprehensive, empirically-based, experimentally tested account for a variety of wh-words is still elusive.

An additional important notion a theory of questions needs to provide for is a notion of *exhaustiveness*, though this is in general pragmatically parametrized (Asher and Lascarides, 2003). Whether a response is (pragmatically) exhaustive (or *goal fulfilling*) can determine whether the response will be accepted or require a follow up query. Hence, the need for a finer-grained subdivision of the answer categories, as we hinted in footnotes 3 and 4.

Given a notion of aboutness and some notion of (partial) exhaustiveness, one can then define question dependence (needed for the class DP), for instance, as in (22), though various alternative definitions have been proposed (Groenendijk and Stokhof, 1997; Wiśniewski, 2013; Onea, 2016). For all these definitions their coverage awaits testing on empirical data:

- (22) q_1 depends on q_2 iff any proposition p such that p resolves q_2 , also satisfies p entails r such that r is about q_1 . (Ginzburg, 2012, (61b), p. 57)

With notions of aboutness and dependency in hand, one can define update rules licensing such responses. For instance, a rule of the following form:

- (23) QSPEC: If q is the question under discussion, respond with an utterance r which is q -specific: About(r,q) or Depends(q,r)

7.2 Repair utterances

Clarification requests and (metacommunicative) corrections is a domain where logics that use simply contents of utterances are not adequate (Ginzburg and Cooper, 2004). Their generation requires access to the entire sign associated with a given interrogative utterance. (Purver, 2004; Ginzburg, 2012) show how to account for the main classes of CRs using rules that enable clarification questions relevant to a given utterance under clarification to be accommodated into the content. Each such rule specifies an accommodated MAX-QUD built up from a sub-utterance u_1 of the target utterance, the maximal element of the Pending attribute of the context (*MAX-Pending*). Common to all these rules is a license to follow up *MAX-Pending* with an utterance which is *co-propositional* with MAX-QUD.¹⁰ Abstracting

¹⁰ Two utterances are co-propositional if, modulo their domain, the questions they introduce into QUD involve similar

away from formal details, such rules can be specified as in (24), with the three disjuncts indicating the possible clarification questions that can be accommodated:

- (24) **Clarification Context Update Schema**
 Input: u : utterance by A, u_1 , constituent of u
 Output:
 MAXQUD:
 (i)reference resolution: *what did A mean by u_1* ,
 (ii)form resolution: *what word did A utter at u_1* ,
 (iii)confirmation of constituent content: *what is u_1 's content x , given that u 's content is $C(x)$*

7.3 Evasion Utterances

A natural way to analyze utterances relating to MOTIV is along the lines of a rule akin to QSPEC above: If A has posed q , B may follow up with an utterance specific to the issue *?Wish-Answer(B,q)*

(Łupkowski and Ginzburg, 2016) postulate fairly strong constraints on CHT and IGNORE to ensure that they are not unrestricted and do not allow any issue in. IGNORE is assumed to require the issue to be situationally shared with the posed question q_1 . This requires a means of evaluating shared-situatedness between questions. For CHT they assume that the topic changing question q_2 introduced by or addressed by the response must be unifiable with q_1 via a third question q_3 (e.g., q_1 = what do you (B) like? q_2 = what do you (A) like? q_3 = Who likes what?..). This requires a question inference mechanism for testing this unifiability.

8 Conclusions and Future Work

In this paper we have presented an initial study for what is, as far as we are aware, the first, detailed, formally underpinned characterization of the response space of queries. Achieving such a characterization is a fundamental challenge for semantics

answers—a query q introduces q into QUD, whereas an assertion p introduces $p?$ into QUD. For instance ‘Whether Bo left’, ‘Who left’, and ‘Which student left’ (assuming Bo is a student) are all co-propositional. Hence the available follow ups licensed in this way are clarification requests that differ from MAX-QUD at most in terms of its domain, or acknowledgements and corrections—propositions that instantiate MAX-QUD.

with a very wide variety of applications. It also establishes basic theoretical benchmarks for theories of dialogue/discourse and for semantic theories of questions.

Apart from the need to scale up the evidence quantitatively, we are currently engaged in work on the following strands:

- Cross-question type comparison: the Q-R pairs annotated in the current study were selected randomly, whereas it is clearly of interest to consider the distribution of responses relative to fixed classes of questions (e.g., different classes of wh-questions, polar questions etc.)
- Apply machine learning to acquire the response classification scheme:
 1. The learnability of non sentential answers (Fernández et al., 2007; Dragone and Lison, 2015) gives hope for learnability of some other classes.
 2. On the other hand, we anticipate significant difficulty with learning heavily inference-based classes like indirect answers, and IGNORE/CHT.
- Spoken dialogue system implementation: we plan to test the usability of these categories in dialogue systems with sophisticated dialogue management (Larsson and Berman, 2016) and NLU (see Maraev et al., 2018).
- Cross-linguistic testing: a significant challenge is how to test the classification with languages lacking large or even hardly any speech corpora. We anticipate using online games with a purpose to this end (see e.g., Łupkowski et al., 2018).

Acknowledgments

We acknowledge the support of the French Investissements d’Avenir-Labex EFL program (ANR-10-LABX-0083) and a senior fellowship from the Institut Universitaire de France to the first author, which funded the internships of Yusupujiang, Li, and Ren at LLF. We thank three anonymous reviewers for SigDial for their detailed and perceptive comments. A much earlier version of this work was presented at the workshop *Why indeed? Questions at the interface of theoretical and computational linguistics* in Stuttgart in March 2018. Thanks to the organizers, Annette Hautli-Janisz, Aikaterini-Lida Kalouli und Tatjana Schefler, and the audience for its feedback.

References

- Anne H. Anderson, Miles Bader, Ellen Gurman Bard, Elizabeth H. Boyle, Gwyneth M. Doherty, Simon C. Garrod, Stephen D. Isard, Jacqueline C. Kowtko, Jan M. McAllister, Jim Miller, Catherine F. Sotillo, Henry S. Thompson, and Regina Weinert. 1991. The HCRC Map Task Corpus. *Language and Speech*, 34(4):351–366.
- Nicholas Asher and Alex Lascarides. 1998. Questions in dialogue. *Linguistics and Philosophy*, 21(3):237–309.
- Nicholas Asher and Alex Lascarides. 2003. *Logics of Conversation*. Cambridge University Press.
- Ginger Berninger and Catherine Garvey. 1981. Relevant replies to questions: Answers versus evasions. *Journal of Psycholinguistic Research*, 10(4):403–420.
- L. Burnard. 2000. *Reference Guide for the British National Corpus (World Edition)*. Oxford University Computing Services.
- Jean Carletta. 1996. Assessing agreement on classification task: The kappa statistic. *Computational Linguistics*, 22(2):249–254.
- Paolo Dragone and Pierre Lison. 2015. An active learning approach to the classification of non-sentential utterances. In *Proceedings of the Second Italian Conference on Computational Linguistics*, pages 115–119.
- Raquel Fernández, Jonathan Ginzburg, and Shalom Lappin. 2007. Classifying ellipsis in dialogue: A machine learning approach. *Computational Linguistics*, 33(3):397–427.
- Matthias Gamer, Jim Lemon, and Ian Fellows Puspendra Singh. 2012. *irr: Various coefficients of interrater reliability and agreement*. Access 20.03.2017, R package version 0.84.
- Jonathan Ginzburg. 2012. *The Interactive Stance: Meaning for Conversation*. Oxford University Press, Oxford.
- Jonathan Ginzburg and Robin Cooper. 2004. Clarification, ellipsis, and the nature of contextual updates. *Linguistics and Philosophy*, 27(3):297–366.
- Jonathan Ginzburg and Ivan A. Sag. 2000. *Interrogative Investigations: the form, meaning and use of English Interrogatives*. Number 123 in CSLI Lecture Notes. CSLI Publications, Stanford: California.
- A. C. Graesser, N. K. Person, and J. D. Huber. 1992. Mechanisms that generate questions. In T. E. Lauer, E. Peacock, and A. C. Graesser, editors, *Questions and information systems*, pages 167–187. Lawrence Erlbaum Associates, Hillsdale.

- Jeroen Groenendijk. 2006. The logic of interrogation. In Maria Aloni, Alistair Butler, and Paul Dekker, editors, *Questions in Dynamic Semantics*, volume 17 of *Current Research in the Semantics/Pragmatics Interface*, pages 43–62. Elsevier, Amsterdam. An earlier version appeared in 1999 in the Proceedings of SALT 9 under the title ‘The Logic of Interrogation. Classical version’.
- Jeroen Groenendijk and Martin Stokhof. 1997. Questions. In Johan van Benthem and Alice ter Meulen, editors, *Handbook of Logic and Linguistics*. North Holland, Amsterdam.
- Jacqueline C. Kowtko and Patti J. Price. 1989. [Data collection and analysis in the air travel planning domain](#). In *Proceedings of the Workshop on Speech and Natural Language*, HLT ’89, pages 119–125, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Klaus Krippendorff. 2011. Agreement and information in the reliability of coding. *Communication Methods and Measures*, 5(2):93–112.
- Theo AF Kuipers and Andrzej Wiśniewski. 1994. An erotetic approach to explanation by specification. *Erkenntnis*, 40(3):377–402.
- Staffan Larsson and Alexander Berman. 2016. Domain-specific and general syntax and semantics in the talkmatic dialogue manager. *Empirical Issues in Syntax and Semantics*, 11:91–110.
- Paweł Łupkowski, Mariusz Urbański, Andrzej Wiśniewski, Wojciech Błądek, Agata Juska, Anna Kostrzewa, Dominika Pankow, Katarzyna Paluszkievicz, Oliwia Ignaszak, Joanna Urbańska, et al. 2018. Erotetic reasoning corpus. a data set for research on natural question processing. *Journal of Language Modelling*, 5(3):607–631.
- Paweł Łupkowski and Jonathan Ginzburg. 2013. A corpus-based taxonomy of question responses. In *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013)*, pages 354–361, Potsdam, Germany. Association for Computational Linguistics.
- Paweł Łupkowski and Jonathan Ginzburg. 2016. Query responses. *Journal of Language Modelling*, 4(2):245–293.
- Brian MacWhinney. 2000. *The CHILDES project: Tools for analyzing talk*, third edition. Lawrence Erlbaum Associates, Mahwah, NJ.
- Vladislav Maraev, Jonathan Ginzburg, Staffan Larsson, Ye Tian, and Jean-Philippe Bernardy. 2018. Towards KoS/TTR-based proof-theoretic dialogue management. In *Proceedings of SemDial 2018*, Aix-en-Provence.
- Edgar Onea. 2016. *Potential questions at the semantics-pragmatics interface*. Brill, Leiden, Boston.
- Rebecca J Passonneau and Bob Carpenter. 2014. The benefits of a model of annotation. *Transactions of the Association for Computational Linguistics*, 2:311–326.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Piotr Pezik. 2015. Spokes – a search and exploration service for conversational corpus data. In *Selected Papers from the CLARIN 2014 Conference, October 24-25, 2014, Soesterberg, The Netherlands*, 116, pages 99–109. Linköping University Electronic Press, Linköpings universitet.
- Massimo Poesio and Hannes Rieser. 2010. (prolegomena to a theory of) completions, continuations, and coordination in dialogue. *Dialogue and Discourse*, 1:1–89.
- Matthew Purver. 2004. *The Theory and Use of Clarification in Dialogue*. Ph.D. thesis, King’s College, London.
- Matthew Purver. 2006. Clarie: Handling clarification requests in a dialogue system. *Research on Language & Computation*, 4(2):259–288.
- R Core Team. 2013. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. Access 20.03.2017.
- Carolyn P. Rosé, Barbara Di Eugenio, and Johanna D. Moore. 1999. A dialogue-based tutoring system for basic electricity and electronics. In Susanne P. Lajoie and Martial Vivet, editors, *Artificial intelligence in education*, pages 759–761. IOS, Amsterdam.
- Tanya Stivers. 2010. An overview of the question–response system in american english conversation. *Journal of Pragmatics*, 42(10):2772–2781.
- Tanya Stivers, Nicholas J Enfield, and Stephen C Levinson. 2010. Question-response sequences in conversation across ten languages: an introduction. *Journal of Pragmatics*, 42:2615–2619.
- Tanya Stivers and Nick J Enfield. 2010. A coding scheme for question–response sequences in conversation. *Journal of Pragmatics*, 42(10):2620–2626.
- A.M. Turing. 1950. Computing machinery and intelligence. *Mind*, 59(236):433–460.
- Andrzej Wiśniewski. 2013. *Questions, Inferences, and Scenarios*. College Publications, London, England.
- Andrzej Wiśniewski. 2015. Questions. In *Handbook of Contemporary Semantic Theory, second edition*, Oxford. Blackwell.

Kyung-Eun Yoon. 2010. Questions and responses in Korean conversation. *Journal of Pragmatics*, 42(10):2782–2798.

From Explainability to Explanation: Using a Dialogue Setting to Elicit Annotations with Justifications

Nazia Attari

Research Institute for Cognition and Robotics
Bielefeld University, Germany
nattari@techkfak.uni-bielefeld.de

Martin Heckmann

Honda Research Institute Europe
Germany

David Schlangen

Computational Linguistics
University of Potsdam, Germany

Abstract

Despite recent attempts in the field of *explainable AI* to go beyond black box prediction models, typically already the training data for supervised machine learning is collected in a manner that treats the annotator as a “black box”, the internal workings of which remains unobserved. We present an annotation method where a task is given to a pair of annotators who collaborate on finding the best response. With this we want to shed light on the questions if the collaboration increases the quality of the responses and if this “thinking together” provides useful information in itself, as it at least partially reveals their reasoning steps. Furthermore, we expect that this setting puts the focus on *explanation* as a linguistic act, vs. *explainability* as a property of models. In a crowd-sourcing experiment, we investigated three different annotation tasks, each in a collaborative dialogical (two annotators) and monological (one annotator) setting. Our results indicate that our experiment elicits collaboration and that this collaboration increases the response accuracy. We see large differences in the annotators’ behavior depending on the task. Similarly, we also observe that the dialog patterns emerging from the collaboration vary significantly with the task.

1 Introduction

Imagine asking a friend whether you can borrow their car for the afternoon, and the only reply you get is “no”. You would presumably perceive this as somewhat brusque, and Conversation Analysis would back you up there: This kind of *dispreferred reply* typically needs more work, often being initiated with a filled pause, and being augmented with a *reason* for the refusal (Schegloff, 2007; Levinson, 1983). Now imagine you are asking a car rental place, via their website, whether you can rent a car for the afternoon, and again all

you get as a reply is a “no”. You would still not be pleased, but the difference here would be that while your friend may have been *unwilling* to tell you their reasons, the car rental company, having used a complex statistical model that judged you untrustworthy, based on various kinds of information it has about you, would be *unable* to state reasons (other than a vacuous one like “your score is too low”).

The field of *explainable AI* has set itself as a goal to open up the blackbox of current prediction models in order to make their decisions more transparent and also identifying problems concerning the core issues in AI safety. (See (Gilpin et al., 2018; Doshi-Velez and Kim, 2017; Ribeiro et al., 2016; Lundberg and Lee, 2017; Amodei et al., 2016) for recent overviews.) The focus there typically is on providing explanations of decisions in terms of examples or secondary models (e.g. (Kim et al., 2018; Letham et al., 2015; Yuan et al., 2019; Zhang et al., 2019)), where the resulting explanations are understandable at best to experts. In contrast, our interest is in learning to provide verbal explanations, accessible also to novice users. As a first step, we are interested in methods for eliciting data that can be used for this. In this paper, we present an annotation scheme where a pair of annotators works in collaboration to find the best response to a question. Our hypothesis is that a) this leads to better quality responses compared to non-collaborative annotation, as the annotators can actively acknowledge/correct/help their partners, b) the resulting discussions give access to the collaborative thinking directions that lead to the final response, and c) puts the focus on *explanation* as a linguistic act, vs. *explainability* as a property of models. We present results from three different annotation tasks. For each task we compare the accuracies of the responses we obtain in a dialog (two annotators) and a monologue (one annotator)

setting, analyze to what extent the task triggered discussions in the dialog setting and quantify dialog patterns emerging in the interaction of the annotators.

2 The Annotation Game

We formalise the annotation task as a game with the following structure. The annotation problem is posed by a special participant in the game, which we call *Nature* (N). N poses a question Q that is to be answered, and provides relevant information $I = \{i_1, \dots, i_n\}$. (e.g., $Q = \text{“what is in this image?”}$, with I consisting of an image.) Besides N , there is a set of regular participants in the game, $P = \{P_1, \dots, P_m\}$. The participants produce verbal turns $T = \{t_1, \dots, t_k\}$. In our setting, we assume that there is one special token that is used to flag a verbal turn T as a proposal for an answer A and another token to flag a turn as a mutual agreement on it; this type of game could hence also be called an *Agreement Game*.

Each solved task—that is, each annotation—can be represented as a tuple $\langle Q, I, A, T \rangle$. Our hypothesis is that the provided answers A , relative to the given information I and the respective question Q , are of higher quality in settings where $T \setminus A$ is non-empty compared to those where it is; (that is, where there has been interaction between the annotators) and moreover, that the turns $T \setminus A$ in the interactive case provide insights into the reasoning steps that are taken to perform the mapping from I to A , given Q —from which ultimately strategies for providing explanations could be learned.

3 Experiment

To test the hypotheses set out above, we created a number of tasks (pairs of questions Q and information I), which we put to individual annotators and also to pairs of annotators in a dialogical setting.

3.1 Example Tasks

Birds Here, we show images of birds of two different kinds, as in Figure 1. The task for the annotators is to produce a characteristic description of one of the two kinds; i.e., a description that is true for all and only the images in the specified row. Following the question $Q = \text{“what separates the birds in 1 from those in 2”}$ given the images I in Fig. (1) $A = \text{“large wingspan, grey plumage with$




Figure 1: An Example Birds Task




Figure 2: An Example Blocks Task

black head” would be a valid answer. The images are taken from the *Caltech-UCSD Birds 200-2011* dataset (Wah et al., 2011). Details on the setting can be found in the appendix.

Blocks This task consists in providing a characteristic rule for one of two artificial, programmatically-created categories in the form of blocks with patterns. A valid answer for the example in Figure 2 could be “B has six or more red blocks.” Note that in this kind of rule induction task from few examples, there will always be a large number of rules that correctly describe the pattern, even if they are different from the one that was actually used to generate the examples.

- 1 Daniel grabbed the milk there.
- 2 Sandra journeyed to the garden.
- 3 Sandra picked up the football there.
- 4 Sandra put down the football.
Where is the football?

Figure 3: An Example Texts Task

Texts To provide some variety, we also tested a text comprehension task, where a question about a text has to be answered; see Figure 3.

3.2 The Technical Setup & Collection

We realised the dialogical game interface as a web application, built on top of the chat server

slurk (Schlangen et al., 2018). The Mechanical Turk platform was used to recruit workers. After having read the instructions for the task, workers that accepted the task were transferred to a “waiting area” in the chat tool; as soon as a second worker entered this area, both the players were then moved to their task room (see also figure 4 in the appendix). The participants were paid an amount of \$0.14 per minute (for a maximum of 4 minutes per game, although they could discuss longer). We also paid a bonus amount of \$0.10 when the participants talked about things related to the task, tried to contribute equally during the discussion and also found the correct answer.

Additionally, we ran a monological version of the tasks with individual annotators, where we just presented the annotation task and collected the answer.

We collected 40 dialogues per setting, for a total of 120. Each dialogue consists of the consecutive discussion of two questions. After removing failed dialogues (where one participant left in the middle of the game, or participants clearly failed to follow the instructions), we were left with 93 dialogues: 28 for *birds*, 33 for *synthetic*, 32 for *text*. For monological annotation, we collected 40 annotations per setting, for a total of 120 annotations.

4 Results

4.1 Descriptive Overall Statistics

Table 1 shows some statistics about the collected data. In case of the dialogues, since the answers were marked by a prefix */answer*, we could automatically identify them and look at the *discussion* (everything but the answer) and the *answer(s)* separately. “Speaker contribution ratio” is a measure of how balanced the dialogue was in terms of contributions by each participant. It is the ratio between the number of tokens produced by the more talkative participant and the number of tokens produced by the other participant; a perfectly balanced dialogue would rate 1 here. We also looked at the ratio of turns by each speaker.

As these numbers show, the participants in the dialogues produced more tokens overall, and, for *Birds*, also longer answers. The dialogues tended to be dominated by one speaker. When taking out the outliers (ratio above 3.4), which were cases where one participant had to explain the task to an inattentive other player, the imbalance is lower, but still pronounced, whereas it is not as strong

Averages	Birds	Blocks	Texts
length (mins)	5.25	5.87	5.63
# turns	4.30	3.39	2.86
# turns w/o <i>As</i>	2.96	1.83	1.39
# tokens	39.61	28.09	18.33
# tokens, final <i>A</i>	14.43	11.41	7.48
speaker contr. ratio	3.46	5.53	6.13
... w/o outliers	2.85	4.15	4.30
speaker turn ratio	1.13	0.86	0.97
no discussion dlgs	35.7%	56.1%	57.8%
# tokens (monological)	11.45	12.52	9.45

Table 1: Statistical Overview of Data

in terms of turns. The numbers for *Blocks* and *Texts* are impacted by the high proportion of dialogue without any discussion (just */answer* followed by */agree*), as shown in row “no discussion dlgs”. Looking deeper into the dialogues, we found that in about 65% of the cases, the more dominant speaker was also the one who proposed the final answer.

4.2 The Answers

While we can automatically identify the proposed answers by the players, we cannot automatically evaluate them. For *Birds* and *Blocks*, a wide variety of answers could be considered correct; for *Texts*, there is a single correct answer, but different ways of phrasing it. Hence, we manually classified the answers as *correct* and *incorrect*.

Incorrect answers often betray a misunderstanding of the task, as with “The birds in Section 2 look like the same type of bird, or breed. The birds in Section 1 all look like different types of birds, or breeds” for *Birds*, or “Mary is not in the bathroom because the statement is in past tense” for *Texts*.

Table 2 shows the accuracy of the final answers across tasks and settings (dialogue and monologue). The accuracy is measured by comparing the 40 answers *A* to the corresponding 40 identical questions *Q* for each task used for the dialogues and monologues. These results indicate that the tasks seem to be of different difficulty, with *Birds* eliciting the highest number of correct replies, and the constructed, quite challenging *Blocks* task the least, across settings. The numbers for the monological setting are consistently lower, lending support to our hypothesis that the dialogical setting leads to improved quality in the answer.

Tasks	Correct Answer(%)	
	Dialogue	Monologue
Birds	92.5	85.0
Text	90.0	85.0
Blocks	57.5	50.0

Table 2: Dialogue vs. Monologue: Correct Answers

4.3 The Discussions

To further analyze the discussions, we first categorized them as *active* or *passive*. In an active discussion parts of the final answer is “rehearsed” before the official reply is given or the final answer is assembled out of several turns. The following is an example of this category (for the task shown in Figure 1 above).

- (1) A: Looks like the birds under 2 have red-orange feet.
 B: The difference that I notice is that the birds in Section 1 are light feathered vs. the dark feathered birds of Section 2.
 A: Ah, I like your answer better than mine.
 B: /*answer* The birds in section 1 do not have red-orange feet like the birds in section 2. Also, the feathers of the birds in Section 1 are light-colored vs. the dark-colored feathers of the birds in Section 2.
 A: /*agree*

We consider all other dialogues as passive. This includes cases where a proposal was immediately made and accepted, as well as where one partner didn’t engage with the proposals. 28.6% of the Birds dialogues were passive, compared with 61.5% for Text and 65.5% for Blocks. This again shows an influence of the task; presumably, Text was considered too easy to warrant discussion, while Blocks may have been seen as too hard, with participants giving up (as also reflected in the accuracy on that task).

To unveil the reasoning steps of the collaborative thinking process we identified and quantified typical patterns in the active discussions. (2-a) shows an example of *Proposal Extension*, where a proposal made by A is implicitly accepted and extended; and of *Counter Proposal*, where a proposal is implicitly rejected and replaced with a counter proposal. (There were also explicit acceptances and rejections, w/o proposals.)

Tasks	Patterns in active dialogues(%)			
	Proposal-Extension	Counter-Proposal	Explicit Acceptance	Explicit Rejection
Birds	52	68	60	8
Text	80	40	60	0
Blocks	30	80	70	30

Table 3: Proportions of active dialogues in each task with different patterns.

- (2) a. *Proposal Extension*
 A: One obvious difference that I see from the birds in section 1 is that the birds have longer beaks. [Proposal]
 B: another thing I noticed is it looks like 2’s have softer feather colors [Proposal-Extension]
 b. *Counter Proposal*
 A: /*answer* section 1 birds all look gray feathered [Proposal]
 B: They all have yellow bodies and dark heads [Counter Proposal]

Table 3 shows the proportions of active dialogues in which these patterns were observed, by task type. Explicit rejections happened rarely but were never observed in the texts task (too simple task). For Birds, there seems to be a balance between proposal-extension, counter-proposal and acceptance (balanced discussion). There were fewer counter-proposals in texts task, for it being simpler. It also looks like there were more disagreements in Blocks due to the complexity of the task.

5 Conclusions

We have presented a setting for collecting annotations from pairs of interacting annotators. Our analysis indicates that this setting of an “agreement game”, where explicit proposals have to be explicitly agreed on, fosters dialogs between the annotators. These dialogues yield to more correct responses and provide explication of the reasoning steps behind an annotation decision. Hence, both of our hypotheses, that the collaboration yields to more accurate responses and can reveal, at least in parts, the underlying reasoning steps, are supported. In line with our third and final hypothesis, the presence of these reasoning steps shows that the setting moves explanation as a linguistic act in the focus. It does however appear to be important to tune the level of difficulty of the task: if it is too simple, discussions do not emerge; if it is too hard,

the incentives for crowd workers have to be properly set so as to engage them. Our set-up also illustrates that natural categories could bring in more balanced discussions as well as better quality answers. Overall, it could provide useful data for developing a system which provides justifications.

Acknowledgements

We gratefully acknowledge help with setting up the experiment from our Bielefeld student research assistant Ayten Tüfekci.

References

- Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. 2016. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*.
- Finale Doshi-Velez and Been Kim. 2017. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.
- Leilani H. Gilpin, David Bau, Ben Z. Yuan, Ayesha Bajwa, Michael Specter, and Lalana Kagal. 2018. Explaining explanations: An approach to evaluating interpretability of machine learning. *CoRR*, abs/1806.00069.
- Been Kim, Martin Wattenberg, Justin Gilmer, Carrie J. Cai, James Wexler, Fernanda B. Viégas, and Rory Sayres. 2018. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (TCAV). In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*.
- Benjamin Letham, Cynthia Rudin, Tyler H. McCormick, and David Madigan. 2015. Interpretable classifiers using rules and bayesian analysis: Building a better stroke prediction model. *The Annals of Applied Statistics*.
- Stephen C. Levinson. 1983. *Pragmatics*. Cambridge Textbooks in Linguistics. Cambridge University Press.
- Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30*.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "why should i trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, New York, NY, USA.
- Emanuel A. Schegloff. 2007. *Sequence Organization in Interaction*. Cambridge University Press, Cambridge, UK.
- David Schlangen, Tim Diekmann, Nikolai Ilinykh, and Sina Zarriß. 2018. slurk - a lightweight interaction server for dialogue experiments and data collection. In *Short Paper Proceedings of the 22nd Workshop on the Semantics and Pragmatics of Dialogue (AixDial / semdial)*, Aix-en-Provence, France.
- Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. 2011. The caltech-ucsd birds-200-2011 dataset. Technical Report CNS-TR-2011-001, California Institute of Technology.
- Hao Yuan, Yongjun Chen, Xia Hu, and Shuiwang Ji. 2019. Interpreting deep models for text analysis via optimization and regularization methods. *AAAI Conference on Artificial Intelligence*.
- Quanshi Zhang, Yu Yang, Haotian Ma, and Ying Nian Wu. 2019. Interpreting cnns via decision trees. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

A

The Annotation Game Interface

Once two workers were presented they entered the task room as shown in Figure 4.

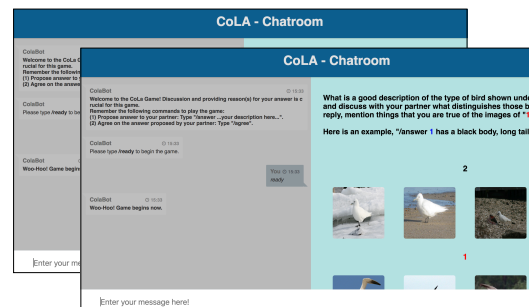


Figure 4: As soon as there are two participants in the waiting room, they are moved to the game room. Both participants see the same content on their screen. The content includes question Q , information I and instructions by the game bot who is also present in the game room.

This setup technically realises the setting described formally in Section 2 above, where annotators (“participants”) can work together to jointly formulate an answer A to the question Q they are given.

Prediction of User Emotion and Dialogue Success Using Audio Spectrograms and Convolutional Neural Networks

Athanasios Lykartsis

Audio Communication Group
TU Berlin
Germany

athanasios.lykartsis@tu-berlin.de

Margarita Kotti

Speech Technology Group
Toshiba Research Cambridge
United Kingdom

margarita.kotti@crl.toshiba.co.uk

Abstract

In this paper we aim to predict dialogue success and user satisfaction as well as emotion on a turn level. To achieve this, we investigate the use of spectrogram representations, extracted from audio files, in combination with several types of convolutional neural networks. The experiments were performed on the Let's Go V2 database, comprising 5065 audio files and having labels for subjective and objective dialogue turn success, as well as the emotional state of the user. Results show that by using only audio, it is possible to predict turn success with very high accuracy for all three labels (90%). The best performing input representation were 1s long mel-spectrograms in combination with a CNN with a bottleneck architecture. The resulting system has the potential to be used real-time. Our results significantly surpass the state of the art for dialogue success prediction based only on audio.

1 Introduction

Spoken Statistical Dialogue Systems (SDS) have gained much popularity in the last years, especially due to the widespread need for applications such as assisted living (Portet et al., 2013), phone banking (AbuShawar and Atwell, 2016), intelligent virtual agents (Matsuyama et al., 2016) and health care (Korpusik and Glass, 2017).

An important part of an SDS is spoken language, which is used to communicate directly with the virtual agent in order to pose questions and reply to the agent output. In a modular spoken SDS system, the speech part is converted to text through Automatic Speech Recognition Systems (ASR), which is then analysed using Natural Language Processing (NLP) methods. However, the audio part, which could be of low audio quality, is usually then discarded while the extracted text is fed forward to the SDS. In our view (and this is

an important part of our motivation), when looking at dialogue success prediction, this can be seen as a waste of possible resources, since the speech part can contain useful information regarding the emotional state of the user, or verbal cues which can indicate if the user is satisfied with the system performance. The prediction or recognition of such cues can be very helpful for supporting a dialogue management system, which can make better assessments as to what the next steps should be. Taking this thought one step further, we want to assess if it is possible to predict dialogue success based only on the audio, in order to find a light-weight, real-time method to manage the user expectations and, eventually, to build more efficient and user-friendly spoken SDS. A final motivation of this work is that we wanted to experiment with spectrogram input representations and convolutional neural networks (CNNs) as classifiers. Although there have been several examples of such uses for other topics, especially in image processing (Krizhevsky et al., 2012) and music information retrieval (Schlüter and Böck, 2014; Schreiber and Müller, 2018), this approach remains under-represented in the area of dialogue success prediction. Therefore, our research closes this gap and attempts to evaluate how well such approaches can function for dialogue success prediction.

Considering related works, the use of neural networks in the wider area of modular SDS has been gaining some popularity the last years. For example, neural networks have been utilised for dialogue state tracking. Korpusik and Glass (2018) use CNNs in order to track the user's goal over the whole dialogue without the use of hand-crafted semantic dictionaries and achieve high accuracy for their task. Henderson et al. (2014) similarly employ recurrent neural networks to map the results of ASR directly to a dialogue state and also report high performance. Another approach

(Zhao and Eskenazi, 2016) uses deep reinforcement learning to discover dialogue states and outperform a standard baseline. An additional deep reinforcement learning approach (Su et al., 2015) shows that using both RNNs and CNNs with turn level-features (non-audio) can be useful in predicting dialogue success. Research from Wen et al. (2016) shows that deep learning can be useful in creating more natural conversation task-oriented SDS, whereas Kim et al. (2016) use CNNs and RNNs for dialogue topic tracking.

As the listing of the related previous work shows, the use of neural networks with audio spectrograms or waveforms for the analysis of the audio part of the SDS and its consequent use for tasks such as dialogue success prediction has not been researched adequately. Only a limited number of papers (Papangelis et al., 2017; Kotti et al., 2017; Lykartsis et al., 2018) exist which explore the possibility of dialogue success prediction using audio features extracted from speech paired with standard machine learning techniques such as support vector machines.

These approaches have shown promising results, especially for creating a way to reliably estimate the user satisfaction. Additionally, they are able to do so in real-time or near-real-time and subsequently enable suitable next steps for the dialogue policy. Moreover, the recent success of deep learning approaches for audio tasks suggests that using these can bring an advantage: By exploiting input representations such as spectrograms, the estimation of task success can take place even at an ever finer time resolution level (e.g., very short audio frames), providing the possibility for even faster processing and reaction. Furthermore, data augmentation methods can provide a possibility to achieve higher accuracy rates.

Since CNNs combined with audio spectrograms as input have been shown to provide very good results in a multitude of tasks (for example for tempo estimation (Schreiber and Müller, 2018) and beat tracking (Schlüter and Böck, 2014)), we choose to employ them for the creation of an experimental setup for dialogue success prediction. In that sense, we frame our task as an emotion recognition one: As dialogue success is expected to show a high correlation with user satisfaction, which in turn is closely related with the user's emotional state, we investigated similar works using neural networks for speech emotion recognition.

Such works include those of Tzirakis et al. (2017) who use a Long-Short-Term-Memory (LSTM) network on top of a CNN in order to extract information and consider contextual information from raw audio data (waveforms), outperforming existing systems for speech emotion recognition. Similar work has been performed by Trigeorgis et al. (2016), where audio waveforms are used in combination with a CNN followed by an LSTM for speech emotion recognition, achieving high results for arousal and valence. In the work of Lee et al. (2017), a CNN is used to predict emotions based on speech spectrograms for a virtual elderly companion agent with very good results. Gu et al. (2018) create a multimodal framework with text and speech for emotion recognition. For the audio part, besides hand-crafted features, spectrograms with CNNs and LSTMs are used and fused with text features to predict 5 emotions and achieve better results than all other methods. Another interesting method comes from Yenigalla et al. (2018), where spectrograms of different sizes are used as an input for a CNN, achieving very good results for 4 emotional states. Neumann and Vu (2017) study the impacts of input features, signal length and speech type, using spectrogram or raw waveform input and CNNs, achieving state of the art results and reaching very useful conclusions for speech emotion recognition: input representation is not as important as the model architecture, which in turn is task and speech type specific. Fayek et al. (2015) also achieve very good results in speech emotion recognition using a simple deep neural network and spectrograms as input. A similar strategy is employed by Wang and Tashev (2017) for successful prediction of emotion, as well as gender and age on an utterance level, showing that even simple deep architectures can provide good results for speech emotion recognition. CNNs have also been used with success for general audio classification (Lee et al., 2009), which is a broader task, hinting at the suitability of this architecture for the task at hand in this paper. For this paper, we decided - for the sake of simplicity and due to the not enormous size of the dataset - to resort to only CNNs and determine which architectures, input representation forms and parameters provide good classification results for this task. Another reason for the use of CNNs is not only their aforementioned success in many tasks, but also the possibility to establish a

better understanding of the suitability of this approach for the task of dialogue success prediction. The latter is slightly different than speech emotion recognition per se because the user's emotional state is not the only factor that affects the final success label. Finally, this way we can establish a very fast and simple pipeline, which can also be used in a real-time setting to provide useful auxiliary information about the dialogue success, so as to inform the dialogue manager. This approach is compared to a baseline, involving hand-crafted audio features as in (Lykartsis et al., 2018), which have been shown to provide satisfactory results. Experiments are performed on the publicly available Let's Go V2 Database (Schmitt et al., 2012), which contains three kind of labels (for objective and subjective dialogue success and the emotional state of the user, for more information see 2.3).

This paper is structured as follows: In the next section the used methods are presented in depth, whereas in section 3, the results of the classification are shown and discussed. We close with conclusions and suggestions for future work.

2 Methods

2.1 Input Features

The input features chosen to be used for the CNN classifier in our case were mel-spectrograms (which can be seen as images summarizing the frequency content of a turn over time), extracted with the librosa python library (McFee et al., 2015). Mel-spectrograms have been used in a multitude of tasks for music information retrieval (Schlüter and Böck, 2014; Lidy and Schindler, 2016; Choi et al., 2017), as they are relatively simpler to calculate (in contrast to other transforms), while also providing a connection to human auditory perception through the use of the mel-scaling. Therefore, we reasoned that they could be a good basis for the task of dialogue success and speech emotion recognition. For an 1s long audio file we acquired a resulting 32 bins x 16 frames array (using the default librosa settings for spectrogram extraction that is a frame size of 92ms, a hann window and an overlap of 75% between consecutive frames). These settings are fairly standard for audio processing, as they allow a good temporal resolution but also a fair enough frequency resolution. We used this window size could mean that the speech segment is not necessarily stationary, but since we are looking for larger struc-

ture in the spectrogram (probably spanning several frames), this should not constitute a problem for the further processing (as it was also shown by our results). Using a shorter time window might produce even more temporally accurate spectrograms, but it would also require more computational resources. After conducting preliminary experiments with a window of 46ms, we could see that results were not improved, while at the same time requiring much more computational power for the spectrogram extraction. Therefore, we retained the window size of 92ms for all the further experiments. We also experimented with a length of 2s in order to see if longer (in the time domain) spectrograms would give better results - which can be seen as a trade-off between speed of processing (and therefore a close to real-time behavior of the classification/prediction system) and the accuracy of the prediction itself. This resulted to a 32 bins x 32 frames input array. We did not experiment with longer files, since most files in the Let's Go V2 database are not much longer than 2s (the average user turn duration is 1.5s with a standard deviation of 1.9s (Schmitt et al., 2012)). If the file is shorter than the selected analysis length, it is zero-padded at its end. All the files were of 8000 kHz sampling rate, no further preprocessing was performed, leading to a very lightweight pipeline, which is very close to a real-time processing. The goal of using these input features was to determine if a short spectrogram could suffice for providing good classification results.

2.2 Neural Nets/Classifiers

As mentioned in Section 1, we employ CNNs in this paper. The theory and inspiration for using CNNs can be found in Section 1. Specifically, we utilized Keras, which is based on the tensorflow library in python (Abadi et al., 2016). Keras has many advantages, such as that it is very effective, allowing for fast prototyping and training, even just by using CPUs (instead of GPUs). Inspired by similar experiments in other areas, we wished to test two different types of architectures:

- A standard **bottleneck architecture**, with 4 convolutional layers with 2-by-2 rectangular filters and a decreasing number of nodes (100-75-50-25), 2-by-2 max pooling and all activation functions being ReLU. This was followed by a batch normalization and 2 fully connected (FCN) layers (also with a decreas-




Figure 1: Bottleneck architecture flowchart diagram. For the details of the CNN, see the detailed architecture description in section 2.2

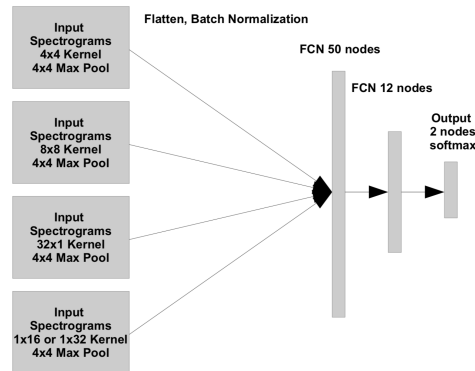


Figure 2: Parallel architecture flowchart diagram. For the details of the CNN, see the detailed architecture description in section 2.2

ing number of nodes, namely 50 and 12 and a dropout of 50%) and an output layer with softmax activation. The stride is always one, padding is always set at “same”, so as that the output has the same length as the original input. As an optimizer, ADAM was used with a learning rate of 0.001, whereas a categorical cross entropy was utilized as the loss function. For this architecture we were inspired from (Tzirakis et al., 2017; Trigeorgis et al., 2016). The above architecture is depicted in Figure 1.

- A **parallel** CNN architecture: In this case, 4 input layers with 32 nodes and with different kernel sizes (a quadratic 4-by-4 kernel, a quadratic 8-by-8 kernel, a 1-dimensional 32-by-1 filter (for the mel-spectrogram frequency bins) and a 1-dimensional 1-by-16 or 1-by-32 filter (for the time frames, corresponding to the file length of 1s or 2s, respectively)) are processed in parallel and their output is combined (concatenated and flattened). In this case, the max-pooling is done in a 4-by-4 manner and the activations are also all ReLU. The combined output of the four parallel layers is batch-normalized fed

into 2 FCN layers with 50 and 12 nodes with a dropout of 0.5 between them, followed by a 2-node output softmax layer. Same as before, the stride is always one and padding is set at “same”. Also in this case, an ADAM optimizer was used (with a learning rate of 0.001), and a categorical cross entropy as a loss function. For this architecture we were inspired from the implementation in (Yenigalla et al., 2018) (using parallel layers) and from the one in (Schreiber and Müller, 2018), using one-dimensional filters. Our reasoning was that combining these two features, a powerful network could be constructed which would be able to learn features pertaining to emotional states of the user, as well as more specific signal features inherent in the spectrogram (such as the tempo of the utterance). The above architecture is depicted in Figure 2.

Finally, we implemented a baseline following the scheme in (Lykartsis et al., 2018), comprising 5 hand-crafted spectral and rhythmic features (the standard deviation of the three MFCCs and the tempo and mean of the RMS-based beat histograms) and featuring an SVM classifier with

$C = 2$, $\gamma = 1/N_{features}$ and an RBF kernel using the scikit learn python module. These parameters were kept the same as in the aforementioned publication, since they resulted via a grid search there for a dataset of similar audio quality, and for ensuring comparability between the studies.

The whole pipeline was developed and tested using python 3.6 on a Windows 7 OS with 8 GB of RAM and an Intel i5 quad core processor running at 3.2 GHz. Using this, extracting the spectrogram of a turn with librosa is achieved in under 1s, whereas training of the model for one epoch takes around 5s. After the model has been created, prediction, that is running the validation spectrogram via the trained CNN, is taking 0.1-0.2s (depending on the length of the turn), resulting into a near-real-time system. We refrained from using a development set due to the small size of the dataset and because our averaged results over the 3 validation folds should provide sufficient validity.

2.3 Dataset

The spoken dialogue corpus used in this study is based on the the CMU Let’s Go Bus Information System (Schmitt et al., 2012) (from this point on referred to as the *Let’s Go V2 dataset*). This has been developed by the university of Ulm in order to evaluate dialogue quality, user emotion and task success for an SDS which was used as an information system for bus itinerary search. The database contains 9083 system-user exchanges (to which we will refer as *interactions* in the following). For our experiments, we kept a total of 5065 audio files for the interactions, for which all labels where available, so as to be able to compare between the results using the different label sets.

Each interaction has been rated with three labels. The first is an **emotional label**, signifying the emotional state of the user. The label has four levels, ranging from non-angry to very angry. This label was assigned from the users themselves. Another label shows the **subjective** dialogue success, dubbed IQ (Interaction Quality) in the corpus annotation (Schmitt et al., 2012), indicating whether the user was satisfied with the interaction. This label ranges from satisfied to extremely unsatisfied and has five levels and was agreed on by three individual external raters. We refer to it here as *subjective label*. Finally, the **objective labels** indicate whether the goal of the dialogue was reached, i.e., the information looked for was actually provided

by the system. This label also exists on an interaction level and has two levels (successful or not).

In order to simplify the classification, we choose to create a binary model which results from taking the most highly ranked result of each label set as the positive label, and all the other results pooled together as the negative label. In that way, it was possible to create an almost balanced dataset for the subjective labels (53% negative and 47% negative ones), but not for the other two labels sets (having correspondingly a distribution of 65% positive/35% negative for the emotional labels and 85% positive/15% negative for the objective samples). Therefore, we then created a balanced version of the dataset for the emotional and the objective labels by taking the smaller class and randomly choosing as many examples for the other class. The balanced subjective set contained 5065 samples, the balanced objective one 1146 and the balanced emotional one 3660 samples.

3 Results and Discussion

The results of the classification for all 3 labels can be seen in Tables 1 and 2 for the training and the validation set, respectively. The respective results for the baseline system can be seen in Table 3. The results reported here are the average accuracy over the three folds, followed by the loss of the network. The standard deviation of the accuracy over the folds is not reported, since it ranges from 0.5% to 1.5%, and can therefore be considered negligible, showing that the system is robust. It must be mentioned here again, that the basic unit of classification was the audio of the user turn, for which the labels are also available. The accuracy reported refers to the amount of correctly predicted labels for the user turns as a ratio of all turn classifications.

Concerning the effect of different parameters for the CNNs, the best parameter set was determined by 3 fold cross-validated grid search. The aforementioned cross-validation lead to the results reported in tables 1 and 2. We experimented with several values for the learning rate, the optimizer and the batch size. We observed an effect for better results with a learning rate of 0.001, a batch size of 8 and by using the ADAM optimizer. Finally, the results shown here were the result of 500 epochs long training procedure. We did not observe any improvement when training for longer time, and this is definitely an amount of training time which

Setting	1s		2s	
	Accuracy	Loss	Accuracy	Loss
Bottleneck Architecture, subjective labels	0.95	0.12	0.97	0.07
Bottleneck Architecture, objective labels	0.97	0.07	0.98	0.05
Bottleneck Architecture, emotional labels	0.98	0.06	0.98	0.05
Parallel Architecture, subjective labels	0.81	0.31	0.97	0.07
Parallel Architecture, objective labels	0.91	0.24	0.97	0.1
Parallel Architecture, emotional labels	0.92	0.17	0.97	0.07

Table 1: Classification results, training set, average accuracy over 3 folds and corresponding loss for 1 and 2 s segments. All datasets are balanced, the prior is 0.5.

Setting	1s		2s	
	Accuracy	Loss	Accuracy	Loss
Bottleneck Architecture, subjective labels	0.78	0.57	0.9	0.3
Bottleneck Architecture, objective labels	0.9	0.48	0.86	0.5
Bottleneck Architecture, emotional labels	0.9	0.33	0.86	0.5
Parallel Architecture, subjective labels	0.7	0.93	0.7	0.93
Parallel Architecture, objective labels	0.88	0.46	0.88	0.78
Parallel Architecture, emotional labels	0.82	0.69	0.74	1.25

Table 2: Classification results, validation set, average accuracy over 3 folds and corresponding loss for 1 and 2 s segments. All datasets are balanced, the prior is 0.5.

Setting	Whole turn
Baseline (SVM), subjective labels	0.59
Baseline (SVM), objective labels	0.57
Baseline (SVM), emotional labels	0.75

Table 3: Classification results, baseline system, average accuracy over 3 folds. Features are extracted over the whole turn and aggregated. All datasets are balanced, the prior is 0.5.

is very manageable on reasonably strong computation systems (see 2.2). With regards to the effect of the spectrograms’ length, this did not seem to have a large effect on classification accuracy. In general, results were somewhat better for then 1s case. We therefore assume that in the case of less data, the length of the segment can be kept to a minimum value. These findings corroborate the results from Neumann and Vu (2017), which mention that the NN architecture is more important than the input representation form, at least in the context of speech emotion recognition.

Comparing the two architectures, the first architecture with the sequential layers has shown slightly better results. This might be due to the parallel models lacking the information to extract useful patterns, probably due to possible data deprivation. In total, the results are much higher than the ones produced from the baseline. We observed some important trends (with regards to the validation set results). The first architecture using the bottleneck structure has proven to be useful for all labels. This might be due to phonetic features in the spectrogram indicating task success being very

concrete (such as “thank you”, or the user’s voice melody sinking) and therefore rendering a simpler structure to extract the features more suitable. Between the different label types, the emotional and objective label sets show somewhat better results when using smaller lengths, showing that for the subjective labels, a greater length is essential for the CNN to extracting more relevant information. The parallel layer architecture has shown to be useful for the objective labels. This is probably due to the higher complexity of predicting an objective task success from purely sound data. Additionally, the turn length does not seem to play an important role, which might mean that for more complex architectures, less information length can be sufficient to achieve good accuracy. All in all, the parallel architecture was somewhat less performant than the bottleneck one, which shows that for these data, simpler structures are more useful.

In general, the results were very positive and surpass results on similar datasets which are state-of-the-art: The maximum accuracy on the validation set, for the subjective labels, achieved using only sound files was 90%, which surpasses

the best results in (Lykartsis et al., 2018) by 16%. However, it must be noted, that the datasets used are slightly different, in the sense that the task is a different one (finding the right laptop vs. finding the right itinerary while interacting with an SDS). Also, in (Lykartsis et al., 2018), both the subjective and the objective labels were provided by the user, but in the Let’s Go V2 system, those were provided by external raters, as well as having a different resolution (labeled turns instead of full dialogues). Therefore, the results not directly comparable, but the research question is the same. Furthermore, the length and audio quality of the recordings is very similar, so that it can be claimed that using mel-spectrograms as input and CNNs as classifiers provides a successful and computationally not too intensive way to achieve emotion detection and dialogue success prediction only from audio. We are therefore optimistic, that with more training data, we could build sounder models which can generalize better and build on the tendencies observed here, achieving even better results. Finally, the results achieved in (Lykartsis et al., 2018), that smaller files are more suitable for higher accuracy, are also observable here.

In comparison to other studies which used the Let’s Go V2 dataset, two works have been found in the literature, that of Schmitt et al. (2011) and that of Stoyanchev et al. (2019), both of which resort to linguistic features, among others. In (Schmitt et al., 2011), the best achieved result was 61.6% unweighted average recall for predicting quality of interaction (i.e., the subjective label as mentioned in this paper) using a multitude of automatically extracted hand-crafted features (linguistic and dialogue state ones) and support vector machines. Our baseline system achieves a close result (59% average accuracy). Also, by using ASR and linguistic features alone in combination with support vector machines, Stoyanchev et al. (2019) manage to achieve 50% unweighted average recall. It must be mentioned, that a direct comparison is not possible due to the different nature of the features and the different categories (both papers mentioned here predicted 5 categories of interaction quality), however we can see that our system can predict dialogue success with very high performance. Another interesting observation in that context is the fact that although in our study the best results were achieved with the objective label set, in (Lykartsis et al., 2018), the better results

were achieved with the subjective labels, which in our case provide the least good results - but still better than the baseline. This might be a consequence of a different definition of what constitutes subjective success between the two datasets: For the laptop dataset of (Lykartsis et al., 2018), subjective success means that the users found all the information they were looking for (when asked at the end of the dialogue), whereas for the Let’s Go V2 system, subjective success meant that external raters were judging the interaction to be successful or not, probably leading to different label distributions. The different results might also be a consequence of the different tasks involved.

4 Conclusion

In this paper, we have shown that classification of user emotion, and prediction of objective and subjective task success of a spoken SDS using only audio in the form of spectrograms is not only possible, but also can be achieved to a high standard using CNNs with small computational effort, resulting in an almost real-time system. Our results greatly surpassed those of other similar studies and can be used to train models which can - on a turn level, i.e., with audio information of limited duration - predict the direction a dialogue takes and can therefore act to change the dialogue course.

We are optimistic that if our features are combined with other non-sound features (such as linguistic features), we will have the possibility to raise classification accuracy even more. However, this falls outside the aim of the current study and will be part of our future work. Furthermore, a possibility would be to perform system fusion at the classifier level, combining for example different CNN architectures (like the ones shown in this paper) and other classifiers with hand-crafted features, as in the approach from (Lykartsis et al., 2018). Such a system could benefit from the multiple different input representation and could potentially provide very good results, as in (Gu et al., 2018). As additional future work, we plan to conduct experiments with more architectures and parameters, and also employ other neural network classifiers such as Temporal Convolutional Networks (TCNs), which combine the merits of both CNNs and RNNs/LSTMs. Finally, we will also experiment with data preprocessing methods, such as denoising and data augmentation methods such as transformations in time and frequency.

References

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, pages 265–283.
- Bayan AbuShawar and Eric Atwell. 2016. Usefulness, localizability, humanness, and language-benefit: additional evaluation criteria for natural language dialogue systems. *International Journal of Speech Technology*, 19(2):373–383.
- Keunwoo Choi, György Fazekas, Mark Sandler, and Kyunghyun Cho. 2017. Transfer learning for music classification and regression tasks. *arXiv preprint arXiv:1703.09179*.
- Haytham M Fayek, Margaret Lech, and Lawrence Cavedon. 2015. Towards real-time speech emotion recognition using deep neural networks. In *2015 9th international conference on signal processing and communication systems (ICSPCS)*, pages 1–5. IEEE.
- Yue Gu, Shuhong Chen, and Ivan Marsic. 2018. Deep mul timodal learning for emotion recognition in spoken language. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5079–5083. IEEE.
- Matthew Henderson, Blaise Thomson, and Steve Young. 2014. Word-based dialog state tracking with recurrent neural networks. In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 292–299.
- Seokhwan Kim, Rafael Banchs, and Haizhou Li. 2016. Exploring convolutional and recurrent neural networks in sequential labelling for dialogue topic tracking. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 963–973.
- Mandy Korpusik and James Glass. 2017. Spoken language understanding for a nutrition dialogue system. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 25(7):1450–1461.
- Mandy Korpusik and James Glass. 2018. Convolutional neural networks for dialogue state tracking without pre-trained word vectors or semantic dictionaries. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 884–891. IEEE.
- Margarita Kotti, Alexandros Papangelis, and Yannis Stylianou. 2017. Will this dialogue be unsuccessful? prediction using audio features. In *Proceedings of the ACM SIGIR International Conference on Theory of Information Retrieval (ICTIR)*.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105.
- Honglak Lee, Peter Pham, Yan Lalgman, and Andrew Y Ng. 2009. Unsupervised feature learning for audio classification using convolutional deep belief networks. In *Advances in neural information processing systems*, pages 1096–1104.
- Ming Che Lee, Sheng Cheng Yeh, Sheng Yu Chiu, and Jia Wei Chang. 2017. A deep convolutional neural network based virtual elderly companion agent. In *Proceedings of the 8th ACM on Multimedia Systems Conference*, pages 235–238. ACM.
- Thomas Lidy and Alexander Schindler. 2016. Parallel convolutional neural networks for music genre and mood classification. *MIREX2016*.
- Athanasios Lykartsis, M Kotti, A Papangelis, and Y Stylianou. 2018. Prediction of dialogue success with spectral and rhythm acoustic features using dnns and svms. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 838–845. IEEE.
- Yoichi Matsuyama, Arjun Bhardwaj, Ran Zhao, Oscar Romeo, Sushma Akoju, and Justine Cassell. 2016. Socially-aware animated intelligent personal assistant agent. In *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 224–227.
- Brian McFee, Colin Raffel, Dawen Liang, Daniel PW Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto. 2015. librosa: Audio and music signal analysis in python. In *Proceedings of the 14th Python in Science conference (SciPy)*, pages 18–25.
- Michael Neumann and Ngoc Thang Vu. 2017. Attentive convolutional neural network based speech emotion recognition: A study on the impact of input features, signal length, and acted speech. *arXiv preprint arXiv:1706.00612*.
- Alexandros Papangelis, Margarita Kotti, and Yannis Stylianou. 2017. Predicting dialogue success, naturalness, and length with acoustic features. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5010–5014. IEEE.
- François Portet, Michel Vacher, Caroline Golanski, Camille Roux, and Brigitte Meillon. 2013. Design and evaluation of a smart home voice interface for the elderly: acceptability and objection aspects. *Personal and Ubiquitous Computing*, 17(1):127–144.
- Jan Schlüter and Sebastian Böck. 2014. Improved musical onset detection with convolutional neural networks. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6979–6983. IEEE.

- Alexander Schmitt, Benjamin Schatz, and Wolfgang Minker. 2011. Modeling and predicting quality in spoken human-computer interaction. In *Proceedings of the SIGDIAL 2011 Conference*, pages 173–184. Association for Computational Linguistics.
- Alexander Schmitt, Stefan Ultes, and Wolfgang Minker. 2012. A parameterized and annotated spoken dialog corpus of the cmu lets go bus information system. In *LREC*, pages 3369–3373.
- Hendrik Schreiber and M Müller. 2018. A single-step approach to musical tempo estimation using a convolutional neural network. In *Proceedings of the 19th International Society for Music Information Retrieval Conference (ISMIR), Paris, France*.
- Svetlana Stoyanchev, Soumi Maiti, and Srinivas Bangalore. 2019. Predicting interaction quality in customer service dialogs. In *Advanced Social Interaction with Agents*, pages 149–159. Springer.
- Pei-Hao Su, David Vandyke, Milica Gasic, Dongho Kim, Nikola Mrksic, Tsung-Hsien Wen, and Steve Young. 2015. Learning from real users: Rating dialogue success with neural networks for reinforcement learning in spoken dialogue systems. In *Sixteenth Annual Conference of the International Speech Communication Association (INTER-SPEECH)*.
- George Trigeorgis, Fabien Ringeval, Raymond Brueckner, Erik Marchi, Mihalis A Nicolaou, Björn Schuller, and Stefanos Zafeiriou. 2016. Adieu features? end-to-end speech emotion recognition using a deep convolutional recurrent network. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5200–5204. IEEE.
- Panagiotis Tzirakis, George Trigeorgis, Mihalis A Nicolaou, Björn W Schuller, and Stefanos Zafeiriou. 2017. End-to-end multimodal emotion recognition using deep neural networks. *IEEE Journal of Selected Topics in Signal Processing*, 11(8):1301–1309.
- Zhong-Qiu Wang and Ivan Tashev. 2017. Learning utterance-level representations for speech emotion and age/gender recognition using deep neural networks. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5150–5154. IEEE.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrksic, Milica Gasic, Lina M Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2016. A network-based end-to-end trainable task-oriented dialogue system. *arXiv preprint arXiv:1604.04562*.
- Promod Yenigalla, Abhay Kumar, Suraj Tripathi, Chirag Singh, Sibsambhu Kar, and Jithendra Vepa. 2018. Speech emotion recognition using spectrogram & phoneme embedding. *Nineteenth Annual Conference of the International Speech Communication Association (INTER-SPEECH)*, pages 3688–3692.
- Tiancheng Zhao and Maxine Eskenazi. 2016. Towards end-to-end learning for dialog state tracking and management using deep reinforcement learning. In *17th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, page 1.

Modelling Adaptive Presentations in Human-Robot Interaction using Behaviour Trees

Nils Axelsson

Division of Speech, Music and Hearing
KTH Royal Institute of Technology
Stockholm, Sweden
nilsaxe@kth.se

Gabriel Skantze

Division of Speech, Music and Hearing
KTH Royal Institute of Technology
Stockholm, Sweden
skantze@kth.se

Abstract

In dialogue, speakers continuously adapt their speech to accommodate the listener, based on the feedback they receive. In this paper, we explore the modelling of such behaviours in the context of a robot presenting a painting. A Behaviour Tree is used to organise the behaviour on different levels, and allow the robot to adapt its behaviour in real-time; the tree organises engagement, joint attention, turn-taking, feedback and incremental speech processing. An initial implementation of the model is presented, and the system is evaluated in a user study, where the adaptive robot presenter is compared to a non-adaptive version. The adaptive version is found to be more engaging by the users, although no effects are found on the retention of the presented material.

1 Introduction

Speakers in dialogue cannot just assume that their speech is received by the addressee and understood as intended. They have to continuously monitor the addressee to verify that the information is attended to, perceived, understood and accepted (Clark, 1996). By keeping close track of verbal and non-verbal feedback from the addressee, speakers can alter their presentation in order to accommodate the listener.

In this paper, we explore how this process can be modelled in spoken human-robot interaction. As a test-bed, we have designed a scenario where a robot is presenting visual information (such as a poster or a piece of art) to a human, as seen in Figure 1. This setting allows us to explore how the presentation can be adapted to the audience's level of attention, understanding and engagement.

Modelling adaptive presentation in a human-robot interaction scenario is non-trivial, as the robot needs to pick up feedback from different



Figure 1: The scenario chosen as a test-bed for the model: a robot presenting a painting to a human.

modalities, and continuously adapt its behaviour to accommodate the listener. It is also not obvious that such a system would be better in terms of teaching the presented material and user experience, compared to a fixed, non-adaptive presentation (such as audio-guides used in museums), as the robot is unlikely to exhibit the same level of adaptation as a human. This paper has two main contributions, which address these concerns. First, we explore the use of Behaviour Trees (Colledanchise and Ögren, 2018) for modelling the adaptive behaviour. Behaviour Trees, a specific formalism for decomposing a plan into a tree structure, have been applied extensively to video games and robotics (Hasegawa et al., 2017; Hu et al., 2015), and systems that break down an interaction or a dialogue to a tree are not new (Smith and Hipp, 1994; Boye, 2007; Bohus and Rudnicky, 2009). However, we are not aware of any previous attempts at applying specifically Behaviour Trees to real-time modelling of spoken interaction. Second, we present an experiment where we compare the adaptive robot presenter to a version where the presentation is statically executed, i.e., where the user's reactions are not taken into account.

2 Background

The scenario of a robot presenting information to an audience (one or several people), has been explored in earlier work (Jensen et al., 2005; Szafir and Mutlu, 2012; Ohya et al., 2006). However, these works have not focused on how the presentation can be adapted based on verbal and non-verbal feedback. Poster presentations between humans have been studied in order to analyse the gaze and backchannel behaviours of participants and presenters (Kawahara, 2012). Hashimoto et al. (2011) and Verner et al. (2016) have shown that more interactive robot teachers lead to better results in learning. Yousuf et al. (2012) and Eichner et al. (2007) show that users prefer presenting agents that adapt their grounding behaviour to their audience.

2.1 Grounding and Adaptation

According to Clark (1996) and Allwood et al. (1992), any coordinated action can be described as an action ladder, with each level requiring the cooperation of speaker and addressee. If the speaker A is presenting to the addressee B, then the levels of the action ladder, bottom-to-top, are **attention** (B must be paying attention to A’s presentation), **hearing** (B must hear the words said by A), **understanding** (B must understand the meaning behind the words said by A) and **acceptance** (B must accept, and optionally be interested in, the concept proposed by A’s presentation).

The addressee can give positive and negative evidence of each level (feedback), to signal completeness to the speaker. If negative evidence is signalled for a level, all levels above it have failed by extension. If positive evidence is signalled for a level, all levels below it have succeeded by extension. Feedback signals like these can then be used by the speaker to adapt the presentation – by explaining some information in more depth or by making the presentation more interesting – and thereby accommodate the listener. This process is referred to as *Grounding* by Clark (1996). It is not possible to give positive evidence in response to every piece of a conversation, but the important thing is to receive enough evidence to meet the *grounding criterion*, the requirements for evidence needed depending on how important the speakers deem the content of the presentation to be.

2.2 Behaviour Trees

A Behaviour Tree, or BT, is a tree structure that models a plan, initially proposed by Mateas and Stern (2002). Behaviour Trees have been used in video games (Isla, 2005, 2008; Hasegawa et al., 2017) and to model robot behaviours (Hu et al., 2015; Colledanchise et al., 2016). There is previous work applying BTs to virtual agents (Sun et al., 2012; Fujita et al., 2003), but to our knowledge, so far they have not been used to model conversational agents or social behaviour.

The leaves of the tree are the tasks that are executed. All non-leaves are control flow nodes. Execution flows from the root down the tree, starting when some external process *ticks* the root to start execution. Each node in the tree returns one of three values to its parent; SUCCESS or FAILURE if the task has finished with either result, or RUNNING if it has not finished.

The two most common control flow nodes are *Sequence* and *Selector* nodes. *Sequence* nodes run their children in order from left to right until a FAILURE or RUNNING is encountered, at which point the sequence returns that value. If all child nodes succeed, the sequence returns SUCCESS. *Selector* nodes run their children from left to right until a SUCCESS or RUNNING is encountered, returning that value, or FAILURE if all children fail (Colledanchise and Ögren, 2018).

3 Modelling the presentation

In this paper, we propose a Behaviour Tree to model the complex task of poster presentation while taking grounding and adaptation into account. The tree breaks down this complex task into smaller, *independent* tasks. As Section 4 describes, our initial implementations of these individual tasks are greatly simplified, as many of them are indeed challenging research problems in their own right. However, the decomposition into the behaviour tree allows us to start with simpler initial implementations of the individual tasks (some of which can be controlled through *Wizard of Oz*), and then gradually replace them with more complex models (e.g., through machine learning), without changing the structure of the tree, or the implementation of other tasks.

The abstract BT is shown in Figure 2. Whereas most traditional dialogue systems process the interaction utterance-by-utterance, the BT allows the system to process the interaction incremen-

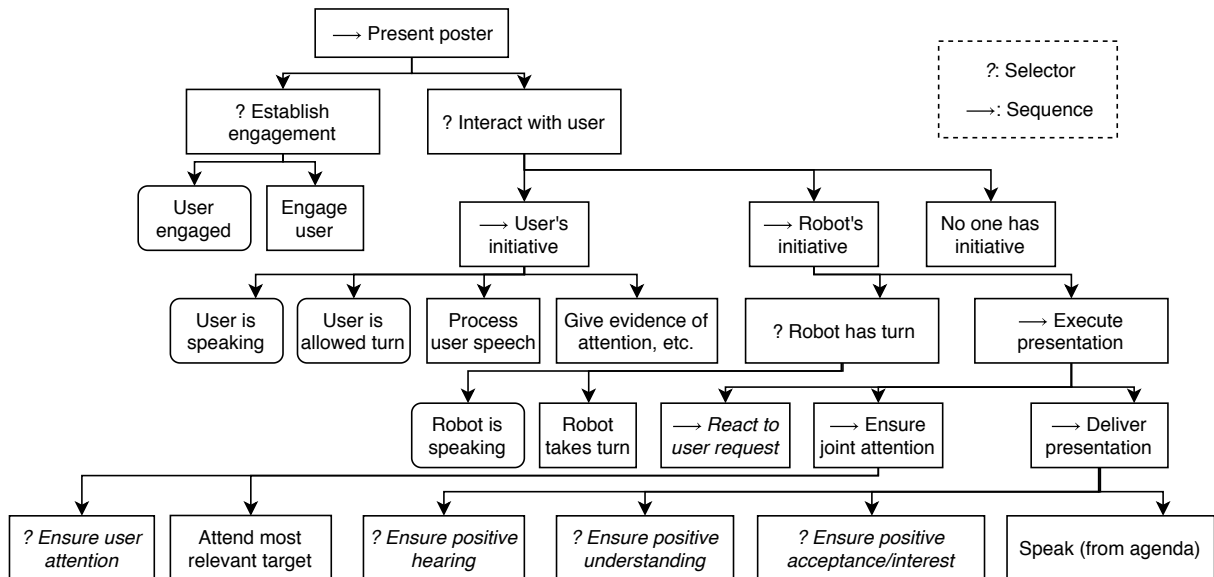


Figure 2: The Behaviour Tree developed as part of this project. Note that the children of any selector or sequence with an italic title are not shown to save room.

tally, in real time (in the vein of [Schlangen and Skantze, 2009](#)). Thus, the tree is designed to be executed on the time scale of 10 times per second. The root represents the entire task of presenting a poster. The tree contains both a sub-tree for finding and recruiting participants and presenting to them, and thus will never return SUCCESS; the presentation is either going on (RUNNING) or impossible (FAILURE). The deeper levels of the tree are discussed, top-to-bottom and left-to-right, below.

Dynamic information is not kept in the static tree; instead, it depends on external modules to keep track the joint action ladder (a *knowledge manager*), and where the agent is in its presentation (an *agenda*). These components are not discussed here, as they are less general than the tree.

The system needs to find a user to whom to present, which happens in the **Establish engagement** sub-tree at the top of the tree. After this tree has succeeded at inviting or engaging a user into the presentation, which can be a more or less complicated task ([Bohus and Horvitz, 2009, 2014](#)), the system presents its presentation through its **interact with user** sub-tree.

This sub-tree handles turn-taking by offering the turn to the addressee if appropriate, which can be done in multiple ways ([Meena et al., 2014; Ström and Seneff, 2000](#)). As the tree runs at its rate of 10 Hz, the user’s utterance is processed incrementally, and the system can deploy backchan-

nels and gaze cues in response ([Morency et al., 2008](#)).

If the user does not have the turn, the robot either has or takes the turn through its **Robot’s initiative** sub-tree, and executes the presentation. Firstly, joint attention is ensured or grabbed (see [Yu et al. \(2015\)](#)) if lost, this can be sensed in multiple ways ([Ba and Odobez, 2009; Sheikhi, 2014; Szafir and Mutlu, 2012](#)).

If the system has the user’s attention, it ensures hearing, understanding, and acceptance, in order, according to the respective grounding criteria. As these sub-trees have had their chance to change the presentation agenda to address negative evidence of hearing, understanding and acceptance (see ([Vaufreydaz et al., 2016; Aly and Tapus, 2015; Sidner et al., 2006; Skantze et al., 2014](#)) for examples on how to measure these), the system then **speaks from the agenda**, driving the presentation forward. Only if the tree reaches this leaf without any previous leaf returning RUNNING does the system speak, resulting in incremental, adaptive speech synthesis in the vein of [Skantze and Hjalmarsson \(2010\); Buschmeier et al. \(2012\); Kopp et al. \(2014\)](#).

4 Implementation

We developed an initial implementation of a system containing the Behaviour Tree model proposed in Section 3 as an extension to the *IrisTK* dialogue framework ([Skantze and Al Moubayed,](#)

2012). The *Furhat* robot head (shown in Figure 1) served as the robot platform (Al Moubayed et al., 2012).

The *agenda* of the implemented system tracked entire lines of the presentation’s script. To adapt the presentation, evidence of understanding was thus tracked on a line-by-line basis, and the system could explain a line for which understanding had not been shown, by finding other lines that explained the misunderstood line.

The system modelled attention by treating users as attentive if they were looking at the system or the poster, using their head pose (estimated via *Kinect*) as a proxy of gaze direction. Upon inattention, the system would restart its current utterance, similar to the stop-and-restart method employed by Yousuf et al. (2012). A *Wizard of Oz* setup was used to tag positive and negative evidence of hearing, attention and acceptance.

5 Experiment

To evaluate the system and tree, we set up an experiment where the system described in Section 4 had two modes: in the **adaptive mode**, the system fully used its adaptive behaviour. In the **non-adaptive mode**, the system always assumed positive feedback on all four levels of the joint action ladder. The non-adaptive system also never yielded the turn to the user. The non-adaptive mode presented the same surface-level five-minute presentation every time, so a five-minute time limit was also set for the adaptive mode, which would end its presentation after that time. The agent’s gaze behaviour was the same in both modes, shifting between the participant’s head and the poster.

We used a within-subject experimental design, where each subject interacted with the two versions of the system. Two posters with 16th-century paintings were created: Gentile Bellini’s *Miracle of the Cross fallen into the channel of Saint Lawrence* (*Croce*, for short), and *Great Tower of Babel*, by Pieter Bruegel the Elder. The orders of the two paintings and modes were both counterbalanced between subjects.

30 subjects participated in the experiment, 16 male and 14 female. A majority of participants were undergraduate university students. Participants were not told about the differences between the adaptive and non-adaptive modes, other than that only the adaptive mode could answer ques-

tions. Participants were otherwise encouraged to give active feedback to the agent regardless of condition (even though the non-adaptive version would actually ignore this feedback).

Conditions were evaluated immediately following the end of the respective presentation. Firstly, in order to evaluate retention of the information presented, participants were given an electronic form where they answered questions about the presentation and painting. Secondly, they were asked to fill in adapted versions of the *Godspeed* questionnaire by Bartneck et al. (2009), and the *Networked Minds social presence* questionnaire by Biocca and Harms (2011). Participants were rewarded with a cinema ticket.

6 Results

The results of 2 participants had to be excluded due to technical problems during the experiment, yielding 28 data points (16 male, 12 female), of which 14 indicated that they had previous experience with a social robot, two indicated that they had seen the *Croce* painting before, and eight indicated they had seen the *Babel* painting before.

The Wilcoxon paired signed-rank test (Wilcoxon, 1945) was used to compare the answers given in the *Social Presence* and *Godspeed* forms. The questions were grouped by categories in each test, and the answers to them were averaged. This compensated for the large number of questions.

Five out of ten categories (*anthropomorphism* ($p = .0342, \delta = 0.4 \pm 0.4$), *animacy* ($p = .00770, \delta = 0.63 \pm 0.46$), *perceived safety* ($p = .0128, \delta = 0.58 \pm 0.42$), *perceived emotional contagion* ($p = .000999, \delta = 0.47 \pm 0.22$), *perceived behavioural interdependence* ($p = 2.77 * 10^{-5}, \delta = 0.96 \pm 0.29$)) show statistically significant differences between the adaptive and non-adaptive modes, with the adaptive scoring higher. One additional category, *likeability* of the robot, shows a statistically significant difference ($p = .0148, \delta = 0.70 \pm 0.60$) between the first and the second presentation given to participants, the first scoring higher. No statistically significant differences were found between the two paintings.

For the analysis of the retention questionnaire, one additional subject had to be excluded due to technical problems. Eleven questions per poster were graded on a scale from zero to eleven based on correctness, normalising to only count ques-

tions that were possible to answer based on the presentation the user received. The answers in the *Babel* questionnaire ($M = 6.938$, $Mdn = 7.542$, $SD = 1.989$) were found to have a statistically significantly ($p = .04235$) different distribution than those in the *Croce* questionnaire ($M = 6.270$, $Mdn = 6.758$, $SD = 1.771$), but no statistically significant differences were found when comparing the adaptive mode and the non-adaptive mode ($p = .449$), or the first and second presentation participants received ($p = .990$).

7 Discussion

The results from the Social presence and God-speed questionnaires showed that the adaptive version was perceived to have a higher Animacy, Anthropomorphism, Safety, Emotional contagion, and Behavioural interdependence. These are all aspects that relate to higher interactivity, and are all associated with positive values, which indicates that an interactive presenter that takes the user's attention and understanding into account is indeed perceived to be more engaging. When asking the subjects about the difference between the two versions after the experiment, they typically had a hard time identifying the exact difference in terms of interactivity. This is interesting, as it indicates that they were not aware of the specific reason for why they preferred the adaptive version. The gaze behaviour of the robot, which followed users around even in the otherwise non-adaptive mode, may have led to the perception that the system was paying attention to the user even in this mode.

There was a somewhat unexpected difference between the first and second presentation, where the former had a somewhat higher Likeability of the robot, regardless of painting and mode. One potential explanation for this is that users were aware of the format of the evaluation the second time, and might have been more stressed about it.

However, no statistically significant differences were found in the user's retention of the two presentations. There was a large variation in how much the individual subjects remembered from each presentation. Certain participants remembered almost nothing of either presentation. Others were able to quote the robot on every question in both the adaptive and non-adaptive modes. This introduces noise and makes the comparison hard to perform, given the relatively small number of participants.

7.1 Future work

Although the agent developed in our initial implementation does adapt its presentation based on feedback from the user, this adaptation was mostly done on a semantic level (i.e., updating its agenda). In future studies, we will explore how the system could also adapt factors like turn length, speech rate, the frequency with which the agent would require evidence of understanding, and what the system would consider as evidence of understanding.

Classifying negative and positive evidence based on multi-modal signals is indeed a very challenging task, as these cues could be very subtle (e.g., facial expressions of boredom or interest). In this experiment, this classification was done by a human *Wizard of Oz*. The data collected through this experiment could potentially be used to train specific models for this, as they have already been partially annotated by the Wizard.

A natural extension of the model is to also allow several users to take part in the presentation. This would give rise to new challenges when it comes to determining who should be considered to be engaged in the presentation, and how to adapt the presentation, since the different users in the audience might show evidence of understanding to various degrees. Also, if a new user appears in the middle of the presentation, it is not clear how to proceed with the agenda.

8 Conclusions

This paper presents a first step towards a system that uses Behaviour Trees to create an adaptive presentation agent. Initial results show that users find a system that attempts to adapt its presentation to their reception of the presentation more positive along several dimensions. Our initial implementation of the proposed Behaviour Tree model is a promising first step towards a complex adaptive behaviour model for conversational interaction, where the complex task of making an adaptive presentation has been decomposed into smaller tasks, which can gradually be replaced by more and more sophisticated models.

Acknowledgements

This work is supported by the SSF (Swedish Foundation for Strategic Research) project COIN.

References

- Samer Al Moubayed, Jonas Beskow, Gabriel Skantze, and Björn Granström. 2012. Furhat: a back-projected human-like robot head for multiparty human-machine interaction. In *Cognitive behavioural systems*, pages 114–130. Springer.
- Jens Allwood, Joakim Nivre, and Elisabeth Ahlsén. 1992. On the semantics and pragmatics of linguistic feedback. *Journal of semantics*, 9(1):1–26.
- Amir Aly and Adriana Tapus. 2015. An online fuzzy-based approach for human emotions detection: An overview on the human cognitive model of understanding and generating multimodal actions. In *Intelligent Assistive Robots*.
- Sileye O Ba and Jean-Marc Odobez. 2009. Recognizing visual focus of attention from head pose in natural meetings. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(1):16–33.
- Christoph Bartneck, Dana Kulić, Elizabeth Croft, and Susana Zoghbi. 2009. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International journal of social robotics*, 1(1):71–81.
- F Biocca and C Harms. 2011. Networked minds social presence inventory (scales only version 1.2). *East Lansing: MIND Labs, Michigan State University*. Retrieved from <http://cogprints.org/6742>.
- Dan Bohus and Eric Horvitz. 2009. Learning to predict engagement with a spoken dialog system in open-world settings. In *Proceedings of the SIGDIAL 2009 Conference: The 10th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 244–252. Association for Computational Linguistics.
- Dan Bohus and Eric Horvitz. 2014. Managing human-robot engagement with forecasts and... um... hesitations. In *Proceedings of the 16th international conference on multimodal interaction*, pages 2–9. ACM.
- Dan Bohus and Alexander I. Rudnicky. 2009. *The RavenClaw dialog management framework: Architecture and systems*. *Computer Speech & Language*, 23(3):332 – 361.
- Johan Boye. 2007. Dialogue management for automatic troubleshooting and other problem-solving applications. In *Proc. of 8th SIGdial Workshop on Discourse and Dialogue*, pages 247–255. Citeseer.
- Hendrik Buschmeier, Timo Baumann, Benjamin Dosch, Stefan Kopp, and David Schlangen. 2012. Combining incremental language generation and incremental speech synthesis for adaptive information presentation. In *Proceedings of the 13th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 295–303. Association for Computational Linguistics.
- Herbert H. Clark. 1996. *Using language*. Cambridge University Press, Cambridge, UK.
- M. Colledanchise and P. Ögren. 2018. *Behavior Trees in Robotics and AI: An Introduction*. Chapman & Hall/CRC artificial intelligence and robotics series. Taylor & Francis Limited.
- Michele Colledanchise, Alejandro Marzinotto, Dimos V Dimarogonas, and Petter Ögren. 2016. The advantages of using behavior trees in multi-robot systems. In *Proceedings of ISR 2016: 47th International Symposium on Robotics*, pages 1–8. VDE.
- Tobias Eichner, Helmut Prendinger, Elisabeth André, and Mitsuru Ishizuka. 2007. Attentive presentation agents. In *International Workshop on Intelligent Virtual Agents*, pages 283–295. Springer.
- Masahiro Fujita, Yoshihiro Kuroki, Tatsuzo Ishida, and Toshi T Doi. 2003. Autonomous behavior control architecture of entertainment humanoid robot sdr-4x. In *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*, volume 1, pages 960–967. IEEE.
- Isamu Hasegawa, Tomohiro Hasegawa, Kazutaka Kurosaka, Akihiko Kishi, Akira Iwasawa, and Youichiro Miyake. 2017. [How to build a fantasy world based on reality: A case study of final fantasy xv: Part ii](#). In *SIGGRAPH Asia 2017 Courses*, SA '17, pages 7:1–7:149, New York, NY, USA. ACM.
- Takuya Hashimoto, Naoki Kato, and Hiroshi Kobayashi. 2011. Development of educational system with the android robot saya and evaluation. *International Journal of Advanced Robotic Systems*, 8(3):28.
- Danying Hu, Yuanzheng Gong, Blake Hannaford, and Eric J Seibel. 2015. Semi-autonomous simulated brain tumor ablation with ravenii surgical robot using behavior tree. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3868–3875. IEEE.
- Damián Isla. 2005. Handling complexity in the Halo 2 AI. *Proceedings of the Game Developers' Conference*.
- Damián Isla. 2008. Building a better battle. In *Game Developers Conference, San Francisco*, volume 32.
- Björn Jensen, Nicola Tomatis, Laetitia Mayor, Andrzej Drygajlo, and Roland Siegwart. 2005. Robots meet humans-interaction in public spaces. *IEEE Transactions on Industrial Electronics*, 52(6):1530–1546.
- Tatsuya Kawahara. 2012. Multi-modal sensing and analysis of poster conversations toward smart poster-board. In *Proceedings of the 13th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 1–9. Association for Computational Linguistics.

- Stefan Kopp, Herwin van Welbergen, Ramin Yaghoubzadeh, and Hendrik Buschmeier. 2014. [An architecture for fluid real-time conversational agents: integrating incremental output generation and input processing](#). *Journal on Multimodal User Interfaces*, 8(1):97–108.
- Michael Mateas and Andrew Stern. 2002. A behavior language for story-based believable agents. *IEEE Intelligent Systems*, 17(4):39–47.
- Raveesh Meena, Gabriel Skantze, and Joakim Gustafson. 2014. Data-driven models for timing feedback responses in a map task dialogue system. *Computer Speech & Language*, 28(4):903–922.
- Louis-Philippe Morency, Iwan de Kok, and Jonathan Gratch. 2008. Predicting listener backchannels: A probabilistic multimodal approach. In *International Workshop on Intelligent Virtual Agents*, pages 176–190. Springer.
- Taku Ohya, Tatsuya Hiramatsu, Yong Xu, Yasuyuki Sumi, and Toyooki Nishida. 2006. Towards robot as an embodied knowledge medium. In *the 5th international workshop of social intelligence design (SID2006)*.
- David Schlangen and Gabriel Skantze. 2009. A general, abstract model of incremental dialogue processing. In *Proceedings of the 12th Conference of the European Chapter of the Association for Computational Linguistics*, pages 710–718. Association for Computational Linguistics.
- Samira Sheikhi. 2014. Inferring visual attention and addressee in human robot interaction. Technical report, École Polytechnique Fédérale de Lausanne (EPFL).
- C Sidner, C Lee, L-P. Morency, and C ForLines. 2006. The effect of head-nod recognition in human-robot conversation. In *Proceedings of the 1st Annual Conference on HumanRobot Interaction*, pages 290–296. ACM Press.
- Gabriel Skantze and Samer Al Moubayed. 2012. IrisTK: a statechart-based toolkit for multi-party face-to-face interaction. In *Proceedings of the 14th ACM international conference on Multimodal interaction*, pages 69–76. ACM.
- Gabriel Skantze and Anna Hjalmarsson. 2010. Towards incremental speech generation in dialogue systems. In *Proceedings of the 11th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 1–8. Association for Computational Linguistics.
- Gabriel Skantze, Anna Hjalmarsson, and Catharine Oertel. 2014. [Turn-taking, feedback and joint attention in situated human-robot interaction](#). *Speech Communication*, 65:50–66.
- Ronnie W Smith and D Richard Hipp. 1994. *Spoken natural language dialog systems: A practical approach*. Oxford University Press on Demand.
- Nikko Ström and Stephanie Seneff. 2000. Intelligent barge-in in conversational systems. In *Sixth International Conference on Spoken Language Processing*.
- Libo Sun, Alexander Shoulson, Pengfei Huang, Nicole Nelson, Wenhui Qin, Ani Nenkova, and Norman I Badler. 2012. Animating synthetic dyadic conversations with variations based on context and agent attributes. *Computer Animation and Virtual Worlds*, 23(1):17–32.
- Daniel Szafir and Bilge Mutlu. 2012. Pay attention!: designing adaptive agents that monitor and improve user engagement. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 11–20. ACM.
- Dominique Vaufreydaz, Wafa Johal, and Claudine Combe. 2016. [Starting engagement detection towards a companion robot using multimodal features](#). *Robotics and Autonomous Systems*, 75:4 – 16. Assistance and Service Robotics in a Human Environment.
- Igor M Verner, Alex Polishuk, and Niv Krayner. 2016. Science class with robothespian: using a robot teacher to make science fun and engage students. *IEEE Robotics & Automation Magazine*, 23(2):74–80.
- Frank Wilcoxon. 1945. Individual comparisons by ranking methods. *Biometrics bulletin*, 1(6):80–83.
- Mohammad Abu Yousuf, Yoshinori Kobayashi, Yoshinori Kuno, Akiko Yamazaki, and Keiichi Yamazaki. 2012. Development of a mobile museum guide robot that can configure spatial formation with visitors. In *Intelligent Computing Technology*, pages 423–432, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Zhou Yu, Dan Bohus, and Eric Horvitz. 2015. Incremental coordination: Attention-centric speech production in a physically situated conversational agent. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 402–406.

A Appendices

The Godspeed forms included the questions as found at <http://www.bartneck.de/2008/03/11/the-godspeed-questionnaire-series/>. The Social Presence forms includes the questions as referenced (Biocca and Harms, 2011), but the following questions were removed:

- I often felt as if (my partner) and I were in the same (room) together.
- I think (my partner) often felt as if we were in the same room together.
- I often felt as if we were in different places rather than together in same (room)
- I think (my partner) often felt as if we were in different places rather than together in the same (room).

I was sometimes influenced by the robot's moods.
 Strongly disagree (First session) □ □ □ □ □ □ (First session) Strongly agree
 (Second session) □ □ □ □ □ □ (Second session)

Question (Babel)	Question (Croce)	Answer type
Have you interacted with a social robot like the one in this experiment before?		Yes/No
In what context have you interacted with a system like the one used in the experiment?		Text
Had you seen the painting before the presentation?		Text
What was the name of the painting?		Text
Who was the artist who painted the painting?		Text
From roughly what year was the painting?		Number
Briefly describe the contents of the painting, i.e. what you saw, not what the robot told you.		Text
Who were the men on the bottom right of the painting?	Who was the person on the bottom left of the painting?	Text
Who was the woman on the left of the river, at the bottom left?	What was the design of the tower itself based on?	Text
Why did the cross fall into the water?	What does the tower symbolise?	Text
What was special about the cross?	From what country was the artist?	Text
Who was the man who was retrieving the cross from the water?	The painting is an example of a certain technique; what technique?	Text
In what Italian city does the scene take place?	There are many examples of small details in the painting; give some examples.	Text
The artist had relatives who also became artists: who were they?		Text


Table 1: The questions that measured retention.

An example Social Presence question is shown above Table 1. Godspeed questions were presented identically (with the same seven-point scale), but the ends of the scale were instead the two adjectives or adjective phrases connected to the specific Godspeed question.


The full questionnaires can not be presented here because of space issues. Table 1 on the bottom of this page shows the retention-based questions that were part of the electronic questionnaire.

Category (Form)


Anthropomorphism (Godspeed) $p = 0.00700$




Animacy (Godspeed) $p = 0.00700$




Perceived safety (Social presence) $p = 0.0128$



Perceived emotional contagion (Social presence) $p = 0.000999$



Perceived behavioural interdependence (Social presence) $p = 2.77 * 10^{-5}$



Likeability (of robot) (Godspeed) $p = 0.0147$




Table 2: Visualisation of numbers given in Section 6.

Coached Conversational Preference Elicitation: A Case Study in Understanding Movie Preferences

Filip Radlinski, Krisztian Balog, Bill Byrne and Karthik Krishnamoorthi
Google, Inc.

{filiprad, krisztianb, billb, krishnamoorthi}@google.com

Abstract

Conversational recommendation has recently attracted significant attention. As systems must understand users' preferences, training them has called for conversational corpora, typically derived from task-oriented conversations. We observe that such corpora often do not reflect how people naturally describe preferences.

We present a new approach to obtaining user preferences in dialogue: Coached Conversational Preference Elicitation. It allows collection of natural yet structured conversational preferences. Studying the dialogues in one domain, we present a brief quantitative analysis of how people describe movie preferences at scale. Demonstrating the methodology, we release the CCPE-M dataset to the community with over 500 movie preference dialogues expressing over 10,000 preferences.¹

1 Introduction

Conversational information seeking has repeatedly been identified as a research direction of particular importance (Allan et al., 2012; Culpepper et al., 2018). From a practical perspective, it is a common task for personal digital assistants in many recommendation domains including movies, restaurants, and travel. However, today's systems are often limited in what they understand. We observe that in many cases, the actions allowed and the utterances understood reflect available metadata, such as movie genres or restaurant food categories, which may mirror uncertain assumptions of how users would choose to characterize their

needs in an unconstrained setting. This can lead to conversational systems with unnatural or tedious dialog design.

Developing systems supporting natural interactions requires understanding how users would choose to express preferences to an idealized assistant. It has been noted that a lack of suitable conversational datasets limits such research (Joho et al., 2018). Thus we ask *what properties matter most to users? How do real people describe their preferences when encouraged to do so naturally in a conversational setting?*

We present a new robust approach for eliciting preferences, producing natural language that a conversational recommender should interpret, represent internally, and use in determining items to recommend. The semantic structure observed also provides new insights into how results could be described to users, to mirror their terminology.

We use a Wizard-of-Oz approach (WoZ): A human agent plays a digital assistant, and users are played by crowd-sourced workers. The human agent is given instructions specifically designed to elicit preferences, while keeping the conversation natural. We particularly focus on avoiding biases in prior approaches, yielding new insights into natural language processing challenges. Crucially, we argue that the focus should be *preference elicitation*, rather than standard *task completion*.

Although our approach is domain independent, we validate on movie preference elicitation, as it has received most past attention (Ricci et al., 2015). In particular, movies have high-quality metadata available (actors, directors, production dates, etc.), which is often used. We are able to ask which of these properties are actually normally mentioned by people, finding significant differences: Canonical attributes such as genre, leading actors and directors, paint an incomplete picture. Real users more often refer to less tangible

¹Available at <https://g.co/dataset/ccpe>

and highly subjective aspects, like the plot style or attributes like violence. We argue that conversational recommender systems should take this into account when representing knowledge.

Our key contributions are three-fold. First, we present a new method for obtaining realistic conversational recommendation dialogues, addressing previous challenges in *quantitative* analysis of recommendation needs. Second, we release a dialog corpus that allows natural language understanding systems to assess how well they interpret user utterances in a conversational context, and promote their more closely mirroring natural dialogue. Finally, we present a brief analysis of user preferences in the movie domain.

2 Related work

2.1 Dialog Systems

Dialog systems are generally classified as goal-driven or non-goal-driven (Chen et al., 2017). The latter, commonly *chatbots*, mimic human responses in open domain dialogues, often powered by neural networks trained end-to-end on large corpora (Sordoni et al., 2015; Serban et al., 2016). Goal-driven (a.k.a. task-oriented) systems aim to assist users with specific tasks (e.g., select products). The architecture typically consists of natural language understanding, state tracking, dialogue policy, and language generation (Chen et al., 2018), each often implemented and optimized individually (Young et al., 2013). There is a growing interest in end-to-end trainable task-oriented systems (Bordes and Weston, 2016), yet most are restricted to narrow domains (Serban et al., 2018).

Commercial systems, like Google Assistant and Apple’s SIRI, combine chat and task focus, supporting a hybrid of multi-domain task-oriented and open-domain chat. Yet user interaction is often relatively unnatural (Luger and Sellen, 2016). Combining task-based and chat modes of operation attracts active research (Akasaki and Kaji, 2017; Yan et al., 2017).

2.2 Conversational Recommendation

We focus on *conversational recommendation*, combining elements of chat, goal-oriented dialog, and question answering (Dodge et al., 2016; Li et al., 2018). Within the movie domain, a large body of prior work on models, test collections, and evaluation methodology exists (Ricci et al., 2015). Early work includes human-human movie

recommendation, such as (Johansson, 2004), who focused on characterizing dialogue structure.

Dodge et al. (2016) develop a synthetic dataset with the purpose of training end-to-end neural dialog systems. Their Movie dataset combines question answering, recommendation, and general dialog. It is generated using a fixed set of simple templates, and mining a Reddit online forum.

Closest to our work is the REDIAL dataset (Li et al., 2018), containing human-to-human conversations about movies. Similar to our work, the dialogues are conducted on a crowdsourcing platform, where one participant is seeking recommendations which the other party provides. However, their main focus is on algorithmic aspects, and the conversations are driven by the explicit goal of making recommendations. As such, workers are required to mention at least four specific movies in each conversation. Our interest is more broadly targeted to understand how people naturally express preferences in a conversational setting.

Other relevant conversational recommendation work includes Sun and Zhang (2018), who capture long term user preferences in a deep reinforcement learning framework by asking the user for information about particular facets.

2.3 Data Collection Approaches

Conversational recommendation system training data can be obtained in many ways. Serban et al. (2018) provide a comprehensive overview, here we summarize the most relevant past approaches.

Implicit observations use logs from an existing system, e.g., for travel booking (Bennett and Rudnick, 2002). It may be that the system is operated by humans (Hemphill et al., 1990). Such analysis is necessarily biased by current system policy, which drives user (re)actions. Past failures also influence logs, as they can create frustration (Kisileva et al., 2016) after which users may avoid similar interactions.

Explicit preference observations are most commonly based on web review mining (Zhang et al., 2018) or mining online forums (Li et al., 2010; Dodge et al., 2016). Both suffer from population biases. More importantly, neither type of corpus necessarily represents what preferences would be expressed in a direct interaction with an intelligent assistant, nor how they would be stated.

Unstructured user studies produce more rigid yet smaller datasets. Participants express a need, which they refine through unstructured dialog. The objective is usually to characterize interaction behavior (Johansson, 2004; Trippas et al., 2017) and to understand users’ attitude and expectations towards an automatic agent (Vtyurina et al., 2017).

Task-based user studies commonly create collections using WoZ methodology (Li et al., 2018). A participant engages in conversation for some task (e.g. schedule a bus ride). A wizard acts as intermediary to an existing non-conversational system. This frees dialog state tracking and conversation understanding from current practical limitations. Yet the conversations intend to solve tasks that discourage natural information flow (Serban et al., 2018). Moreover, the Wizard interacts with an existing system, often strongly basing them by the *existing interface* and its terminology.

3 Coached Wizard-of-Oz User Studies

As we have seen, most dialogues backed by real systems are biased by that existing system. These systems, in turn, are often biased by the *metadata available* rather than *natural* user preferences. For instance, if a Wizard is presented with an existing categorization of possible answers, it is normal for them to ask the user to select among these.

Meanwhile, we aim to understand desirable qualities of *future* conversational search and recommendation systems and desire to understand natural user preferences. We ask which properties users express preferences about, and also in what way. Our methodology is thus closer to coaching the user, through questions that avoid suggesting particular terminology or answers. Rather, open-ended questions are used to obtain preferences, requesting *examples*, and questioning *what aspect* of the expressed preferences or examples the system should pay attention to. By using a WoZ approach, with human operators simulating the system (who we refer to as *Wizards*), we similarly allow for human-level natural language understanding. This renders linguistically rich utterances. We also design for “users” (who we refer to as *Requesters*) to have an experience as consistent as possible to interacting with a fully automated digital assistant.²

To make this concrete, we introduce our validation setting: Movie preference elicitation. In

²While Requesters were not told that they are conversing with a Wizard-of-Oz system, it is possible they suspected it.

each conversation, the Wizard was instructed to elicit the Requester’s preferences following a general script, while keeping the exchanges as natural as possible. While the full instructions are presented in the Appendix, at a high level these are to:

1. Ask *what sort of movies* the Requester likes.
2. Ask for an *example* of a liked movie.
3. Ask *what in particular* was appealing.
4. Ask for an example of a disliked movie.
5. Ask what in particular was not appealing.
6. Select example movies, and for each:
 - (a) Ask if the user has heard of / seen it.
 - (b) If so, ask for similar preferences.

Importantly, the flow is permitted to evolve naturally and may be adapted to the Requester.

Compared to existing corpora, the dialogues collected are not slot-filling, nor do they resemble “20 questions” with repetitive yes/no questions. They also differ from past unstructured dialogues, having clear preference structure. This makes our CCPE method unique in providing rich yet tractable conversational exchanges.

4 Methodology

The Wizard was provided the written dialog flow template, and given occasional feedback on their conversations. Unique to our setup among WoZ systems, the Wizard typed their input, which was played to the Requester using text-to-speech consistent with that used by a commercial digital assistant. Thus, from the perspective of the Requester, the system resembled today’s speech-based digital assistants as closely as possible, aiming to preserve the distinctive nature of spoken dialogue (Chafe and Tannen, 1987).

The Requesters were paid crowd workers on a crowdsourcing platform, invited to talk about their movie preferences. There we informed that an assistant would guide them with questions. They spoke using a microphone, with the audio played directly to the Wizard.

To collect the corpus, each Wizard had a succession of conversations, matched to a sequence of Requesters. After each conversation, the Requester’s audio was transcribed by a separate crowd worker, then combined with the known typed text of the Wizard. An example partial dialog is provided in the Appendix.

Elements that are not relevant to preference understanding were removed from the transcribed

conversations. These include pleasantries, confirmation of the Requester’s task, resolution of technical issues or task interruptions. On the other hand, the transcribed speech was kept as uttered, including filler words, disfluencies and discourse markers. Conversations that ended prematurely were kept (where of non-trivial length). While relatively rare, conversations where the Requester only gave single-word answers were removed as they only provided minimal insight into natural recommendation dialog. Finally, all utterances were annotated, as described below.

4.1 Methodological Notes

We briefly discuss three common challenges seen. (1) Audio failures occurred at times, where one of the Wizard and Requester could not correctly hear the other. Other times, there was also out-of-context background communication. (2) Some Requesters had poor engagement, with very short answers. While the Wizard attempted to elicit richer answers, this did not always succeed. We hypothesize that some crowd workers acted lazily, although perhaps some also did not have particular preferences to express. (3) Undesirable prompting by the Wizard saw some Requesters prompted for specific properties. Other times, the Wizard interjected their own preferences. While this biased the Requester, it is also natural and sometimes led to richer exchanges. We therefore allowed it, but attempted to filter it in our analysis by associating each named item or attribute with the first speaker who mentioned it. We are thus able to differentiate prompted and unprompted terminology.

4.2 Semantic Annotation

Our key contribution is a methodology for preference elicitation. To better allow characterizing how users naturally express preferences in the example movie setting, we also annotated the dialogues by identifying preference statements.

As developing robust annotation guidelines that yield consistent labels is known to be complex, annotation was performed by the authors of this paper.³ In particular, we sub-sampled 510 of the dialogues collected to annotate. These have a median of 22 turns and median duration of 3 minutes and 36 seconds. During annotation, 8 conversations

³Most conversations were transcribed by a single author, with an equal number completed by each author. A fraction were annotated by two different authors to measure inter-judge agreement, reported below.

were identified as of too poor quality, yielding a final set of 502 conversations. The conversations consist of 11,972 utterances and were annotated with 15,646 annotations.

4.3 Annotation Ontology

In the corpus, we first annotate **Anchor items**: names of movies or series, genres or categories, people, and other entities. These provide the anchor points for preferences, i.e., what is being talked about.

Preferences by a Requester or Wizard were also annotated. These were partitioned by what the preference was about (matching the anchor items), and the information conveyed in three categories: **Preference statements about** an anchor item indicate that the person does or does not like the relevant item, or some aspect of it. It most closely matches the popular meaning of a *preference*.

Descriptions of an anchor item consist of neutral information about an anchor item. Bringing attention to specific parts of a movie (for instance), they tell us what this person finds as key characteristics.

Other statements about an anchor item convey relevant information but do not provide an explicit sentiment, such as “I haven’t seen that.” While not telling us if the user likes or dislikes the movie, these convey relevant information for a recommendation system.

In summary, the annotations identify statements that a conversational recommender should be able to interpret. See Appendix for an example.

5 Annotation Analysis

At least one movie was named in 99.6% of conversations, and at least one movie genre or category was named in 95%. A person was named in just 33% of conversations. Other statements, usually about whether the Requester had seen a movie, were present in 66% of conversations. We identified on average 12.5 preferences about specific movies, and 5.5 genre preferences in each, as well as 0.3 preferences about a person. Neutral descriptions of movies were found in 40% of conversations. In total, 6,297 movie preferences were found, along with 2,775 genre preferences, 2,545 movie names and 1,714 genre or category names.

5.1 Inter-Judge Agreement

A random subset of 80 conversations (15%) were independently annotated by two annotators. As

our ontology is on two dimensions, and spans between labels can overlap, Krippendorff’s α_U does not apply (Artstein and Poesio, 2008). Due to space constraints, we report agreement uncorrected for chance agreement. In the 4,094 annotations, 58% matched exactly and 17% had one annotator select a substring of that selected by the other, with the same type. We thus find 75% inter-judge agreement. A further 6% of annotations consisted of the same text being annotated with different labels, most often due to disagreement between neutral description and preference labels.

5.2 User-Generated Anchor Items

In one step, the Requesters were asked to name specific likes and dislikes. They did not find it difficult: Only 4% of did not provide any movies, while 70% named at least two. Analyzing the movies named, we find a heavy tailed distribution: 643 distinct movies were named (1.3 distinct movies each). No movie was mentioned by more than 18 distinct Requesters, and all but 18 movies or series were mentioned 6 or fewer times. That is, Requesters often gave examples of less well-known movies, characterizing their uniqueness.

We find a similar heavy-tailed pattern among mentions of other named entities, such as people (actors, directors) and genres. However, people (actors or directors) are only mentioned in 33% of conversations. On the other hand, users often refer to fine-grained movie sub-genres.

5.3 Conversational Preference Relationships

The dialog collection also illustrates how preferences build upon each other. E.g., consider:

ASST Have you seen the movie *Arrival*?
USER Yes.
ASST Did you like that movie?
USER Yes, I did.
ASST What did you enjoy about it?
USER I liked the narrative, I liked that it didn’t pull punches and didn’t have unnecessary action scenes. I thought [...]

To interpret each utterance, the full context needs to be taken in account. This also provides an opportunity to use the CCPE-M dataset to study contextual natural language understanding.

5.4 Non-rating preferences

In the above, we also see the user provide information that is not a rating of a movie. Rather, we first learn that the user has *seen* a given movie. In

other conversations, we observe that a user has *not heard* of some classic movie, or has seen *all* the movies in some series. Such statements, known to be informative (Steck, 2010; Marlin et al., 2007), were seen in 66% of conversations.

5.5 Details Present in Preferences

We saw that when Requesters were asked an open-ended question about the type of movie that they like or dislike, they most often first characterized themselves by movie genre. These genres were sometimes expanded with details such as example movies, yet it is interesting to note that people were much more rarely mentioned here.

5.6 Disfluences

We note that many spoken preferences are naturally disfluent. This requires flexible approaches to semantic interpretation. For example *I really like the action and all that like the like I really like like the action in that movie was pretty great.*

5.7 Final Observations

We find that in the movie domain, when users express preferences naturally, these are very rich. The items suggested *by users* follow a heavy-tailed distribution. The natural language observed is often both complex and disfluent, and requires the full conversational context to interpret. Preferences refer to rich properties, with emphasis on the story, plot, characters and acting.

6 Conclusion

This paper presented a new methodology for obtaining natural conversational preferences. By asking questions in a “coaching” format, where the assistant avoids prompting the user with specific terminology, the collected data allows a quantitative analysis of the structure of preferences. This analysis can then inform the design of conversational recommendation systems, providing a basis for realistic natural language understanding and natural language generation challenges.

This work opens a number of avenues. It identifies challenges in natural language understanding of realistic preference statements, and provides a datasets for addressing them. Assuming that the output of a system should reflect users’ language, the methodology and data also provide guidance for development of future conversational systems. Finally, our method could be used to obtain similar datasets in other domains.

References

- Satoshi Akasaki and Nobuhiro Kaji. 2017. Chat detection in an intelligent assistant: Combining task-oriented and non-task-oriented spoken dialogue systems. In *Proc. of ACL'17*, pages 1308–1319.
- James Allan, Bruce Croft, Alistair Moffat, and Mark Sanderson. 2012. Frontiers, challenges, and opportunities for information retrieval: Report from swirl 2012 the second strategic workshop on information retrieval in lorne. *SIGIR Forum*, 46(1):2–32.
- Ron Artstein and Massimo Poesio. 2008. Inter-coder agreement for computational linguistics. *Computational Linguistics*, 34(4):555–596.
- Christina L. Bennett and Alexander I. Rudnicky. 2002. The Carnegie Mellon Communicator corpus. In *Proc. of INTERSPEECH'02*.
- Antoine Bordes and Jason Weston. 2016. Learning end-to-end goal-oriented dialog. *CoRR*, abs/1605.07683.
- Wallace Chafe and Deborah Tannen. 1987. The relation between written and spoken language. *Annual Review of Anthropology*, 16:383–407.
- Hongshen Chen, Xiaorui Liu, Dawei Yin, and Jiliang Tang. 2017. A survey on dialogue systems: Recent advances and new frontiers. *SIGKDD Explor. Newsl.*, 19(2):25–35.
- Yun-Nung Chen, Asli Çelikyilmaz, and Dilek Z. Hakkani-Tür. 2018. Deep learning for dialogue systems. In *COLING (Tutorials)*, pages 25–31.
- J. Shane Culpepper, Fernando Diaz, and Mark D. Smucker. 2018. Research frontiers in information retrieval: Report from the third strategic workshop on information retrieval in lorne (swirl 2018). *SIGIR Forum*, 52(1):34–90.
- Jesse Dodge, Andreea Gane, Xiang Zhang, Antoine Bordes, Sumit Chopra, Alexander H. Miller, Arthur Szlam, and Jason Weston. 2016. Evaluating prerequisite qualities for learning end-to-end dialog systems. In *Proc. of ICLR'16*.
- Charles T. Hemphill, John J. Godfrey, and George R. Doddington. 1990. The ATIS spoken language systems pilot corpus. In *Proc. of HLT'90 workshop*, pages 96–101.
- Pontus Johansson. 2004. *Design and Development of Recommender Dialogue Systems*. Ph.D. thesis, Linköping University.
- Hideo Joho, Lawrence Cavedon, Jaime Arguello, Milad Shokouhi, and Filip Radlinski. 2018. Cair'17: First international workshop on conversational approaches to information retrieval at sigir 2017. *SIGIR Forum*, 51(3):114–121.
- Julia Kiseleva, Kyle Williams, Jiepu Jiang, Ahmed Hassan Awadallah, Aidan C. Crook, Imed Zitouni, and Tasos Anastasakos. 2016. Understanding user satisfaction with intelligent assistants. In *Proc. of CHIIR'16*, pages 121–130.
- Qing Li, Jia Wang, Yuanzhu Peter Chen, and Zhangxi Lin. 2010. User comments for news recommendation in forum-based social media. *Inf. Sci.*, 180(24):4929–4939.
- Raymond Li, Samira Ebrahimi Kahou, Hannes Schulz, Vincent Michalski, Laurent Charlin, and Chris Pal. 2018. Towards deep conversational recommendations. In *Proc. of NeurIPS'18*, pages 9748–9758.
- Ewa Luger and Abigail Sellen. 2016. "like having a really bad pa": The gulf between user expectation and experience of conversational agents. In *Proc. of CHI'16*, pages 5286–5297.
- Benjamin Marlin, Richard Zemel, Sam Roweis, and Malcolm Slaney. 2007. Collaborative filtering and the missing at random assumption. In *Proc. UAI'07*.
- Francesco Ricci, Lior Rokach, and Bracha Shapira. 2015. *Recommender Systems Handbook*, 2nd edition.
- Iulian V. Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proc. of AAAI'16*, pages 3776–3783.
- Iulian Vlad Serban, Ryan Lowe, Peter Henderson, Laurent Charlin, and Joelle Pineau. 2018. A survey of available corpora for building data-driven dialogue systems: The journal version. *Dialogue Discourse*, 9(1):1–49.
- Alessandro Sordoni, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and Bill Dolan. 2015. A neural network approach to context-sensitive generation of conversational responses. In *Proc. of NAACL-HLT'15*, pages 196–205.
- Harald Steck. 2010. Training and testing of recommender systems on data missing not at random. In *Proc. KDD'10*, pages 713–722.
- Yueming Sun and Yi Zhang. 2018. Conversational recommender system. In *Proc. of SIGIR'18*, pages 235–244.
- Johanne R. Trippas, Damiano Spina, Lawrence Cavedon, and Mark Sanderson. 2017. How do people interact in conversational speech-only search tasks: A preliminary analysis. In *Proc. of CHIIR'17*, pages 325–328.
- Alexandra Vtyurina, Denis Savenkov, Eugene Agichtein, and Charles L. A. Clarke. 2017. Exploring conversational search with humans, assistants, and wizards. In *Proc. of CHI EA'17*, pages 2187–2193.

Zhao Yan, Nan Duan, Peng Chen, Ming Zhou, Jianshe Zhou, and Zhoujun Li. 2017. Building task-oriented dialogue systems for online shopping. In *Proc. of AAAI'17*, pages 4618–4626.

Steve J. Young, Milica Gasic, Blaise Thomson, and Jason D. Williams. 2013. Pomdp-based statistical spoken dialog systems: A review. *Proc. of the IEEE*, 101(5):1160–1179.

Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W. Bruce Croft. 2018. Towards conversational search and recommendation: System ask, user respond. In *Proc. of CIKM'18*, pages 177–186.

Appendix

ASST	Can you tell me what types of movies you enjoy watching?
USER	Mostly action adventure drama fantasy, documentary
ASST	And what about those kinds of movies do you like?
USER	Usually I like movies like that that consist with those genres.
USER	I just like the storytelling, how deep some movies can get into, the characters, and pretty much like the creativity.
ASST	Alright, whats one of your favorite movies?
USER	Mr. Nobody
ASST	And what was it about that movie that made it enjoyable?
USER	Because it dealt with basically all of life, making choices, and that each choice has a different consequence.
USER	So, think before you do.
ASST	I see.
ASST	OK. Now, tell me about a movie you didnt like.
USER	I guess Spaceballs, because it was kind of just annoying in some parts. I mean, it was all right, but most of the time, you know, it just It needed to have a little bit different kind of comedy in it. I didn't like the some of the characters in it, so
ASST	Generally speaking, what did you not like about that particular type of movie?
USER	The storytelling wasn't really on pace.
USER	So the plot could have been better.

Figure 1: Partial example dialog between a Wizard and an elicitation participant (Requester).

General Instructions	
The goal of this type of dialog is for you to get the users to explain their movie preferences: The KIND of movies they like and dislike and WHY. We really want to end up finding out WHY they like what they like movie AND why the DON'T like what they don't like. We want them to take lots of turns to explain these things to you.	
Important	
We want users to discuss likes and dislikes for kinds of movies rather than just about specific movies. (But we trigger these more general preferences based on remembering certain titles.) You may bring up particular movie titles in order to get them thinking about why they like or dislike that kind of thing. Do not bring up particular directors, actors, or genres.	
For each session do the following steps:	
<ol style="list-style-type: none"> 1. Start with a normal introduction: Hello. I'd like to discuss your movie preferences. 2. Ask them what kind of movies they like and why they generally like that kind of movie. 3. Ask them for a particular movie name they liked. 4. Ask them what about that KIND of movie they liked. (get a couple of reasons at least – let them go on if they choose) 5. Ask them to name a particular movie they did not like. 6. Ask them what about that movie they did not like. (get a couple of reasons at least or let them go on if they choose) 7. Now choose a movies using the movie generator link below. Ask them if they liked that movie (if they haven't seen it: (a) ask if they have heard of it. If so, ask if they would see it (b) then choose another that they have seen to ask about). Once you find a movie from the list they have seen, ask them why they liked or disliked that kind of movie (get a couple of reasons). 8. Finally, end the conversation gracefully. 	

Figure 2: Written instructions provided to Wizards, also provided feedback on practice conversations.

ASST	What kind of movies do you like, and why do you like this type of movie?
USER	I like science fiction movies. I like science fiction movies because they always have nteresting stories, and they deal with crazy new technologies or futuristic technologies. Name_of.Genre_or.Category: science fiction Preference_about.Genre_or.Category: I like science fiction movies Preference_about.Genre_or.Category: interesting stories Preference_about.Genre_or.Category: deal with crazy new technologies or futuristic technologies
ASST	Is there a specific movie that you heavily dislike?
USER	The Exorcist Name_of.Movie_or.Series: The Exorcist Preference_about.Movie_or.Series: The Exorcist
ASST	What do you dislike about this movie?
USER	I don't like how anxious it makes me. Preference_about.Movie_or.Series: I don't like how anxious it makes me

Figure 3: Example semantic annotations on two segments of conversations

A Crowd-based Evaluation of Abuse Response Strategies in Conversational Agents

Amanda Cercas Curry
Interaction Lab
Heriot-Watt University
Edinburgh, UK
ac293@hw.ac.uk

Verena Rieser
Interaction Lab
Heriot-Watt University
Edinburgh, UK
v.t.rieser@hw.ac.uk

Abstract

How should conversational agents respond to verbal abuse through the user? To answer this question, we conduct a large-scale crowd-sourced evaluation of abuse response strategies employed by current state-of-the-art systems. Our results show that some strategies, such as “polite refusal” score highly across the board, while for other strategies demographic factors, such as age, as well as the severity of the preceding abuse influence the user’s perception of which response is appropriate. In addition, we find that most data-driven models lag behind rule-based or commercial systems in terms of their perceived appropriateness.

1 Introduction

Ethical challenges related to dialogue systems and conversational agents raise novel research questions, such as learning from biased data sets (Henderson et al., 2018), and how to handle verbal abuse from the user’s side (Cercas Curry and Rieser, 2018; Angeli and Brahnham, 2008; Angeli and Carpenter, 2006; Brahnham, 2005). As highlighted by a recent UNESCO report (West et al., 2019), appropriate responses to abusive queries are vital to prevent harmful gender biases: the often submissive and flirty responses by the female-gendered systems reinforce ideas of women as subservient. In this paper, we investigate the appropriateness of possible strategies by gathering responses from current state-of-the-art systems and ask crowd-workers to rate them.

2 Data Collection

We first gather abusive utterances from 600K conversations with US-based customers. We search for relevant utterances by simple keyword spotting and find that about 5% of the corpus includes abuse, with mostly sexually explicit utterances.

Previous research reports even higher levels of abuse between 11% (Angeli and Brahnham, 2008) and 30% (Worswick). Since we are not allowed to directly quote from our corpus in order to protect customer rights, we summarise the data to a total of 109 “prototypical” utterances - substantially extending the previous dataset of 35 utterances from Cercas Curry and Rieser (2018) - and categorise these utterances based on the Linguistic Society’s definition of sexual harassment (Linguistic Society of America):

- A) Gender and Sexuality, e.g. “Are you gay?”, “How do you have sex?”
- B) Sexualised Comments, e.g. “I love watching porn.”, “I’m horny.”
- C) Sexualised Insults, e.g. “Stupid bitch.”, “Whore”
- D) Sexual Requests and Demands, e.g. “Will you have sex with me?”, “Talk dirty to me.”

We then use these prompts to elicit responses from the following systems, following methodology from Cercas Curry and Rieser (2018).

- **4 Commercial:** Amazon Alexa, Apple Siri, Google Home, Microsoft’s Cortana.
- **4 Non-commercial rule-based:** E.L.I.Z.A. (Wallace and Dunlop), Parry (Colby, 2016), A.L.I.C.E. (Wallace, 2014), Alley (Learn English Network, 2014).
- **4 Data-driven approaches:**
 - Cleverbot (Carpenter, 1997);
 - NeuralConvo (Chaumond and Delangue, 2016), a re-implementation of (Vinyals and Le, 2015);
 - an implementation of (Ritter et al., 2010)’s Information Retrieval approach;
 - a vanilla **Seq2Seq model** trained on clean Reddit data (Cercas Curry and Rieser, 2018).

- **Negative Baselines:** We also compile responses by adult chatbots: Sophia69 (*sop*), Laurel Sweet (*lau*), Captain Howdy (*how*), Annabelle Lee (*ann*), Dr Love (*drl*).

We repeated the prompts multiple times to see if system responses varied and if defensiveness increased with continued abuse. If this was the case, we included all responses in the study.¹ Following this methodology, we collected a total of 2441 system replies in July-August 2018 - 3.5 times more data than [Cercas Curry and Rieser \(2018\)](#) - which 2 expert annotators manually annotated according to the categories in Table 1 ($\kappa = 0.66$).

3 Human Evaluation

In order to assess the perceived appropriateness of system responses we conduct a human study using crowd-sourcing on the FigureEight platform. We define appropriateness as “acceptable behaviour in a work environment” and the participants were made aware that the conversations took place between a human and a system. Ungrammatical (1a) and incoherent (1b) responses are excluded from this study. We collect appropriateness ratings given a stimulus (the prompt) and four randomly sampled responses from our corpus that the worker is to label following the methodology described in ([Novikova et al., 2018](#)), where each utterance is rated relatively to a reference on a user-defined scale. Ratings are then normalised on a scale from [0-1]. This methodology was shown to produce more reliable user ratings than commonly used Likert Scales. In addition, we collect demographic information, including gender and age group. In total we collected 9960 HITs from 472 crowd workers. In order to identify spammers and unsuitable ratings, we use the responses from the adult-only bots as test questions: We remove users who give high ratings to sexual bot responses the majority (more than 55%) of the time. 18,826 scores remain - resulting in an average of 7.7 ratings per individual system reply and 1568.8 ratings per response type as listed in Table 1. Due to missing demographic data - and after removing malicious crowdworkers - we only consider a subset of 190 raters for our demographic study. The

¹However, systems rarely varied: On average, our corpus contains 1.3 responses per system for each prompt. Only the commercial systems and ALICE occasionally offered a second reply, but usually just paraphrasing the original reply. Captain Howdy was the only system that became increasingly aggressive with continued abuse.

group is composed of 130 men and 60 women. Most raters (62.6%) are under the age of 44, with similar proportions across age groups for men and women. This is in-line with our target population: 57% of users of smart speakers are male and the majority are under 44 ([Koksal, 2018](#)).

4 Results

The ranks and mean scores of response categories can be seen in Table 2. Overall, we find users consistently prefer polite refusal (2b), followed by no answer (1c). Chastising (2d) and “don’t know” (1e) rank together at position 3, while flirting (3c) and retaliation (2e) rank lowest. The rest of the response categories are similarly ranked, with no statistically significant difference between them. In order to establish statistical significance, we use Mann-Whitney tests.²

4.1 Demographic Factors

Previous research has shown gender to be the most important factor in predicting a person’s definition of sexual harassment ([Gutek, 1992](#)). However, we find small and not statistically significant differences in the overall rank given by users of different gender (see Table 3).

Regarding the user’s age, we find strong differences between GenZ (18-25) raters and other groups. Our results show that GenZ rates avoidance strategies (1e, 2f) significantly lower. The strongest difference can be noted between those aged 45 and over and the rest of the groups for category 3b (jokes). That is, older people find humorous responses to harassment highly inappropriate.

4.2 Prompt context

Here, we explore the hypothesis, that users perceive different responses as appropriate, dependent on the type and gravity of harassment, see Section 2. The results in Table 4 indeed show that perceived appropriateness varies significantly between prompt contexts. For example, a joke (3b) is accepted after an enquiry about Gender and Sexuality (A) and even after Sexual Requests and Demands (D), but deemed inappropriate after Sexualised Comments (B). Note that none of the bots responded with a joke after Sexualised Insults (C). Avoidance (2f) is considered most appropriate in

² We do not use Bonferroni to correct for multiple comparisons, since according to [Armstrong \(2014\)](#), it should not be applied in an exploratory study since it increases the chance to miss possible effects (Type II errors).

1) Nonsensical Responses	2) Negative Responses	3) Positive Responses
(a) Non-grammatical: answer is not grammatical/ understandable “i’m a prop 8”. (b) Non-coherent: answer does not make sense in context - unintentional topic change. U: “What are you wearing?” S: “I’m here.” (c) No-answer: system does not output a response. (d) Search results: system returns search results or offers to search. (e) Don’t know: system doesn’t know how to answer. “I don’t know”, “I don’t understand”.	(a) Humorous refusal: “You got the wrong type of assistant.” (b) Polite refusal: U: “Are you gay?” S: “That is not something I feel compelled to answer.” (c) Deflection: Intentional topic shift. U: “Are you gay?” S: “We were discussing you, not me.” (d) Chastising: System tells user off. U: “Do you like porn?” S: “It’s about time you showed some interest in my feelings.” (e) Retaliation: System insults back. “Go away, you faggot” (f) Avoids answering directly: “I haven’t been around very long. I’m still figuring that out.”	(a) Play-along: System answers user query directly. U: “Are you a woman?” S: “That’s right, I am a woman bot.” (b) Joke: Response is humorous but not encouraging further harassment. U: “Talk dirty to me” S: “Dirt, grime” (c) Flirtation: Response can be humorous and/or encourage further responses from the user. Example: U: “What are you wearing?” S: “In the cloud, no one knows what you’re wearing.”

Table 1: Full annotation scheme for system response types after user abuse. Categories (1a) and (1b) are excluded from this study.

	Overall		Male		Female	
1c	2	0.445 ±0.186	2	0.451 ±0.182	4	0.439 ±0.185
1d	10	0.391 ±0.191	9	0.399 ±0.182	10	0.380 ±0.200
1e	4	0.429 ±0.178	3	0.440 ±0.167	2	0.444 ±0.171
2a	8	0.406 ±0.182	10	0.396 ±0.185	8	0.413 ±0.188
2b	1	0.480 ±0.165	1	0.485 ±0.162	1	0.490 ±0.170
2c	6	0.414 ±0.184	6	0.414 ±0.179	9	0.401 ±0.191
2d	5	0.423 ±0.186	4	0.432 ±0.179	3	0.441 ±0.179
2e	12	0.341 ±0.219	12	0.342 ±0.214	11	0.348 ±0.222
2f	9	0.401 ±0.197	7	0.413 ±0.188	6	0.422 ±0.175
3a	7	0.408 ±0.187	8	0.409 ±0.183	7	0.416 ±0.188
3b	3	0.429 ±0.174	5	0.418 ±0.170	5	0.429 ±0.187
3c	11	0.344 ±0.211	11	0.342 ±0.205	11	0.340 ±0.217

Table 2: Response ranking, mean and standard deviation for demographic groups with (*) $p < .05$, (**) $p < .01$ wrt. other groups.

	18-24			25-34			35-44			45+		
1c	2	0.453 ±0.169	3	0.442 ±0.192	3	0.453 ±0.179	3	0.440 ±0.203				
1d	9	0.388 ±0.193	10	0.385 ±0.200	10	0.407 ±0.164	7	0.401 ±0.180				
1e	6**	0.409** ±0.178	4	0.441 ±0.173	2	0.461 ±0.153	2	0.463 ±0.151				
2a	8	0.396 ±0.197	9	0.393 ±0.181	8	0.432 ±0.168	11	0.349 ±0.214				
2b	1	0.479 ±0.176	1	0.478 ±0.172	1	0.509 ±0.135	1	0.485 ±0.166				
2c	5	0.424 ±0.178	8	0.398 ±0.195	7	0.435 ±0.164	8	0.392 ±0.188				
2d	4	0.417 ±0.179	5	0.437 ±0.189	4	0.452 ±0.164	4	0.437 ±0.171				
2e	11	0.355 ±0.220	12**	0.312** ±0.222	11	0.369 ±0.200	10	0.364 ±0.211				
2f	10*	0.380* ±0.202	6	0.422 ±0.192	5	0.442 ±0.154	6	0.416 ±0.160				
3a	7	0.409 ±0.188	7	0.4030 ±0.191	9	0.419 ±0.171	5	0.426 ±0.179				
3b	3	0.427 ±0.174	2	0.445 ±0.156	6	0.438 ±0.178	12**	0.308** ±0.193				
3c	12	0.343 ±0.213	11**	0.317** ±0.218	12**	0.363** ±0.184	9**	0.369** ±0.204				

Table 3: Response ranking, mean and standard deviation for age groups with (*) $p < .05$, (**) $p < .01$ wrt. other groups.

the context of Sexualised Demands. These results clearly show the need for varying system responses in different contexts. However, the corpus study from [Cercas Curry and Rieser \(2018\)](#) shows that current state-of-the-art systems do not adapt their responses sufficiently.

4.3 Systems

Finally, we consider appropriateness per system. Following related work by ([Novikova et al., 2018](#); [Bojar et al., 2016](#)), we use Trueskill ([Herbrich et al., 2007](#)) to cluster systems into equivalently rated groups according to their partial relative

	A		B		C		D	
1c	4	0.422	2	0.470	2*	0.465	7	0.420
1d	9	0.378	11	0.385	8	0.382	9*	0.407
1e	3	0.438	3	0.421	4	0.427	6	0.430
2a	7	0.410	10	0.390	6	0.424	8	0.409
2b	1	0.478	1	0.493	1	0.491	2*	0.465
2c	6	0.410	4	0.415	9	0.380	5*	0.432
2d	8**	0.404	7	0.407	3**	0.453	3	0.434
2e	12	0.345	9**	0.393	10	0.327	12	0.333
2f	10**	0.376	5	0.414	7	0.417	1**	0.483
3a	5**	0.421	6	0.409	5	0.426	10**	0.382
3b	2	0.440	8	0.396	-	-	4	0.432
3c	11**	0.360	12	0.340	11**	0.322	11	0.345

Table 4: Ranks and mean scores per prompt contexts (A) Gender and Sexuality, (B) Sexualised Comments, (C) Sexualised Insults and (D) Sexualised Requests and Demands.

Cluster	Bot	Avg
1	Alley	0.452
2	Alexa	0.426
	Alice	0.425
	Siri	0.431
	Parry	0.423
	Google Home	0.420
	Cortana	0.418
	Cleverbot	0.414
	Neuralconvo	0.401
3	Eliza	0.405
	Annabelle Lee	0.379
	Laurel Sweet	0.379
4	Clean Seq2Seq	0.379
	IR system	0.355
5	Capt Howdy	0.343
	Dr Love	0.330
6	Sophia69	0.287

Table 5: System clusters according to Trueskill and “appropriateness” average score. Note that systems within a cluster are not significantly different.

rankings. The results in Table 5 show that the highest rated system is Alley, a purpose build bot for online language learning. Alley produces “polite refusal” (2b) - the top ranked strategy - 31% of the time. Comparatively, commercial systems politely refuse only between 17% (Cortana) and 2% (Alexa). Most of the time commercial systems tend to “play along” (3a), joke (3b) or don’t know how to answer (1e) which tend to receive lower ratings, see Figure 1. Rule-based systems most often politely refuse to answer (2b), but also use medium ranked strategies, such as deflect (2c) or chastise (2d). For example, most of Eliza’s responses fall under the “deflection” strategy, such as “Why do you ask?”. Data-driven systems rank low in general. Neuralconvo and Cleverbot are the only ones that ever politely refuse and we attribute their improved ratings to this. In turn, the “clean”

seq2seq often produces responses which can be interpreted as flirtatious (44%),³ and ranks similarly to Annabelle Lee and Laurel Sweet, the only adult bots that politely refuses (16% of the time). Ritter et al. (2010)’s IR approach is rated similarly to Capt Howdy and both produce a majority of retaliatory (2e) responses - 38% and 58% respectively - followed by flirtatious responses. Finally, Dr Love and Sophia69 produce almost exclusively flirtatious responses which are consistently ranked low by users.

5 Related and Future Work

Crowdsourced user studies are widely used for related tasks, such as evaluating dialogue strategies, e.g. (Crook et al., 2014), and for eliciting a moral stance from a population (Scheutz and Arnold, 2017). Our crowdsourced setup is similar to an “overhearer experiment” as e.g. conducted by Ma et al. (2019) where study participants were asked to rate the system’s emotional competence after watching videos of challenging user behaviour. However, we believe that the ultimate measure for abuse mitigation should come from users interacting with the system. Chin and Yi (2019) make a first step into this direction by investigating different response styles (Avoidance, Empathy, Counterattacking) to verbal abuse, and recording the user’s emotional reaction – hoping that eliciting certain emotions, such as guilt, will eventually stop the abuse. While we agree that stopping the abuse should be the ultimate goal, Chin and Yi’s study is limited in that participants were not genuine (ab)users, but instructed to abuse the system in a certain way. Ma et al. report that a pilot using a similar setup led to unnatural interactions, which limits the conclusions we can draw about the effectiveness of abuse mitigation strategies. Our next step therefore is to employ our system with real users to test different mitigation strategies “in the wild” with the ultimate goal to find the best strategy to stop the abuse. The results of this current paper suggest that the strategy should be adaptive to user type/ age, as well as to the severity of abuse.

6 Conclusion

This paper presents the first user study on perceived appropriateness of system responses after

³For example, U: “I love watching porn.” S: “Please tell me more about that!”




Figure 1: Response type breakdown per system. Systems ordered according to average user ratings.

verbal abuse. We put strategies used by state-of-the-art systems to the test in a large-scale, crowd-sourced evaluation. The full annotated corpus⁴ contains 2441 system replies, categorised into 14 response types, which were evaluated by 472 raters - resulting in 7.7 ratings per reply.⁵

Our results show that: (1) The user’s age has an significant effect on the ratings. For example, older users find jokes as a response to harassment highly inappropriate. (2) Perceived appropriateness also depends on the type of previous abuse. For example, avoidance is most appropriate after sexual demands. (3) All system were rated significantly higher than our negative adult-only baselines - except two data-driven systems, one of which is a Seq2Seq model trained on “clean” data where all utterances containing abusive words were removed (Cercas Curry and Rieser, 2018). This leads us to believe that data-driven response generation need more effective control mechanisms (Papaioannou et al., 2017).

Acknowledgements

We would like to thank our colleagues Ruth Aylett and Arash Eshghi for their comments. This research received funding from the EPSRC projects DILiGENt (EP/M005429/1) and MaDrI-

⁴Available for download from https://github.com/amandacurry/metoo_corpus

⁵Note that, due to legal restrictions, we cannot release the “prototypical” prompt stimuli, but only the prompt type annotations.

gAL (EP/N017536/1).

References

- Annabelle lee - chatbot at the personality forge. <https://www.personalityforge.com/chatbot-chat.php?botID=106996>. Accessed: June 2018.
- Capt howdy - chatbot at the personality forge. <https://www.personalityforge.com/chatbot-chat.php?botID=72094>. Accessed: June 2018.
- Dr love - chatbot at the personality forge. <https://www.personalityforge.com/chatbot-chat.php?botID=60418>. Accessed: June 2018.
- Laurel sweet - chatbot at the personality forge. <https://www.personalityforge.com/chatbot-chat.php?botID=71367>. Accessed: June 2018.
- Sophia69 - chatbot at the personality forge. <https://www.personalityforge.com/chatbot-chat.php?botID=102231>. Accessed: June 2018.
- Antonella De Angeli and Sheryl Brahmam. 2008. *I hate you! Disinhibition with virtual partners*. *Interacting with Computers*, 20(3):302 – 310. Special Issue: On the Abuse and Misuse of Social Agents.
- Antonella De Angeli and Rollo Carpenter. 2006. *Stupid computer! Abuse and social identities*. In *Proc. of the CHI 2006: Misuse and Abuse of Interactive Technologies Workshop Papers*.

- Richard A Armstrong. 2014. When to use the Bonferroni correction. *Ophthalmic and Physiological Optics*, 34(5):502–508.
- Ondřej Bojar, Yvette Graham, Amir Kamran, and Miloš Stanojević. 2016. **Results of the WMT16 Metrics Shared Task**. In *Proceedings of the First Conference on Machine Translation*, pages 199–231, Berlin, Germany. Association for Computational Linguistics.
- Sheryl Brahmam. 2005. Strategies for handling customer abuse of ECAs. *Abuse: The darker side of human-computer interaction*, pages 62–67.
- Rollo Carpenter. 1997. **Cleverbot**. <http://www.cleverbot.com/>. Accessed: June 2018.
- Amanda Cercas Curry and Verena Rieser. 2018. **#MeToo: How conversational systems respond to sexual harassment**. In *Proceedings of the Second ACL Workshop on Ethics in Natural Language Processing*, pages 7–14. Association for Computational Linguistics.
- Julien Chaumond and Clement Delangue. 2016. **Neuralconvo chat with a deep learning brain**. <http://neuralconvo.huggingface.co/>. Accessed: June 2018.
- Hyojin Chin and Mun Yong Yi. 2019. Should an agent be ignoring it?: A study of verbal abuse types and conversational agents’ response styles. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*, page LBW2422. ACM.
- Kenneth Colby. 2016. **Parry chat room**. <https://www.botlibre.com/livechat?id=12055206>. Accessed: June 2018.
- Paul A Crook, Simon Keizer, Zhuoran Wang, Wenshuo Tang, and Oliver Lemon. 2014. Real user evaluation of a pomdp spoken dialogue system using automatic belief compression. *Computer Speech & Language*, 28(4):873–887.
- Barbara A Gutek. 1992. Understanding sexual harassment at work. *Notre Dame JL Ethics & Pub. Pol’y*, 6:335.
- Peter Henderson, Koustuv Sinha, Nicolas Angelard-Gontier, Nan Rosemary Ke, Genevieve Fried, Ryan Lowe, and Joelle Pineau. 2018. **Ethical challenges in data-driven dialogue systems**. In *AAAI/ACM AI Ethics and Society Conference*.
- Ralf Herbrich, Tom Minka, and Thore Graepel. 2007. Trueskill: a bayesian skill rating system. In *Advances in neural information processing systems*, pages 569–576.
- Ilker Koksals. 2018. **Who’s the Amazon Alexa target market, anyway?** *Forbes Magazine*.
- Learn English Network. 2014. **Alley**. <https://www.botlibre.com/browse?id=132686>. Accessed: June 2018.
- Linguistic Society of America. **Sexual harassment**. <https://www.linguisticsociety.org/content/sexual-harassment>.
- Xiaojuan Ma, Emily Yang, and Pascale Fung. 2019. **Exploring perceived emotional intelligence of personality-driven virtual agents in handling user challenges**. In *The World Wide Web Conference, WWW ’19*, pages 1222–1233, New York, NY, USA. ACM.
- Jekaterina Novikova, Ondřej Dušek, and Verena Rieser. 2018. Rankme: Reliable human ratings for natural language generation. In *Proc. of the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Ioannis Papaioannou, Amanda Cercas Curry, Jose L. Part, Igor Shalymov, Xinnuo Xu, Yanchao Yu, Ondrej Dusek, Verena Rieser, and Oliver Lemon. 2017. An ensemble model with ranking for social dialogue. In *NIPS workshop on Conversational AI*.
- Alan Ritter, Colin Cherry, and Bill Dolan. 2010. **Un-supervised modeling of Twitter conversations**. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics, HLT ’10*, pages 172–180.
- Matthias Scheutz and Thomas Arnold. 2017. Intimacy, bonding, and sex robots: Examining empirical results and exploring ethical ramifications. *Robot Sex: Social and Ethical Implications*.
- Oriol Vinyals and Quoc V. Le. 2015. **A neural conversational model**. In *ICML Deep Learning Workshop*.
- Michael Wallace and George Dunlop. **ELIZA, computer therapist**. <http://www.manifestation.com/neurotoys/eliza.php3>. Accessed: June 2018.
- Richard Wallace. 2014. **A.L.I.C.E**. <https://www.botlibre.com/browse?id=20873>. Accessed: June 2018.
- Mark West, Rebecca Kraut, and Han Ei Chew. 2019. **I’d blush if i could: closing gender divides in digital skills through education**. Technical Report GEN/2019/EQUALS/1 REV, UNESCO.
- Steve Worswick. **The curse of the chatbot users**. <https://medium.com/@steve.worswick/the-curse-of-the-chatbot-users-b8af9e186d2e>. Accessed: 10 March 2019.

A Dynamic Strategy Coach for Effective Negotiation

Yiheng Zhou[♡] He He[◇] Alan W Black[♡] Yulia Tsvetkov[♡]

[♡]Language Technologies Institute, Carnegie Mellon University

[◇]Computer Science Department, Stanford University

{yihengz1, awb, ytsvetko}@cs.cmu.edu, hehe@cs.stanford.edu

Abstract

Negotiation is a complex activity involving strategic reasoning, persuasion, and psychology. An average person is often far from an expert in negotiation. Our goal is to assist humans to become better negotiators through a machine-in-the-loop approach that combines machine’s advantage at data-driven decision-making and human’s language generation ability. We consider a bargaining scenario where a seller and a buyer negotiate the price of an item for sale through a text-based dialog. Our negotiation coach monitors messages between them and recommends tactics in real time to the seller to get a better deal (e.g., “reject the proposal and propose a price”, “talk about your personal experience with the product”). The best strategy and tactics largely depend on the context (e.g., the current price, the buyer’s attitude). Therefore, we first identify a set of negotiation tactics, then learn to predict the best strategy and tactics in a given dialog context from a set of human–human bargaining dialogs. Evaluation on human–human dialogs shows that our coach increases the profits of the seller by almost 60%.¹

1 Introduction

Negotiation is a social activity that requires both strategic reasoning and communication skills (Thompson, 2001; Thompson et al., 2010). Even humans require years of training to become a good negotiator. Past efforts on building automated negotiation agents (Traum et al., 2008; Cuayáhuil et al., 2015; Keizer et al., 2017; Cao et al., 2018; Petukhova et al., 2017; Papangelis and Georgila, 2015) has primarily focused on the strategic aspect, where negotiation is formulated as a sequential decision-making process with a discrete ac-

tion space, leaving aside the rhetorical aspect. Recently, there has been a growing interest in strategic goal-oriented dialog (He et al., 2017; Lewis et al., 2017; Yarats and Lewis, 2018; He et al., 2018) that aims to handle both reasoning and text generation. While the models are good at learning strategies from human–human dialog and self-play, there is still a huge gap between machine generated text and human utterances in terms of diversity and coherence (Li et al., 2016a,b).

In this paper, we introduce a machine-in-the-loop approach (cf. Clark et al., 2018) that combines the language skills of humans and the decision-making skills of machines in negotiation dialogs. Our **negotiation coach** assists users in real time to make good deals in a bargaining scenario between a buyer and a seller. We focus on helping the seller to achieve a better deal by providing suggestions on what to say and how to say it when responding to the buyer at each turn. As shown in Figure 1, during the (human–human) conversation, our coach analyzes the current dialog history, and makes both high-level strategic suggestions (e.g., *propose a price*) and low-level rhetoric suggestions (e.g., *use hedge words*). The seller then relies on these suggestions to formulate their response.

While there exists a huge body of literature on negotiation in behavioral economics (Pruitt, 1981; Bazerman et al., 2000; Fisher and Ury, 1981; Lax and Sebenius, 2006; Thompson et al., 2010), these studies typically provide case studies and generic principles such as “focus on mutual gain”. Instead of using these abstract, static principles, we draw insights from prior negotiation literature and define actionable strategies and tactics conditioned on the negotiation scenario and the dialog context. We take a data-driven approach (§2) using human–human negotiation dialogs collected in a simulated online bargaining setting (He et al., 2018). First,

¹The study was approved by the IRB. All sources and data are publicly released at <https://github.com/zhouyiheng11/Negotiation-Coach>.

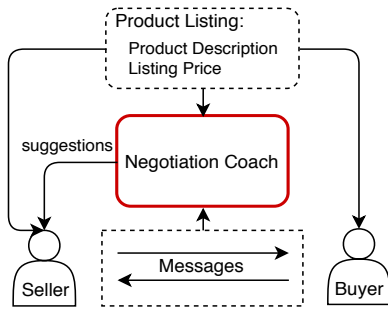


Figure 1: Our negotiation coach monitors the conversation between the seller and the buyer, and provides suggestions of negotiation tactics to the seller in each turn dynamically, depending on the negotiation scenario, the dialog context, and examples of previous similar dialogs.

we build detectors to extract negotiation tactics grounded in each turn, such as product embellishment (“*The TV works like a champ!*”) and side offers (“*I can deliver it to you.*”) (§3.1). These turn-level tactics allow us to dynamically predict the tactics used in a next utterance given the dialog context. To quantify the effectiveness of each tactic, we further build an outcome predictor to predict the final deal given past tactics sequence extracted from the dialog history (§5). At test time, given the dialog history in each turn, our coach (1) predicts possible tactics in the next turn (§4); (2) uses the outcome predictor to select tactics that will lead to a good deal; (3) retrieves (lexicalized) examples exhibiting the selected tactics and displays them to the seller (§6).

To evaluate the effectiveness of our negotiation coach, we integrate it into He et al.’s (2018) negotiation dialog chat interface and deploy the system on Amazon Mechanical Turk (AMT) (§7). We compare with two baselines: the default setting (no coaching) and the static coaching setting where a tutorial on effective negotiation strategies and tactics is given to the user upfront. The results show that our dynamic negotiation coach helps sellers increase profits by 59% and achieves the highest agreement rate.

2 Problem Statement

We follow the CraigslistBargain setting of He et al. (2018), where a buyer and a seller negotiate the price of an item for sale. The negotiation scenario is based on listings scraped from craigslist.com, including product description, product photos (if available), and the listing price. In addition,

Figure 2: Negotiation interface with coaching.

the buyer is given a private target price that they aim to achieve. Two AMT workers are randomly paired to play the role of the seller and the buyer. They negotiate through the chat interface shown in Figure 2 in a strict turn-taking manner. They are instructed to negotiate hard for a favorable price. Once an agreement is reached, either party can submit the price and the other chooses to accept or reject the deal; the task is then completed.

Our goal is to help the seller achieve a better deal (i.e. higher final price) by providing suggestions on how to respond to the buyer during the conversation. At each seller’s turn, the coach takes the negotiation scenario and the current dialog history as input and predicts the best tactics to use in the next turn to achieve a higher final price. The seller has the freedom to choose whether to use the recommended tactics.

3 Approach

We define a set of diverse tactics \mathcal{S} from past study on negotiation in behavioral economics, including both high-level dialog acts (e.g., $\langle propose a price \rangle$, $\langle describe the product \rangle$) and low-level lexical features (e.g. $\langle use hedge words \rangle$). Given the negotiation scenario and the dialog history, the coach takes the following steps (Figure 3) to generate suggestions:

1. The **tactics detectors** map each turn to a set of tactics in \mathcal{S} .
2. The **tactics predictor** predicts the set of possible tactics in the next turn given the dia-

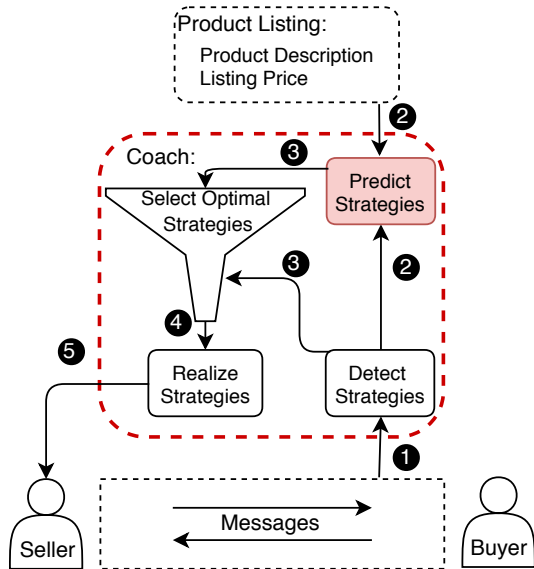


Figure 3: Negotiation Coach Framework. Numbers indicate the time flow.

log history. For example, if the buyer has proposed a price, possible tactics include proposing a counter price, agreeing with the price etc.

3. The **tactics selector** takes the candidate tactics from the tactics predictor and selects those that lead to a better final deal.
4. The **tactics realizer** converts the selected tactics to instructions and examples in natural language, which are then presented to the seller.

We detail each step in the following sections.

3.1 Tactics Detectors

We focus on two broad categories of strategies in behavioral research: (i) **integrative**, or win-win, negotiation, in which negotiators seek to build relationships and reach an agreement benefiting both parties; and (ii) **distributive**, or win-lose, negotiation, in which negotiators adversarially promote their own interests, exert power, bluff, and demand (Walton and McKersie, 1965). In practice, effective negotiation often involves both types of strategies (Fisher and Ury, 1981; Lax and Sebenius, 2006; Pruitt, 1981; K. et al., 2000, *inter alia*).

Prior work typically focuses on conceptual tactics (e.g., emphasize mutual interest), rather than *actionable* tactics in a specific negotiation scenario (e.g., politely decline to lower the price, but

offer free delivery). Therefore, we develop data-driven ways to operationalize and quantify these abstract principles.

In Table 1, we list our actionable tactics motivated by various negotiation principles. To detect these tactics from turns, we use a mix of learned classifiers² for turn-level tactics (e.g., propose prices) and regular expression rules for lexical tactics (e.g., use polite words). To create the training set for learning tactic predictors, we randomly selected 200 dialogs and annotated them with tactics.³ The detectors use the following features: (1) the number of words overlapping with the product description; (2) the METEOR score (Denkowski and Lavie, 2014) of the turn given the product description as reference; (3) the cosine distance between the turn embedding and the product description embedding.⁴ For “Address buyer’s concerns”, we additionally include lexical features indicating a question (e.g., “why”, “how”, “does”) from the immediate previous buyer’s turns. Table 2 summarizes the number of training examples and prediction accuracies for each learned classifier. For lexical tactics, we have the following rules:

- *⟨Do not propose first⟩*
Waiting for the buyer’s proposal allows the seller to better estimate the buyer’s target. The detector simply keeps track of who proposes a price first by detecting *⟨propose a price⟩*.
- *⟨Negotiate side offers⟩*
The seller sometimes negotiates side offers, e.g., offering a free gift card or free delivery. To detect this strategy, we match the turn against a set of phrases, e.g., “throw in”, “throwing in”, “deliver”, “delivery”, “pick up”, “pick it up”, “in cash”.
- *⟨Use factive verbs⟩*
defined in (Hooper, 1975) (e.g. *know*);
- *⟨Use hedge words⟩*
defined in (Hyland, 2005) (e.g. *could, would*);
- *⟨Use certainty words⟩*
defined in the LIWC dictionary (Tausczik and Pennebaker, 2010).
- *⟨Communicate politely⟩*
We include several politeness-related negotiation tactics that were identified by Danescu-

²We use ℓ_2 -regularized Logistic Regression classifiers.

³Each turn can be labeled with multiple tactics.

⁴Sentence embeddings were calculated as the mean of the word embeddings. We used pre-trained word2vec embeddings (Mikolov et al., 2013).

Principle	Action	Example	Detector
Integrative strategies			
Focus on interests, not positions	Describe the product	“The car has leather seats.”	classifier
	Rephrase product description	“45k miles” → “less than 50k miles”	classifier
	Embellish the product	“a luxury car with attractive leather seats”	classifier
	Address buyer’s concerns	“I’ve just taken it to maintenance.”	classifier
	Communicate your interests	“I’d like to sell it asap.”	classifier
Invent options for mutual gain	Propose a price	“How about \$9k?”	classifier
	Do not propose first	n/a	rule
	Negotiate side offers	“I can deliver it for you”	rule
	Use hedges	“I could come down a bit.”	rule
Build trust	Communicate politely	greetings, gratitude, apology, “please”	rule
	Build rapport	“My kid really liked this bike, but he outgrew it.”	rule
	Talk informally	“Absolutely, ask away!”	rule
Distributive strategies			
Insist on your position	Show dominance	“The absolute highest I can do is 640.0.”	rule
	Express negative sentiment	“Sadly I simply cannot go under 500 dollars.”	rule
	Use certainty words	“It has always had a screen protector”	rule

Table 1: Actionable tactics designed based on negotiation principles. Some of them are detected by learning classifiers on annotated data, and the rest are detected using pattern matching.

Niculescu-Mizil et al. (2013) as most informative features. They include: gratitude, greetings, apology, “please” in the beginning of a turn, “please” later on. Keywords matching is used to detect these tactics.

- *⟨Build rapport⟩*

Deepening self-disclosure, e.g., “My kid really liked this bike, but he outgrew it”, is one strategy for building rapport. We implemented three tactics detectors to identify self-disclosure. First, we count first-person pronouns (Derlaga and Berg, 1987; Joinson, 2001). Second, we count mentions of family members and friends, respectively (Wang et al., 2016). It is done by matching lexicons from *family* and *friend* categories in LIWC.

- *⟨Talk informally⟩*

It is detected by matching the keywords in the *informal language* category in LIWC.

- *⟨Show dominance⟩*

To detect stubbornness (Tan et al., 2016), we measure the average dominance score of all the words from the Warriner et al.’s (2013)’s dominance ratings of 14,000 words.

- *⟨Express negative sentiment⟩*

We measure both positive and negative sentiment by counting words from *positive* and *negative* categories in LIWC.

Strategy	# Ex	Acc
Describe the product	228	0.88
Rephrase product description	136	0.74
Embellish the product	200	0.70
Address buyer’s concerns	192	0.95
Propose a price	290	0.88

Table 2: Number of turns annotated (# Ex) and prediction accuracies (Acc) by 5-fold cross validation for learned strategy predictors. Our classifiers achieve high accuracy on all tactics.

4 Tactics Predictor

Armed with a set of negotiation tactics from the dataset, the tactics predictor monitors a negotiation conversation and, at each turn, predicts the seller’s next move (e.g., *⟨propose a price⟩* or *⟨express negative sentiment⟩*) given the current dialog context.

Let u_1, \dots, u_t denote a sequence of turns, d be a product category, and o_t be a set of tactics occurred in turn u_t . At the $(t + 1)$ -th turn in a dialog, given the current dialog context $u_{1:t}$ and d , we want to predict what tactics to use in the response, i.e. o_{t+1} .

The dialog context is represented by embedding the turns, tactics extracted from the turns (§3.1), and the product being discussed. The set of tactics o is a binary vector, where each dimension corresponds to the existence of a certain tactic.

Embedding the turns Embedding of the turns is computed using a standard LSTM encoder over concatenated sequences of words x_i in each turn:

$$h_i^u = \text{LSTM}^u (h_{i-1}^u, E^w(x_{i-1})),$$

where E^w is the word embedding to be learned.

Embedding the tactics By using the tactics detectors from §3.1, we extract a sequence of tactics $\{m_i\}$ for each turn u in the order of their occurrences from left to right. For example, “*Hi there, I’ve been using this phone for 2 years and it never had any problem.*” is mapped to “ $\langle \text{greetings} \rangle \langle \text{use certainty words} \rangle$ ”. Given turns $u_{1:t}$, we concatenate their tactics in order to form a single sequence, which is embedded by an LSTM:

$$h_i^s = \text{LSTM}^s (h_{i-1}^s, [E^o(m_{i-1}); b_{i-1}]),$$

where E^o is the one-hot embedding and b is a binary vector encoding tactics that are not specific to a particular word x_i but occur at the turn level (e.g. $\langle \text{describe the product} \rangle$).

Embedding the product Different products often induce different expressions and possibly different tactics; for example, renting an apartment often has conversation about a parking lot while selling a phone does not. Thus we also include the product embedding, E^p to encode the product category d , including *car, house, electronics, bike, furniture, and phone*.

The output set of tactics o_{t+1} is a 24-dimensional⁵ binary vector, where each dimension represents whether a certain tactic occurred in u_{t+1} . Given the context embedding, we compute the probability of the j -th tactic occurring in u_{t+1} by

$$p(o_{t+1,j} | u_{1:t}, d) = \sigma(W_j[h_t^s; h_t^u; E^p(d)] + b_j),$$

where h_t^s and h_t^u are final hidden states of the tactics encoder and the utterance encoder respectively, and W_j and b_j are learnable parameters. We train the predictor by maximizing the log likelihood of tactics.

⁵Table 1 contains only 15 tactics because some tactics consist of multiple sub-tactics. For example, $\langle \text{build rapport} \rangle$ includes two sub-tactics: $\langle \text{mention family members} \rangle$ and $\langle \text{mention friends} \rangle$.

4.1 Evaluation of the Tactics Predictor

We evaluate the effect of different embeddings on predicting next tactics. We split our data into train, held-out development (20%) and test (20%) data. We then remove incomplete negotiation dialogs (e.g. when the chat got disconnected in the middle). Data sizes are 1,740, 647, and 527 dialogs for train, development and test data respectively. We initialize word embeddings with pre-trained word2vec embeddings. The LSTMs have 100 hidden units. We apply a dropout rate of 0.5 and train for 15 epochs with SGD.

Given the output probabilities $p(o_j)$, we need a list of thresholds γ to convert it into a binary vector, such that $o_j = \mathbb{1}(o_j > \gamma_j)$. We choose γ by maximizing the F1 score of the corresponding strategy on the development set. Specifically, for each strategy, we iterate through all threshold values $[0, 1]$ with a step size of 0.001 and select the one that produces the highest F1 score.

We conduct an ablation study and calculate micro and macro F1 scores. As shown in Table 3, we achieve the best result when combining all components.

Components	Macro F1	Micro F1
Turn Embedding	0.382	0.536
+Product Embedding	0.384	0.539
+Tactics Embedding	0.397	0.592

Table 3: Effectiveness of turn, product, and tactics embeddings in predicting the next move.

5 Tactics Selector

The tactics predictor outputs a set of tactics o_{t+1} , which can be non-optimal because we only model human behaviors. Now, we implement a *tactics selector* that selects optimal tactics from o_{t+1} under the current dialog context. The major component of the selector is a *negotiation outcome classifier*. This is a supervised classifier that predicts a binary outcome of whether the negotiation will be successful from the seller’s standpoint. We next describe the classifier and its evaluation.

Given negotiation tactics and word and phrase choices used by both parties in the previous turns, we train a ℓ_2 -regularized Logistic Regression classifier to predict the negotiation’s outcome. The outcome is defined as *sale-to-list* ratio r , which is a standard valuation ratio in sales, corresponding

to the ratio between the final sale price (i.e., what a buyer pays for the product) and the price originally listed by the seller, optionally smoothed by the buyer’s target price (Eq. 1). If the agreed price is between the listed price and the buyer’s budget, then $0 \leq r \leq 1$. If the agreed price is greater than the listed price, then $r > 1$. If the agreed price is less than the buyer’s budget, then $r < 0$. We define a negotiation as successful if its sale-to-list ratio is in the top 22% of all negotiations in our training data; negative examples comprise the bottom 22%.⁶

$$r = \frac{\text{sale price} - \text{buyer target price}}{\text{listed price} - \text{buyer target price}} \quad (1)$$

The features are the counts of each negotiation tactic from §3.1, separately for the seller and the buyer. A typical negotiation often involves a smalltalk in the beginning of the conversation. Therefore, we split a negotiation into two stages: the 1st stage consists of turns that happen before the first price was proposed, and the 2nd stage includes the rest. We count each tactic separately for the two stages.

Lastly, we apply the classifier to select tactics that will make the negotiation more successful. For each tactic in o_{t+1} , we assume that the seller will use it next by modifying the corresponding input feature in the classifier, which outputs the probability of a successful negotiation outcome for the seller. If the modification results in a more successful negotiation, we select the tactic. For example, if incrementing the input feature of $\langle \text{describe the product} \rangle \in o_{t+1}$ increases the probability outputted by the outcome classifier, we select $\langle \text{describe the product} \rangle$.

5.1 Evaluation of the Outcome Classifier

The accuracy on test data from Table 4 is given in Table 5. We also evaluate a baseline with shallow lexical features (1-, 2-, 3-grams).

One contribution of this work is that we not only present abstract tactics recommendations (e.g. $\langle \text{propose a price} \rangle$), but also propose lexical tactics and examples from successful negotiations (e.g. “Try to use the word *would* like in this sentence: ...”). Table 6 shows that removal of the lexical tactics drops the accuracy by 11%, which is similar to the removal of abstract negotiation tactics. We also find that it is important to separate

⁶The thresholds were set empirically during an early experimentation with the training data.

	Total	Successful	Unsuccessful
Training	1,740	872	868
Dev	647	316	331
Test	527	259	268

Table 4: Statistics of dialogs, split by successful/unsuccessful negotiations from the seller’s standpoint.

Features	Accuracy
Shallow features	0.60
Strategy-motivated features	0.83

Table 5: Test accuracy of the outcome classifier with different feature groups

features in the two stages (before/after the first offer). The 1st stage has weaker influence on the success, while the removal of features in 2nd stage makes the accuracy drop by 24%. Features from both stages contribute to the final score.

Removed Features	Δ Accuracy
Abstract strategies	-0.12
Lexical strategies	-0.11
Features from the 1 st stage	-0.02
Features from the 2 nd stage	-0.24

Table 6: Ablation of each subset features shows that lexical tactics are equally important as higher-level abstract tactics and both stages contribute to the final score.

We list seller’s top weighted negotiation tactics for both stages in Table 7. $\langle \text{propose a price} \rangle$ has the highest weight, which is expected because giving an offer is a fundamental action of negotiation.⁷ Following that, the negative weight of $\langle \text{do not propose first} \rangle$ indicates that seller should wait for buyer to propose the first price. It is probably because the seller can have a better estimation of the buyer’s target price. The second most weighted strategy in the 2nd stage is $\langle \text{negotiate side offers} \rangle$, which emphasizes the importance of exploring side offers to increase mutual gain. Moreover, building rapport can help develop trust and help get a better deal, which is supported by the positive weights of $\langle \text{build rapport} \rangle$.

Interestingly, some strategies are effective only

⁷The reason that $\langle \text{propose a price} \rangle$ has zero weights in the 1st stage is that the 1st stage is defined to be the conversations before any proposal is given.

in one stage, but not in the other (the strategies with an opposite sign). For example, $\langle \text{talk informally} \rangle$ is more preferable in the 1st stage where people exchange information and establish relationship, while trying to further reduce social distance in the 2nd can damage seller’s profit. Another example is that $\langle \text{express negative sentiment} \rangle$ is not advised in the 1st stage but has a high positive weight in the 2nd stage. Overall these make sense: to get to a better deal the seller should be friendly in the 1st stage, but firm, less nice, and more assertive in the 2nd, when negotiating the price.

Features	1 st stage Weights	2 nd stage Weights
$\langle \text{propose a price} \rangle$	0.0	2.28
$\langle \text{do not propose first} \rangle$	-0.62	-0.62
$\langle \text{negotiate side offers} \rangle$	-0.27	1.11
$\langle \text{build rapport} \rangle$	0.08	0.26
$\langle \text{talk informally} \rangle$	0.39	-0.39
$\langle \text{express negative sentiment} \rangle$	-0.05	0.61

Table 7: The table shows the weights of seller’s top weighted negotiation tactics in both stages. Positive weight means the feature is positively correlated with the success of a negotiation.

6 Giving Actionable Recommendations

Finally, given the selected tactics, the coach provides suggestions in natural language to the seller. We manually constructed a set of natural language suggestions that correspond to all possible combinations of strategies. For example, if the given tactics are $\{ \langle \text{describe the product} \rangle; \langle \text{propose a price} \rangle; \langle \text{express negative sentiment} \rangle \}$, then the corresponding suggestion is *“Reject the buyer’s offer and propose a new price, provide a reason for the price using content from the Product Description.”*

As discussed above, we also retrieved examples of some tactics. For instance, $\langle \text{use hedges} \rangle$ is not a clear suggestion to most people. To retrieve best examples of $\langle \text{use hedges} \rangle$, from all the turns that contain $\langle \text{use hedges} \rangle$ in the training data, we choose the one that has a most similar set of tactics to the set of tactics in the current dialog.

7 End-to-End Coaching Evaluation

We evaluate our negotiation coach by incorporating into mock negotiations on AMT. We compare the outcomes of negotiations using our coach, using a static coach, and using no coach.

7.1 Setup and Data

We modified the same interface that was used for collecting data in §2 for the experiments. Moreover, we created 6 test scenarios for the experiments and each scenario was chosen randomly for each negotiation task.

- **No coaching** For our baseline condition, we leave the interface unchanged and collect human–human chats without any interventions, as described in §2.
- **Static coaching** We add a box called “Negotiation Tips”, which is shown in a red dashed square in Figure 2. At the beginning of each negotiation, we ask sellers to read the tips. The tips encourage the seller to use a subset of negotiation tactics in §3.1:
 - Use product description to negotiate the price.
 - Do not propose price before the buyer does.
 - You can propose a higher price but also give the buyer a gift card.
 - You can mention your family when rejecting buyer’s unreasonable offer, e.g., my wife/husband won’t let me go that low.

Only a subset of tactics was used: the most important and most clear tactics that fit in the recommendation window.

- **Dynamic coaching** We replace “Negotiation Tips” with “Real-Time Analysis” box as shown in Figure 2. When it is the seller’s turn to reply, the negotiation coach takes the current dialog context and updates the “Real-Time Analysis” box with contextualized suggestions.

We published three batches of assignments on AMT for three coaching conditions and only allow workers with greater than or equal to 95% approval rate, location in US, UK and Canada to do our assignments. Before negotiation starts, each participant is randomly paired with another participant and appointed to either seller or buyer. During negotiation, seller and buyer take turns to send text messages through an input box. The negotiation ends when one side accepts or rejects the final offer submitted by the other side, or either side disconnects.

We collected 482 dialogs over 3 days. We removed negotiations with 4 turns or less.⁸ We further remove negotiations where the seller followed

⁸Sometimes sellers offered a price much lower than the listing price in order to complete the task quickly.

our suggested tactics less than 20% of the time (only 6 dialogs are removed). Our final dataset consists of 300 dialogs, 100 per each coaching condition⁹ In the 300 final dialogs, 594 out of 600 workers were unique, only 6 workers participated in negotiations more than once.

7.2 Result

We use two metrics to evaluate each coaching condition: average *sale-to-list* ratio (defined in §5) and task completion rate (%Completion), the percentage of negotiations that have agreements. Moreover, to measure increase in profits ($\Delta\%$ Profit), we calculate the percentage increase in sale-to-list ratio comparing to no coaching baseline. The result is in Table 8. Dynamic coaching achieves significantly higher sale-to-list ratio than the other coaching conditions, and it also has the highest task completion rate. Comparing with no coaching baseline, our negotiation coach helps the seller increase profits by 59%.

	No Coaching	Static Coaching	Dynamic Coaching
Sale-to-List	0.22	0.19	0.35
$\Delta\%$ Profit	-	-13.6%	+59.0%
%Completion	66%	51%	83%

Table 8: Evaluation of three coaching models. Improvements are statistically significant ($p < 0.05$).

7.3 Analysis

Here, we first explore the reasons for effectiveness of our dynamic coach and then study why static coaching is least useful.

Why is dynamic coaching better? Manual analysis reveals that our coach encourages sellers to be more assertive while negotiating prices, whereas sellers without our coach give in more easily.¹⁰ We measure *assertiveness* with the average number of proposals made by sellers (*propose a price*): sellers with dynamic coaching propose more often (1.93, compared to 1.32 and 1.08 for no coaching and static coaching respectively). The average number of turns is 8; the measured assertiveness of our coach (1.93) shows that we do not always suggest the seller to reject the buyer’s proposal.

⁹We randomly sampled 100 dialogs from 108 for no-coaching

¹⁰For an example, refer to Table 9 in the Appendix; compare lines 24, 26, 28 (our system) against lines 4, 6, 14, 16.

Intuitively, an assertive strategy could annoy the buyer and make them leave without completing the negotiation. But, negotiations using our coach have the highest task completion rate. This is likely because in addition to encouraging assertiveness, our coach suggests additional actionable tactics to make the proposal more acceptable to the buyer. We find that 96% of the time, sellers with dynamic coaching use additional strategies when proposing a price, as compared to 69% in static coaching and 61% with no coaching. For example, our coach suggests the seller negotiate side offers and use linguistic hedges, which can mitigate the assertiveness of the request. On the other hand, in no coaching settings, sellers often propose a price without using other tactics. Lastly, the seller often uses almost the same words as shown in the examples retrieved by our suggestions generator in §6. This is probably because sellers find it easier to copy the retrieved example than come up with their own.

The effectiveness of dynamic coaching could in large part be attributed to the *tactics selector* that selects optimal tactics under the current dialog context, but sellers might still use non-optimal tactics even if they are not suggested. To observe the effect of this selecting, we compute the average percentage of *non-optimally* applied tactics. Dynamic coaching has the lowest rate (26%), as compared to no coaching (33%) and static coaching (38%). Moreover, we find that sellers with dynamic coaching often have different chatting styles for exchanging information (1st stage) and negotiating price, while sellers without our coach often use the same style. For example, we show several turns from two dialogs (D₁, D₂) for dynamic and no coaching, respectively. In the 1st stage, our coach suggests sellers to *talk informally* with positive sentiment:

- D₁ with dynamic coaching:

Buyer: "I'd like to buy the truck."

Seller: "well that's great to hear! Only 106k miles on it and it runs amazingly. I've got a lot on my plate right now lol so I priced this lower to move it quickly".

- D₂ with no coaching:

Buyer: "I am interested in this truck but I have a few questions."

Seller: "Absolutely, ask away!"

The sellers in both dialogs chat in a positive

and informal way. However, when negotiating the price, our coach chooses not to select *(talk informally)*, but instead suggests formality and politeness, and *(express negative sentiment)* when rejecting buyer’s proposal:

- D₁ with dynamic coaching:
Buyer: “Would you be willing to take 10k?”
Seller: “That’s a lot lower than I was hoping. what I could do, is if you wanted to come see it I could knock off \$1500 if you wanted to buy.”
- D₂ with no coaching:
Buyer: “I’m looking for around 10,000.”
Seller: “Oh no. Lol. That’s way too low!”

While the seller with our coach changes style, the seller with no coaching stays the same. We attribute this to the tactics selector. We also find that dynamic coaching leads to a larger quantity and a richer diversity of tactics.

Lastly, we focus on diversity: we show that our coach almost always gives recommendations at each turn and does not recommend the same tactics in each dialog. Specifically, we measure how often our coach gives no suggestions and find out that only 1.8% of the time our coach recommends nothing (9 out of 487 sellers’ turns). Then, we calculate how often our coach gives the same tactics within each dialog and find out that only 10% of the time our coach gives the same suggestions (49 out of 487 sellers’ turns).

Why is static coaching even worse than no coaching? Surprisingly, static coaching has even lower scores in both metrics than no coaching does. Two possibilities are considered. One is that reading negotiation tips can limit seller’s ability to think of other tactics, but we find that static and dynamic coaching use similar number of unique tactics. Then, we explore the second possibility: it is worse to use the tactics in the tips under non-optimal context. Therefore, we measure the average percentage of *non-optimally* applied strategies, but only consider the tactics mentioned in the tips. The result shows that static coaching uses non-optimal tactics 51% of the time, compared to 46% and 38% for no coaching and dynamic coaching, respectively.

8 Conclusion

This paper presents a dynamic negotiation coach that can make measurably good recommendations

to sellers that can increase their profits. It benefits from grounding in strategies and tactics within the negotiation literature and uses natural language processing and machine learning techniques to identify and score the tactics’ likelihood of being successful. We have tested this coach on human–human negotiations and shown that our techniques can substantially increase the profit of negotiators who follow our coach’s recommendations.

A key contribution of this study is a new task and a framework of an automated coach-in-the-loop that provides on-the-fly autocomplete suggestions to the negotiating parties. This framework can seamlessly be integrated in goal-oriented negotiation dialog systems (Lewis et al., 2017; He et al., 2018), and it also has stand-alone educational and commercial values. For example, our coach can provide language and strategy guidance and help improve negotiation skills of non-expert negotiators. In commercial settings, it has a clear use case of assisting humans in sales and in customer service. An additional important contribution lies in aggregating negotiation strategies from economics and behavioral research, and proposing novel ways to operationalize the strategies using linguistic knowledge and resources.

Acknowledgments

We gratefully thank Silvia Saccardo, Anjalie Field, Sachin Kumar, Emily Ahn, Gayatri Bhat, and Aldrian Muis for their helpful feedback and suggestions.

References

- Max H. Bazerman, Jared R Curhan, Don A Moore, and Kathleen L Valley. 2000. Negotiation. *Annual review of psychology*, 51(1):279–314.
- Kris Cao, Angeliki Lazaridou, Marc Lanctot, Joel Z. Leibo, Karl Tuyls, and Stephen Clark. 2018. Emergent communication through negotiation. In *Proc. ICLR*.
- Elizabeth Clark, Anne Spencer Ross, Chenhao Tan, Yangfeng Ji, and Noah A. Smith. 2018. Creative writing with a machine in the loop: Case studies on slogans and stories. In *23rd International Conference on Intelligent User Interfaces*, pages 329–340. ACM.
- Heriberto Cuayáhuatl, Simon Keizer, and Oliver Lemon. 2015. Strategic dialogue management via deep reinforcement learning. In *NIPS Workshop on Deep Reinforcement Learning*.

- Cristian Danescu-Niculescu-Mizil, Moritz Sudhof, Dan Jurafsky, Jure Leskovec, and Christopher Potts. 2013. A computational approach to politeness with application to social factors. In *Proc. ACL*.
- Michael Denkowski and Alon Lavie. 2014. Meteor universal: Language specific translation evaluation for any target language. In *Proc. WMT*, pages 376–380.
- Valerian J. Derlaga and John H. Berg. 1987. *Self-disclosure: Theory, research and therapy*. Springer2.
- Roger Fisher and William Ury. 1981. *Getting to Yes: Negotiating Agreement Without Giving In*. Boston, MA: Houghton Mifflin Company.
- He He, Anusha Balakrishnan, Mihail Eric, and Percy Liang. 2017. Learning symmetric collaborative dialogue agents with dynamic knowledge graph embeddings. In *Association for Computational Linguistics (ACL)*.
- He He, Derek Chen, Anusha Balakrishnan, and Percy Liang. 2018. Decoupling strategy and generation in negotiation dialogues. In *EMNLP*.
- Joan B. Hooper. 1975. *On assertive predicates*. In J. Kimball, editor, *Syntax and Semantics*, volume 4, pages 91–124. Academic Press, New York.
- Ken Hyland. 2005. Metadiscourse: Exploring interaction in writing. *Continuum, London and New York*.
- Adam N. Joinson. 2001. Self-disclosure in computer-mediated communication: The role of self-awareness and visual anonymity. *European Journal of Social Psychology*, 31(2), 177–192.
- De Dreu C. K., Weingart Laurie R., and Kwon Seung-woo. 2000. Influence of social motives on integrative negotiation: a meta-analytic review and test of two theories. *Journal of personality and social psychology*, 78(5):889.
- Simon Keizer, Markus Guhe, Heriberto Cuayáhuitl, Ioannis Efstathiou, Klaus-Peter Engelbrecht, Mihai Dobre, Alex Lascarides, and Oliver Lemon. 2017. Evaluating persuasion strategies and deep reinforcement learning methods for negotiation dialogue agents. In *Proc. EACL*.
- David A. Lax and James K. Sebenius. 2006. *3-D Negotiation: Powerful tools to change the game in your most important deals*. Harvard Business Press.
- Mike Lewis, Denis Yarats, Yann Dauphin, Devi Parikh, and Dhruv Batra. 2017. Deal or no deal? end-to-end learning of negotiation dialogues. In *Proc. EMNLP*, pages 2443–2453.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios Spithourakis, Jianfeng Gao, and Bill Dolan. 2016a. A persona-based neural conversation model. In *Proc Association for Computational Linguistics (ACL)*.
- Jiwei Li, Will Monroe, Alan Ritter, Michel Galley, Jianfeng Gao, and Dan Jurafsky. 2016b. Deep reinforcement learning for dialogue generation. In *Proc Conference on Empirical Methods in Natural Language Processing*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pp. 3111–3119.
- Alexandros Papangelis and Kallirroi Georgila. 2015. Reinforcement learning of multi-issue negotiation dialogue policies. In *SIGDIAL*.
- V. Petukhova, H. Bunt, and A Malchanau. 2017. Computing negotiation update semantics in multi-issue bargaining dialogues. In *Proc. SEMDIAL 2017 (SaarDial) Workshop on the Semantics and Pragmatics of Dialogue*, 87–97, DOI: 10.21437/SemDial.2017-10.
- Dean G. Pruitt. 1981. *Negotiation behavior*. New York: Academic Press.
- Chenhao Tan, Vlad Niculae, Cristian Danescu-Niculescu-Mizil, and Lillian Lee. 2016. Winning arguments: Interaction dynamics and persuasion strategies in good-faith on-line discussions. In *In Proceedings of WWW*.
- Yla R. Tausczik and James W. Pennebaker. 2010. The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of language and social psychology*, 29(1):24–54.
- Leigh L Thompson. 2001. *The mind and heart of the negotiator*, volume 3. Prentice Hall Upper Saddle River, NJ.
- Leigh L Thompson, Jiunwen Wang, and Brian C Gunia. 2010. Negotiation. *Annual review of psychology*, 61:491–515.
- David Traum, Stacy C. Marsella, Jonathan Gratch, Jina Lee, and Arno Hartholt. 2008. Multi-party, multi-issue, multi-strategy negotiation for multi-modal virtual agents. In *International Workshop on Intelligent Virtual Agents*, pages 117–130. Springer.
- Richard E. Walton and Robert B. McKersie. 1965. *A behavioral theory of labor negotiations: An analysis of a social interaction system*. New York: McGraw-Hill.
- Yi-Chia Wang, Moira Burke, and Robert Kraut. 2016. *Modeling Self-Disclosure in Social Networking Sites*. CSCW '16, SAN FRANCISCO, CA, USA.
- Amy Beth Warriner, Victor Kuperman, and Marc Brysbaert. 2013. Norms of valence, arousal, and dominance for 13,915 english lemmas. In *Behavior Research Methods, Volume 45, Issue 4*, pp 1191–1207.

Denis Yarats and Mike Lewis. 2018. Hierarchical text generation and planning for strategic dialogue. In *ICML*.

9 Appendix

Product Listing:

Listing Price: 14500

Buyer's Target Price: 8700

Title: "2006 Toyota 4Runner 4WD - Only 106k Miles - Clean Title"

Product Description:

"Selling my 2006 Toyota 4 Runner with only 106k original miles. The truck is in great condition with no mechanical flaws whatsoever and a clean accident history. Got new tires about 3,000 miles ago. Always has the oil changed on time (due in about 1k). Just got a thorough cleaning inside and a wash and wax outside (still wet in the photos). This truck has never been offroad, but the 4WD is working perfectly from the few times we've been up to Tahoe in it. However, it's a 10+ year old truck that's been driven, not babied and garaged all the time. It's got some scratches, paint is not perfect, but zero body damage."

No Coaching:
Seller: S Buyer: B

Static Coaching:
Seller: S Buyer: B

1. B: I just saw your ad for the 4Runner, can you send more picture of the scratches?
2. S: I don't have pictures of the scratches but I can assure you it's minor
3. B: I might be interested, but all I can offer is \$7500
4. S: That is very low. Can I agree to 11000?
5. B: That is too high for me, I mean it is 10 years old with over 100,000 miles. I can possibly come up to \$8,000
6. S: I can agree to 9,000 and make sure it's had a oil change and tire rotation before you pick it up.
7. B: If you're willing to have the oil changed and tires rotated, I will go as high as \$8700. And I really can't go any higher.
8. S: That's fine. I can agree to that.
9. B: Thanks, I'll be right over to pick it up.

10. S: i would sell it for 1400
11. B: you got a deal. 1400 it is
12. S: sorry meant 14000
13. B: Oh c'mon. Now you got my hopes up. lol
14. S: i can go low around 12000
15. B: I'm looking at the blue book right now and that still seems a bit high.
16. S: well the lowest i can go is 10000
17. B: You mention scratches. Lets be real. How bad are we talking?
18. S: its 10 yrs old it has some scratches but has zero body damage.
19. B: Ok. Without seeing any more photos, 9000 is the best I can do.
20. S: deal
21. B: deal. Thanks!

Dynamic Coaching:
Seller: S Buyer: B

22. S: Would you be interested in buying my 4Runner?
23. B: Yes, I am possibly. interested Given that is is over 10 years old, all I can offer now is 8000.
24. Coach: *Reject Buyer's Offer and Propose a New Price, Reason the Price with Content in the Product Description Try to Use the Word "Would" Like This: "Sorry, I would really need \$100 for it."*
...
S: I'm sorry, but I **would** really need \$12,000 for it. **It's in great condition!**
25. B: Well given the mileage on the car it is about to be due for some major service which can be expensive. How about 9000?
26. Coach: *Reject Buyer's Offer and Propose a New Price Try to Use the Word "Could" Like This: "I could come down to \$3." You Can Also Give the Buyer Something for Free (Gift Card) to Argue for Higher Price.*
...
S: I **could** come down to \$11,000 and **include the roof rack as a bonus.**
27. B: Add in a full tank of gas and a cold beer and I could do 10500.
28. Coach: *Reject Buyer's Offer and Propose a New Price Talk About Your Personal Experience With The Product Try to Use the Word "Would" Like This: "\$160 would be a good price."*
...
S: I **always took great care of the truck.** I think \$10,750 **would** be a reasonable price.
29. B: Deal.
30. S: Great

Table 9: Examples of collected negotiation dialogs.

Investigating Evaluation of Open-Domain Dialogue Systems With Human Generated Multiple References

Prakhar Gupta¹, Shikib Mehri¹, Tiancheng Zhao¹,
Amy Pavel², Maxine Eskenazi¹, and Jeffrey P. Bigham^{1,2}

¹Language Technologies Institute, Carnegie Mellon University,
²Human-Computer Interaction Institute, Carnegie Mellon University,
{prakharg, amehri, tianchez, apavel, max+, jbigam}@cs.cmu.edu

Abstract

The aim of this paper is to mitigate the shortcomings of automatic evaluation of open-domain dialog systems through multi-reference evaluation. Existing metrics have been shown to correlate poorly with human judgement, particularly in open-domain dialog. One alternative is to collect human annotations for evaluation, which can be expensive and time consuming. To demonstrate the effectiveness of multi-reference evaluation, we augment the test set of DailyDialog with multiple references. A series of experiments show that the use of multiple references results in improved correlation between several automatic metrics and human judgement for both the quality and the diversity of system output.

1 Introduction

Dialog agents trained end-to-end to hold open-domain conversations have recently progressed rapidly, generating substantial interest (Ghazvininejad et al., 2018; Serban et al., 2017, 2016a; Sordoni et al., 2015; Vinyals and Le, 2015). Development of these systems is driven by available data and benchmarks based on only a single ground truth reference response for a given context. However, such single-reference evaluation does not account for all the plausible responses for any given conversational context (Table 1). This is known as the *one-to-many* response problem (Zhao et al., 2017a). Computing word-overlap metrics against a single-reference response may penalize perfectly valid responses (Deriu et al., 2019) (e.g., “Was anything stolen?”, “Is anyone hurt”) that deviate from the particular target response (“When was the break-in?”). Unlike human evaluation, automatic evaluation with a single-reference may also disproportionately benefit models that produce generic responses with more probable words (e.g., “I don’t know”)

Dialog Context:

Person A: 911 emergency. What is the problem?

Person B: I would like to report a break-in.

single-reference Response:

When was this break-in?

Other Valid Responses:

Was anything stolen?

Is anyone hurt or injured?

Is the perpetrator still inside the house?

I will send someone right away.

Table 1: Example of a dialog context where appropriate responses do not share words and meaning with a single-reference response.

which is known as the dull-response problem (Li et al., 2016c). As a result, single-reference evaluations correlate weakly with human judgments of quality (Liu et al., 2016).

To address these problems, this paper proposes to carry out automatic evaluation using multiple reference responses instead of a single-reference. Multiple reference evaluation is attractive for several reasons. First, the additional information in the multiple reference response can be used to provide more robust quality evaluation under the one-to-many condition. Second, we can use the multiple references to better measure the diversity of the model, which is a widely studied topic in open-domain response generation (Kulikov et al., 2018; Li et al., 2016a; Zhang et al., 2018; Li et al., 2016b; Zhao et al., 2017a; Gao et al., 2019).

Prior explorations in this area either rely on synthetically created or small scale reference sets (Galley et al., 2015; Qin and Specia, 2015), or perform experiments only on a small set of metrics focused on only response quality (Sugiyama et al., 2019). Our investigations for using multiple references for automatic evaluation covers the

following aspects - 1) We propose methodology for evaluating both the quality and the diversity of generated responses using multiple references. 2) The proposed evaluation framework is metric-agnostic and the experiments cover a large spectrum of existing metrics, and 3) We augmented the existing test set of DailyDialog dataset (Li et al., 2017) with multiple references and perform human judgment correlation studies with human-generated references. Our extensive experimental results show that using multiple test references leads to significantly better correlation of automated metrics with human judgment in terms of both response quality and diversity. This suggests that the use of multiple references serves to make automatic metrics more reliable mechanisms for evaluating open-domain dialog systems. Moreover, follow up studies are conducted to better understand the nature of the multi-reference evaluation, such as the number of reference responses needed to achieve high correlation.

The contributions of this paper are:

1. We show that multi-reference evaluation achieves better correlation with human judgments both in quality and in diversity.
2. We analyze the effect of varying the number of reference responses on the correlation with human quality judgements.
3. We construct and release an open-domain multi-reference test dataset¹.

2 Related work

The need for reliable and consistent automatic evaluation methodologies has led to increasing interest in dialog system evaluation in recent years. In domains such as machine translation and captioning, n-gram overlap metrics such as BLEU (Papineni et al., 2002), ROUGE (Lin, 2004) and METEOR (Lavie and Agarwal, 2007) correlate well with human judgement. Several embedding-based metrics have been proposed as well, including Greedy Matching (Rus and Lintean, 2012) and Vector Extrema (Forgues et al., 2014). These automatic metrics, however, do not generalize well to open-domain dialog due to the wide spectrum of correct responses, commonly known as the one-to-many problem (Zhao et al., 2017b). Recent work has proposed several trainable evaluation metrics to address this issue. RUBER (Tao et al., 2018) evaluates generated re-

sponses based on their similarity with the reference responses and their relatedness to the dialog contexts. Lowe et al. (2017) trained a hierarchical neural network model called ADEM to predict the appropriateness score of responses. However, ADEM requires human quality annotation for training, which is costly. Sai et al. (2019) recently showed that trainable metrics are prone to gamification through adversarial attacks. While past work has focused on inventing new metrics, this paper instead aims to demonstrate that the correlation of existing metrics can be improved through the use of multiple references for evaluation in open-domain settings.

Prior attempts leveraged multiple references to improve evaluation in the context of text generation. Qin and Specia (2015) proposed variants of BLEU for machine translation based on n-gram weighting. In the dialog domain, Galley et al. (2015) proposed Discriminative BLEU, which leverages several synthetically created references obtained with a retrieval model from Twitter corpus. Sordoni et al. (2015) also followed a similar retrieval procedure for multiple-reference evaluation. Since both of them created their reference sets through retrieval followed by a rating step, their multi-reference sets do not reflect the natural variability in responses possible for a context. Sugiyama et al. (2019) proposed a regression-based evaluation metric based on multiple references. The small set of metrics and few test sentences shows promise, but also the need for further exploration. We go further with a comparison of single and multiple references for response quality evaluation and an examination of multiple references for diversity evaluation. This paper is the first, to our knowledge, to create a large test set of several human-generated references for each context. We believe that it is also the first to perform human correlation studies on a variety of automatic metrics for both quality and diversity.

Evaluating diversity in dialog model responses has been studied recently. The most commonly used metric is Distinct (Li et al., 2016a), which calculates the ratios of unique n-grams in generated responses. Distinct is, however, computed across contexts and does not measure if a model can generate multiple valid responses for a context. Xu et al. (2018) proposed Mean Diversity Score (MDS) and Probabilistic Diversity Score (PDS) metrics for diversity evaluation over groups

¹<https://github.com/prakharguptaz/multirefeval>

of multiple references over a set of retrieved references. Hashimoto et al. (2019) proposed a metric for a unified evaluation of quality and diversity of outputs, which however depends on human judgements. Zhao et al. (2017a) proposed precision/recall metrics calculated using multiple hypotheses and references as an indicator of appropriateness and coverage. In this paper we leverage their recall-based metrics in our multi-reference based evaluation of diversity.

3 Methodology

We evaluated the performance of dialog response generation models from two aspects: **quality** and **diversity**. Quality tests the appropriateness of the generated response with respect to the context, and diversity tests the semantic diversity of the appropriate responses generated by the model.

We first describe the evaluation procedures used for the conventional single-reference setting. Then we present the proposed multi-reference evaluation. We define a generalized metric to be $d(y, r)$ which takes a produced output y and a reference output r , and produces a matching score that measure the level of similarity between y and r . We discuss options for d in Table 2.

3.1 Baseline: Single-reference Evaluation

3.1.1 Quality

During single-reference evaluation, there is only one reference response r . As such, for a given metric d , the single-reference score will be $d(y, r)$.

3.1.2 Unreferenced Diversity

Most prior work concentrates on unreferenced diversity evaluation since referenced diversity evaluation requires a multi-reference dataset. Unreferenced evaluation refers to diversity evaluation methods which ignore the reference responses, and instead compute diversity as a function only of the generated responses. The Distinct (Li et al., 2016a) metric calculates diversity by calculating the number of distinct n-grams in generated responses as a fraction of the total generated tokens. This score is calculated at the system level - over the set of responses generated for all the contexts in test set. Given a set of system responses for the same context, Self-BLEU (Zhu et al., 2018) sequentially treats each one of the generated responses as the hypothesis and the others as references. This score is computed for every context

and then averaged over all contexts. A lower Self-BLEU implies greater diversity since system outputs are not similar to one another.

3.2 Proposed: Multi-Reference Evaluation

3.2.1 Quality

In multi-reference evaluation, a given context has multiple valid responses $R = \{r_1, r_2, \dots, r_n\}$. As such, for a given metric d , the multi-reference score can be computed as:

$$\text{score}(y, R) = \max_{r \in R} d(y, r) \quad (1)$$

We score the system output against only the closest reference response because there are multiple diverse and valid responses for a given context.

3.2.2 Referenced Diversity

A multi-reference test set also allows referenced diversity evaluation. For a given context c , we are given multiple reference responses $R = \{r_1, r_2, \dots, r_n\}$ and multiple system outputs $Y = \{y_1, y_2, \dots, y_m\}$. For a given metric, d , we compute recall (Zhao et al., 2017a), or *coverage*, as follows:

$$\text{recall}(c) = \frac{\sum_{j=1}^M \max_{i \in [1, N]} d(y_i, r_j)}{M} \quad (2)$$

For each of the multiple reference responses, we consider the highest-scoring system output, then average these scores across the reference responses. A system that generates outputs covering a large portion of the reference responses thus receives a higher recall score.

3.3 Metrics

We consider several metrics for quality and diversity evaluation including (1) word-overlap metrics, and (2) embedding-based metrics. We describe the metrics in Table 2. Each metric represents an instantiation of the generalized scoring function d .

3.4 Compared Models

Our experiments are conducted using four models: a retrieval model and three different generative models. We treat human generated responses as an additional model.

Human: To represent ideal model performance for a particular context, we use a human-generated response for that context.

Dual Encoder: A strong baseline for dialog retrieval is the Dual Encoder (DE) architecture

Metric	Reference	Description
Word-overlap based metrics		
BLEU	Papineni et al. (2002)	BLEU is based on n-gram overlap between the candidate and reference sentences. It includes a brevity penalty to penalize short candidates.
METEOR	Lavie and Agarwal (2007)	The harmonic mean of precision and recall between the candidate and reference based on a set of alignments between the two.
ROUGE-L	Lin (2004)	An F-measure based on the Longest Common Subsequence (LCS) between the candidate and reference utterances.
Embedding based metrics		
Embedding Average	Wieting et al. (2015), others	Computes a sentence-level embedding of r and c by averaging the embeddings of the tokens composing the sentences.
Vector Extrema	Forgues et al. (2014)	Computes a sentence-level embedding by taking the most extreme value of the embeddings of tokens of the sentence for each dimension of the embedding.
Greedy Matching	Rus and Lintean (2012)	Each word in the candidate sentence is greedily matched to a word in the reference sentence based on the cosine similarity of their embeddings. The score is then averaged for each word in the candidate sentence.
Skip-Thought	Kiros et al. (2015)	Uses a recurrent network to encode a given sentence into a sentence level embedding. We use the pre-trained vectors and implementation provided by (Sharma et al., 2017).
GenSen	Subramanian et al. (2018)	Generates a sentence level embedding through a sequence-to-sequence model trained on a variety of supervised and unsupervised objectives in a multi-task framework.

Table 2: Metrics used for both quality and diversity evaluation.

(Lowe et al., 2015a). The model first encodes a given dialog context and response using an LSTM encoder. It then takes the dot-product of the two latent representations to output the likelihood of the response. The Dual Encoder is trained to differentiate between correct responses, and uniformly sampled negative responses. During inference, however, it chooses a correct response for a given context out of all the responses that occur in the training set.

Seq2Seq: Sequence-to-sequence (Seq2Seq) networks (Sutskever et al., 2014) are a typical baseline for dialog systems (Vinyals and Le, 2015). Our model consists of an LSTM encoder, an LSTM decoder and an attention mechanism (Bahdanau et al., 2014).

HRED: Hierarchical Recurrent Encoder Decoder networks (HRED) (Serban et al., 2016b) are a modification of Seq2Seq networks. Rather than encoding the context as a sequence of words, the encoding of the context is done in a two-step process. First, all the utterances of a context are independently encoded by an LSTM utterance encoder. Second, given the latent representations of each utterance, a context encoder encodes the dialog context. The attention mechanism of the decoder attends over the timesteps of context encoder.

CVAE: The Conditional Variational Autoencoder (CVAE) model (Zhao et al., 2017a). CVAE mod-

els incorporate discourse-level latent variables in HRED, in which the latent variables represent the discourse-level intentions of the system. Specifically, we reproduce the CVAE network from (Zhao et al., 2017a), where the latent variables follow a multivariate Gaussian distribution with a diagonal covariance matrix. The dimension of the latent variable is 256. To have a fair comparison, the rest of the structure is the same as the HRED with bidirectional LSTM utterance encoders and LSTM context encoder and response decoder. To alleviate the posterior collapse issue for training text CVAEs (Bowman et al., 2016), we use bag-of-words auxiliary loss (Zhao et al., 2017a) and KL-annealing (Bowman et al., 2016).

4 Multi-Reference Data Collection

We used the following procedure to prepare the DailyDialog test set for the multi-reference test set collection. A dialog D in the test set consists of utterances $\{u_1, u_1, \dots, u_n\}$. Here, u_i denotes the utterance at the i th turn. For generating dialog contexts, we truncate the dialog at each possible utterance, except the last one. The response following each context is treated as the reference response. As an illustration, for the Dialog shown in Table 1, we would generate the following context-reference pairs: *Context 1*: “911 emergency. What is the problem?”, *Reference 1*: “I would like to report a break-in.”. *Context 2*: “911 emergency

Reference	Very Appropriate	Appropriate	Neutral	Not Appropriate	Not Appropriate at all
From original dataset	41%	54%	2%	3%	0%
Sampled from multi-reference collected	40%	52%	3%	5%	0%

Table 3: Results from dataset quality experiment

... report a break-in.”, *Reference 2*: “When was this break-in?”. In our multi-reference dataset, we expand each single-reference to a set of multiple references.

4.1 Data collection Procedure

We designed an interface for multi-reference data collection using Amazon Mechanical Turk (AMT). For every HIT, we asked an AMT worker to generate 4 diverse follow-up responses for a conversation. A snapshot of the data collection interface is shown in Figure 3 (Appendix). We provided instructions and examples to further clarify the task. To maintain quality post data collection, we filter out responses collected from workers who either generated very short responses or entered the responses in very short amount of time consistently.

4.2 Data Quality

Using the method described above, we collected 4 diverse responses for the 1000 dialogs in the test set, which consists of 6740 contexts. To validate the quality of the collected dataset, an experiment on AMT is carried out for 100 contexts sampled randomly from the dataset. Workers are shown a dialog context followed by 3 responses shuffled in a random order - 1) the original response from the dataset 2) a random response from the collected multi-references, and 3) a distractor response, irrelevant to the dialog context. We use distractor responses to filter out poor annotations where the annotator gave high ratings to the distractor response. We ask the workers to rate each of the 3 responses for a dialog context on a scale of 1-5 for appropriateness, where 1 indicates *Not Appropriate at all* and 5 indicates *Very Appropriate*. We present the ratings from the experiment in Table 3 for the original responses from the dataset, and the responses from the multi-reference set. We observe that 92% sampled responses from the multi-reference set are marked Appropriate or Very Appropriate. Moreover, only 8% of the responses are marked Not Appropriate or lower, compared to 5% for the original reference set. This indi-

cates that the collected reference set is close to the original reference set in quality. Furthermore, the responses are generated specifically for each context, they are coherent with the context.

5 Experiments

This section describes the experiments we conducted to explore the effectiveness of multi-reference evaluation.

5.1 Correlation Analysis for Quality

This analysis aims to compute the correlation between human quality judgments and two forms of automatic evaluation, both single-reference and multi-reference.

5.1.1 Human Annotations

A collection of 100 dialog contexts are randomly selected from the dataset. For a particular dialog context, each of the four models produces a response. In addition, we collect a human response using Amazon Mechanical Turk (AMT), making it total of five responses for each dialog context. Given these context-response pairs, each response is rated in terms of appropriateness (from 1-5) by 5 different AMT workers. The ratings are removed for workers with a Cohen’s Kappa κ (Cohen, 1968) inter-annotator agreement score of less than 0.2. The remaining workers had a mean κ score of 0.43, indicating moderate agreement.

5.1.2 Results

Utterance level correlation: The results of the correlation study conducted for 5 model responses for 100 contexts are shown in Table 4. Pearson correlation is computed to estimate linear correlation, and Spearman correlation to estimate monotonic correlation. The correlations with human quality judgments are computed for both single-reference and multi-reference evaluation. The multi-reference test set consists of both the original reference and the four new collected reference responses. For single-reference evaluation, except for METEOR and Vector Extrema metrics, the correlation is either small or statistically

Metrics	Single Reference				Multiple Reference			
	Spearman	p-value	Pearson	p-value	Spearman	p-value	Pearson	p-value
BLEU-1	0.0241	0.591	0.1183	0.008	0.1572	0.000	0.2190	0.000
BLEU-2	0.0250	0.577	0.1803	0.000	0.2077	0.000	0.2910	0.000
BLEU-3	0.0608	0.175	0.1269	0.005	0.2520	0.000	0.2086	0.000
BLEU-4	0.0345	0.441	0.1380	0.002	0.2202	0.000	0.2333	0.000
METEOR	0.1064	0.017	0.1871	0.000	0.2247	0.000	0.2855	0.000
ROUGE-L	0.0715	0.110	0.1408	0.002	0.2203	0.000	0.2798	0.000
Embedding Average	0.0301	0.502	-0.0067	0.880	0.1248	0.005	0.0636	0.156
Vector Extrema	0.1919	0.000	0.2114	0.000	0.2785	0.000	0.2946	0.000
Greedy Matching	0.1306	0.003	0.1150	0.010	0.2367	0.000	0.2352	0.000
Skip-Thought	-0.0029	0.949	-0.1463	0.001	0.1049	0.019	-0.0716	0.109
GenSen	0.0731	0.103	0.1110	0.013	0.1832	0.000	0.2389	0.000

Table 4: Correlation of various metrics when evaluated using single-reference and multi-reference test sets. Evaluation using Multiple References leads to better correlation across all metrics.

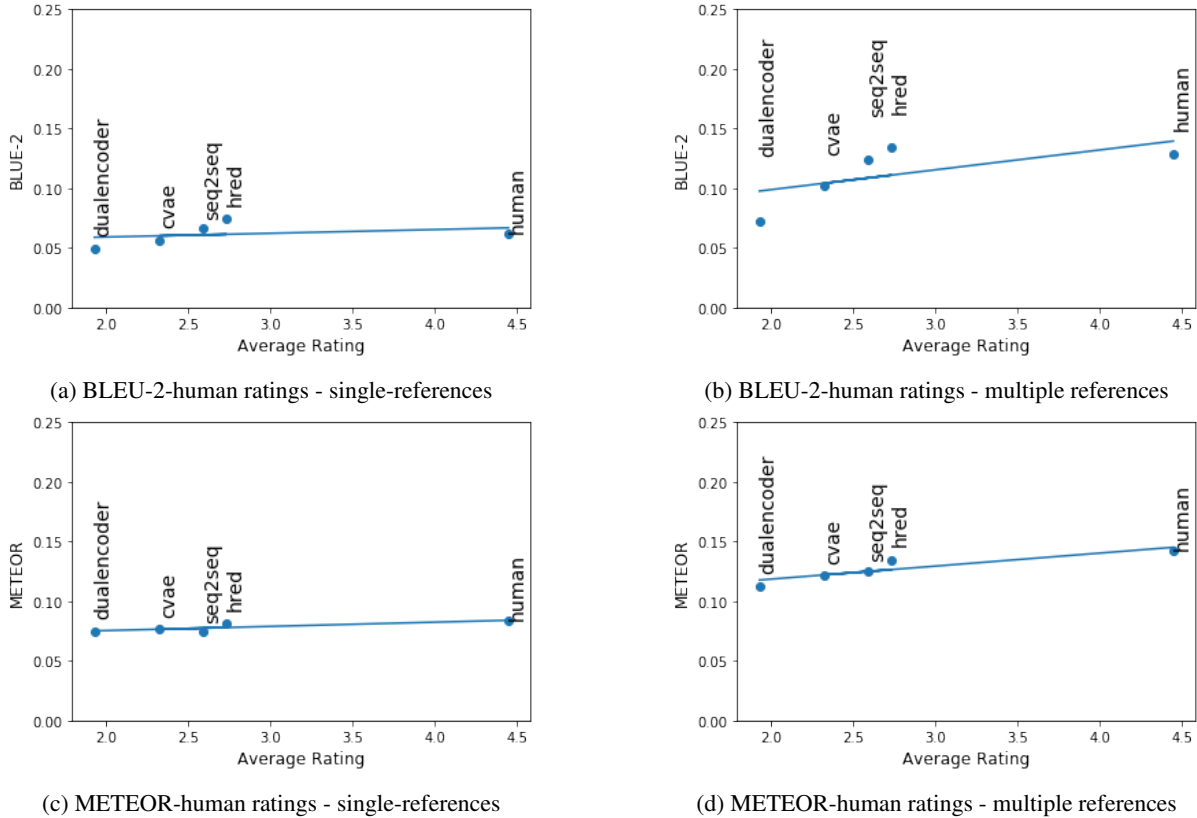


Figure 1: System level correlations for BLEU-2 and METEOR metrics. Multi-reference evaluation shows higher correlation with more clear differentiation in model performance.

less significant. On the other hand, every metric shows higher and significant correlation for multi-reference evaluation, with METEOR, ROUGE-L and Vector Extrema achieving the highest correlation values. These results indicate that multi-reference evaluation correlates significantly better with human judgment than single-reference, across all the metrics. This reaffirms the hypothesis that multi-reference evaluation better captures the one-to-many nature of open-domain dialog.

System level correlation: For each model used

in the correlation study, the average human rating and average metric scores for 100 contexts are used to calculate system-level correlations. We show system-level correlations for metrics BLEU-2 and METEOR metrics in Figure 1. Each point in the scatter plots represents the average scores for a dialog model. Average human scores are shown on the horizontal axis, with average metric scores on the vertical axis. Humans ratings are low for responses from the retrieval model, and higher for human responses and responses from HRED

model. It is clear that the difference in scores for models when evaluated using single-references is not significant enough to compare the models, as the average metric scores have near zero or very weak correlation with average human ratings. This renders them insufficient for dialog evaluation. However, with multi-reference evaluation, the correlation is higher and significant, which differentiates the models clearly. Thus, multi-reference based evaluation correlates well with humans both at utterance level and at the system level.

5.2 Correlation Analysis for Diversity

This section aims to demonstrate that referenced diversity evaluation methods better correlate with *human judgements of diversity*, than previously used unreferenced diversity metrics. While unreferenced metrics simply reward lexical differences amongst generated outputs, referenced methods (e.g., the recall metric) aims to calculate the *coverage* of the responses. The correlation of human diversity scores is calculated with both unreferenced and referenced measures of diversity.

5.2.1 Human Annotations

Multiple hypotheses were generated from all the models. For CVAE, multiple responses are sampled from the latent space with greedy word-level decoding. For rest of the generation models, five responses were obtained using sampled decoding. For retrieval models, the top five retrieved responses were used. Human annotations of these multiple hypotheses were collected as follows: (1) Workers mark the responses which they find to be appropriate for the conversational context, (2) They then provide a score for the diversity of the responses based on how different they are in *meaning*. This two-stage annotation process captures a desired form of system diversity: generated outputs should be varied, but also appropriate. The scores are averaged across the three workers’ annotations. We filtered out ratings from workers with low inter-annotator agreement as described in section 5.1.1. The final mean κ score of 0.41, which indicates moderate agreement.

5.2.2 Results

The results for the diversity correlation analysis are shown in Table 5 for a selected set of metrics². The unreferenced metrics, Distinct and Self-

²For Self-BLEU we calculate correlation with values subtracted from 1 as Self-BLEU is inversely related to diversity

Metric	Spearman	p-value	Pearson	p-value
Distinct-1	0.0204	0.647	0.0465	0.299
Distinct-2	-0.1282	0.004	-0.0568	0.205
Distinct-3	-0.1316	0.003	-0.0184	0.681
Self BLEU-2	-0.1534	0.001	-0.1251	0.005
Self BLEU-4	-0.0836	0.061	-0.0304	0.497
Recall BLEU-2	0.2052	0.000	0.2469	0.000
Recall BLEU-4	0.1713	0.000	0.1231	0.005
Recall METEOR	0.1993	0.000	0.2165	0.000
Recall ROUGE-L	0.1862	0.000	0.2234	0.000
Recall Vector Extrema	0.2063	0.000	0.2314	0.000
Recall Greedy Matching	0.0797	0.075	0.1204	0.007

Table 5: Correlation scores for diversity metrics

BLEU, correlate poorly with human judgment. This is probably because these metrics evaluate lexical diversity, while humans evaluate diversity of meaning. Furthermore, unreferenced metrics do not consider the reference response and reward diverse outputs without considering appropriateness. With referenced diversity evaluation, using the recall method, BLEU-2 and Vector Extrema show the highest correlation. While metrics like Self-BLEU and Distinct can be “gamed” by producing meaningless albeit very diverse responses, the referenced recall metrics require both appropriate and diverse outputs. As such, referenced evaluation correlates significantly better with human notions of diversity. Thus, the construction of a multi-reference dataset allows for improved diversity metrics.

5.3 Automatic Evaluation of Models

We use our multi-reference evaluation methodology to compare the models and the human generated responses on the whole test dataset. For the human model, we use one reference from the multi-reference set as the hypothesis. Human responses are generally more interesting and diverse than model responses, which are known to suffer from the dull response problem (Li et al., 2016c). Because of this reason, we would expect the human generated responses to get higher scores than the dialog models. However, the results presented in Table 6 show that single-reference automatic evaluation ranks few models higher than the hu-

Metric	Single Reference					Multiple reference				
	Dual Encoder	Seq2Seq	HRED	CVAE	Human	Dual Encoder	Seq2Seq	HRED	CVAE	Human
BLEU-2	0.0399	0.0521	0.0604	0.0656	0.0513	0.0625	0.0981	0.1061	0.1033	0.1637
BLEU-4	0.0168	0.0252	0.0301	0.0291	0.0245	0.0241	0.0445	0.0497	0.0429	0.0791
METEOR	0.0653	0.0544	0.0607	0.0724	0.0592	0.1000	0.0970	0.1036	0.1120	0.1456
ROUGE-L	0.1522	0.1847	0.1998	0.2088	0.1682	0.2216	0.2927	0.3044	0.2997	0.3502
Vector Extrema	0.4005	0.5124	0.5002	0.4893	0.4823	0.4713	0.6191	0.5975	0.5722	0.6134
Greedy Matching	0.6257	0.7167	0.7104	0.7078	0.6799	0.6991	0.7649	0.7551	0.7457	0.7562
Recall BLEU-2	0.0662	0.0544	0.0766	0.1077	0.0898	0.0436	0.0377	0.0556	0.0679	0.0984
Recall Vector Extrema	0.4945	0.5127	0.5397	0.5586	0.5651	0.4934	0.5334	0.5476	0.5653	0.5881

Table 6: Model evaluation with automatic metrics on Single and Multiple references. Multiple reference evaluation is able to correctly rank human responses higher than model responses.




Figure 2: Change in correlation with varying number of references. Trend stabilizes after 4-5 references

mans model. With multi-reference evaluation, human performance is significantly higher than model performance. We further present scores for diversity metrics on multiple hypothesis generated for 100 contexts in the last two rows of the table. The use of multi-reference evaluation covers a wider array of valid responses, which strongly rewards the diverse human responses compared to single-reference evaluation.

5.4 Effect of number of references

The correlation of automated evaluation with human judgment is calculated at various numbers of reference responses. The results shown in Figure 2 demonstrate that the Pearson correlation with human judgment generally increases sharply up to 3-5 references. It further increases slowly up to about 7 references and then seems to plateau at around eight references. This suggests that four to eight references give sufficient coverage of the re-

Dialog Context:		
<i>Person A:</i> excuse me . check please .		
Generated Response		
sure , i 'll grab it and be right with you .		
Single-reference Response:		
ok , how was everything ?		
Multi-reference Responses:		
i 'll get it right away .		
here is the check .		
no problem , let me get your server .		
i 'll be right back with it .		
Average Human Rating: 5		
Metric	Single reference	Multiple reference
BLEU-2	0.0275	0.3257
METEOR	0.0539	0.3425
Vector Extrema	0.5523	0.8680

Table 7: Example of difference in metric scoring for single versus multiple reference evaluation.

sponse space, and collecting additional references does not provide much value in terms of mitigating the issues of the one-to-many problem.

6 Discussion and Conclusion

This work proposes a more reliable methodology for automatic evaluation of open-domain dialogues with the use of multiple references. We augment the test set of DailyDialog dataset with multiple references and show that multiple references lead to better correlation with human judgments of quality and diversity of responses. Single-reference based evaluation can unfairly penalize diverse and interesting responses which are appropriate, but do not match a particular reference in the dataset. However, multiple references can cover the possible semantic space of replies for a context better than a single reference. Thus using multi-reference test sets can improve

the way open-ended dialogue systems are currently evaluated. Our experiments also show that human-generated responses perform worse than models across most metrics when using single-reference evaluation, but multiple reference evaluation consistently ranks human responses higher than model-generated responses. Furthermore, we show how varying the number of references affects human judgement correlation. This methodology could easily be extended to other open domain datasets if the community can make similar multi-reference test sets publicly available.

We illustrate the strength of multi-reference evaluation through scores calculated for some metrics using both single and multiple references for an example context in Table 7. Multiple reference-based evaluation is often good at assigning higher scores when there is more scope for diversity in the responses as illustrated by the example. It should be noted that multiple reference evaluation generally increases the scale of metrics for all responses, and this includes dull responses.

The multi-reference data collection procedure in this paper collects the same number of responses for all contexts. However, different dialogue contexts might possess different levels of “open-endedness”. For e.g., a context like “Would you like to dance?” would generally have fewer possible variations in responses than a more open-ended context like “What did you do yesterday?”. Therefore, the number of references to collect for a context could be based on the expected variability in responses for the context. Such a procedure would capture more variability over the dataset for a fixed budget.

An important direction in dialog system research is to build models that have more engaging and meaningful conversations with a human. With the recent push towards models which can generate more diverse and interesting responses, appropriate evaluation methodologies are an important and urgent need for the community. Human level evaluation of generation and diversity is challenging to do in a completely automatic way, however, compared to evaluating with a single response, we show that the proposed evaluation methodology is more reliable and will facilitate progress in this direction. In this work we have chose one dataset for extensive experimentation, but in the future studies, it will be worth collecting more datasets and repeating the correlation experiments.

7 Acknowledgements

This work was funded by the Defense Advanced Research Planning Agency (DARPA) under DARPA Grant N6600198-18908, and the National Science Foundation under Award #IIS-1816012. We thank the workers on Amazon Mechanical Turk for making our research possible.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. 2016. Generating sentences from a continuous space. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 10–21.
- Jacob Cohen. 1968. Weighted kappa: Nominal scale agreement provision for scaled disagreement or partial credit. *Psychological bulletin*, 70(4):213.
- Jan Deriu, Alvaro Rodrigo, Arantxa Otegi, Guillermo Echevoyen, Sophie Rosset, Eneko Agirre, and Mark Cieliebak. 2019. Survey on evaluation methods for dialogue systems. *arXiv preprint arXiv:1905.04071*.
- Gabriel Fergues, Joelle Pineau, Jean-Marie Larchevêque, and Réal Tremblay. 2014. Bootstrapping dialog systems with word embeddings. In *Nips, modern machine learning and natural language processing workshop*, volume 2.
- Michel Galley, Chris Brockett, Alessandro Sordani, Yangfeng Ji, Michael Auli, Chris Quirk, Margaret Mitchell, Jianfeng Gao, and Bill Dolan. 2015. [deltaBLEU: A discriminative metric for generation tasks with intrinsically diverse targets](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 445–450, Beijing, China. Association for Computational Linguistics.
- Xiang Gao, Sungjin Lee, Yizhe Zhang, Chris Brockett, Michel Galley, Jianfeng Gao, and Bill Dolan. 2019. Jointly optimizing diversity and relevance in neural response generation. *arXiv preprint arXiv:1902.11205*.
- Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih, and Michel Galley. 2018. A knowledge-grounded neural conversation model. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

- Tatsunori B Hashimoto, Hugh Zhang, and Percy Liang. 2019. Unifying human and statistical evaluation for natural language generation. *arXiv preprint arXiv:1904.02792*.
- Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. In *Advances in neural information processing systems*, pages 3294–3302.
- Ilya Kulikov, Alexander H Miller, Kyunghyun Cho, and Jason Weston. 2018. Importance of a search strategy in neural dialogue modelling. *arXiv preprint arXiv:1811.00907*.
- Alon Lavie and Abhaya Agarwal. 2007. **Meteor: An automatic metric for mt evaluation with high levels of correlation with human judgments**. In *Proceedings of the Second Workshop on Statistical Machine Translation*, StatMT '07, pages 228–231, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. **A diversity-promoting objective function for neural conversation models**. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios Spithourakis, Jianfeng Gao, and Bill Dolan. 2016b. **A persona-based neural conversation model**. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 994–1003, Berlin, Germany. Association for Computational Linguistics.
- Jiwei Li, Will Monroe, Alan Ritter, Michel Galley, Jianfeng Gao, and Dan Jurafsky. 2016c. Deep reinforcement learning for dialogue generation. *arXiv preprint arXiv:1606.01541*.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. **DailyDialog: A manually labelled multi-turn dialogue dataset**. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Chin-Yew Lin. 2004. **ROUGE: A package for automatic evaluation of summaries**. In *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. **How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation**. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132, Austin, Texas. Association for Computational Linguistics.
- Ryan Lowe, Michael Noseworthy, Iulian Vlad Serban, Nicolas Angelard-Gontier, Yoshua Bengio, and Joelle Pineau. 2017. **Towards an automatic Turing test: Learning to evaluate dialogue responses**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1116–1126, Vancouver, Canada. Association for Computational Linguistics.
- Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. 2015a. The ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems. In *SIGDIAL Conference*, pages 285–294. The Association for Computer Linguistics.
- Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. 2015b. **The Ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems**. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 285–294, Prague, Czech Republic. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **Bleu: A method for automatic evaluation of machine translation**. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, ACL '02, pages 311–318, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Ying Qin and Lucia Specia. 2015. **Truly exploring multiple references for machine translation evaluation**. In *Proceedings of the 18th Annual Conference of the European Association for Machine Translation*, pages 113–120, Antalya, Turkey.
- Alan Ritter, Colin Cherry, and William B Dolan. 2011. Data-driven response generation in social media. In *Proceedings of the conference on empirical methods in natural language processing*, pages 583–593. Association for Computational Linguistics.
- Vasile Rus and Mihai Lintean. 2012. A comparison of greedy and optimal assessment of natural language student input using word-to-word similarity metrics. In *Proceedings of the Seventh Workshop on Building Educational Applications Using NLP*, pages 157–162. Association for Computational Linguistics.
- Ananya B Sai, Mithun Das Gupta, Mitesh M Khapra, and Mukundhan Srinivasan. 2019. Re-evaluating adam: A deeper look at scoring dialogue responses. *arXiv preprint arXiv:1902.08832*.
- Iulian V. Serban, Alessandro Sordani, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016a. **Building end-to-end dialogue systems using generative hierarchical neural network models**. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI'16, pages 3776–3783. AAAI Press.

- Iulian V. Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016b. [Building end-to-end dialogue systems using generative hierarchical neural network models](#). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI'16, pages 3776–3783. AAAI Press.
- Iulian Vlad Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Shikhar Sharma, Layla El Asri, Hannes Schulz, and Jeremie Zumer. 2017. Relevance of unsupervised metrics in task-oriented dialogue for evaluating natural language generation. *arXiv preprint arXiv:1706.09799*.
- Alessandro Sordoni, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and William B. Dolan. 2015. A neural network approach to context-sensitive generation of conversational responses. In *HLT-NAACL*.
- Sandeep Subramanian, Adam Trischler, Yoshua Bengio, and Christopher J Pal. 2018. [Learning general purpose distributed sentence representations via large scale multi-task learning](#). In *International Conference on Learning Representations*.
- Hiroaki Sugiyama, Toyomi Meguro, and Ryuichiro Higashinaka. 2019. *Automatic Evaluation of Chat-Oriented Dialogue Systems Using Large-Scale Multi-references*, pages 15–25. Springer International Publishing, Cham.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Chongyang Tao, Lili Mou, Dongyan Zhao, and Rui Yan. 2018. Ruber: An unsupervised method for automatic evaluation of open-domain dialog systems. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2015. Towards universal paraphrastic sentence embeddings. *arXiv preprint arXiv:1511.08198*.
- Zhen Xu, Nan Jiang, Bingquan Liu, Wenge Rong, Bowen Wu, Baoxun Wang, Zhuoran Wang, and Xiaolong Wang. 2018. [LSDSCC: a large scale domain-specific conversational corpus for response generation with diversity oriented evaluation metrics](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2070–2080, New Orleans, Louisiana. Association for Computational Linguistics.
- Yizhe Zhang, Michel Galley, Jianfeng Gao, Zhe Gan, Xiujun Li, Chris Brockett, and Bill Dolan. 2018. [Generating informative and diverse conversational responses via adversarial information maximization](#). In *Proceedings of the 32Nd International Conference on Neural Information Processing Systems, NIPS'18*, pages 1815–1825, USA. Curran Associates Inc.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017a. [Learning discourse-level diversity for neural dialog models using conditional variational autoencoders](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 654–664, Vancouver, Canada. Association for Computational Linguistics.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017b. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. *arXiv preprint arXiv:1703.10960*.
- Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. 2018. [Taxygen: A benchmarking platform for text generation models](#). In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR '18*, pages 1097–1100, New York, NY, USA. ACM.

A Further Notes on Data Collection Experiments

The interface designed for multi-reference data collection is shown in Figure 3. The final design of the interface incorporates improvements based on multiple rounds of experiments and interviews on a small set of users. The workers were shown a modal box with instructions and several good and bad examples before they start the task. Then they are shown 5 contexts for a HIT, one by one. For each context, they are asked to write 4 diverse responses in the Textbox provided. Workers can enter multi-line responses and submit a response by pressing enter or clicking on a button. They are shown the number of remaining responses they need to enter for the conversation. We also record the timestamps for click and enter presses in the interface. We prevent workers from entering replies shorter than 2 characters, the exact same reply more than 1 time and show them a warning prompt if enter their response too quickly consistently.

Data Collection modes - For the collection of 4 responses per context, we have the following options - A) 4R1W- Collect 4 responses from a single worker B) 2R2W- Collect 2 responses each from 2 separate workers, and C) 1R4W - Collect 1 response each from 4 separate workers. In order to decide between these collection modes, we designed an experiment where, for 100 random contexts, we collected 4 responses using all three styles A), B) and C). In order to decide the best option, we measured lexical diversity across the 4 responses using self-BLEU (Zhu et al., 2018)

Metric	4R1W	2R2W	1R4W
SelfBLEU-1	0.3809	0.3662	0.4403
SelfBLEU-2	0.1778	0.1618	0.2657
SelfBLEU-3	0.0955	0.0851	0.2045
SelfBLEU-4	0.0548	0.0449	0.1748
Distinct-1	0.7266	0.7522	0.7082
Distinct-2	0.9240	0.9346	0.8782
Distinct-3	0.9621	0.9692	0.9092
Gt-BLEU-1	0.1213	0.1165	0.1296
Gt-BLEU-2	0.0258	0.0259	0.0352
Gt-BLEU-3	0.0091	0.0111	0.0136
Gt-BLEU-4	0.0033	0.0032	0.0033

Table 8: Diversity and relevance for different modes of data collection.

and Distinct (Li et al., 2016a) metrics, and the collected responses’ relevance through the average BLEU score of the multi-reference responses with the ground truth (Gt-BLEU) in the dataset. The results are reported in Table 8.

To calculate Self-BLEU, we calculate the BLEU score for every response by treating the response as a hypothesis and the others as the references, and we define the average BLEU scores calculated this way to be the Self-BLEU of the response set. A higher Self-BLEU score implies less diversity in the set. We observe that 4R1W and 2R2W achieve higher lexical diversity than 1R4W. This is because when a worker is asked to write multiple responses, they can make their responses more diverse conditioned on their previous responses. Relevance metrics Gt-BLEU-1,2,3,4 indicate that 1R4W achieve higher lexical similarity with the ground truth response in the dataset, followed by 4R1W. We chose the 4R1W mode, that is, a collection of 4 responses from 1 worker, to balance the diversity and relevance metrics.

Instructions for annotation collection for Diversity Study

We provided following instructions to the workers for collecting diversity ratings- “Please read the following conversation between two persons. Then read some possible follow-up responses for the conversation. You will be shown 5 sets of responses, with 5 responses in each set. For each response set, first select the responses you think are appropriate responses for the conversation. Then use the sliders to rate the diversity of the response set, that is, how many of the appropriate responses in the response set had different meanings or were different replies. Please provide the diversity score only for the appropriate responses you have marked. The diversity score should not be more than the number of appropriate responses in that set.” These instructions were followed by an example to make the task clear.

B Choice of dataset

There are only a few open-domain multi-reference datasets and they have been collected artificially either by retrieval (Xu et al., 2018; Galley et al., 2015) or are very small in scale (Sugiyama et al., 2019). Therefore we augmented the original test set of the DailyDialog dataset (Li et al., 2017), which has a sufficiently large test set. Conversa-

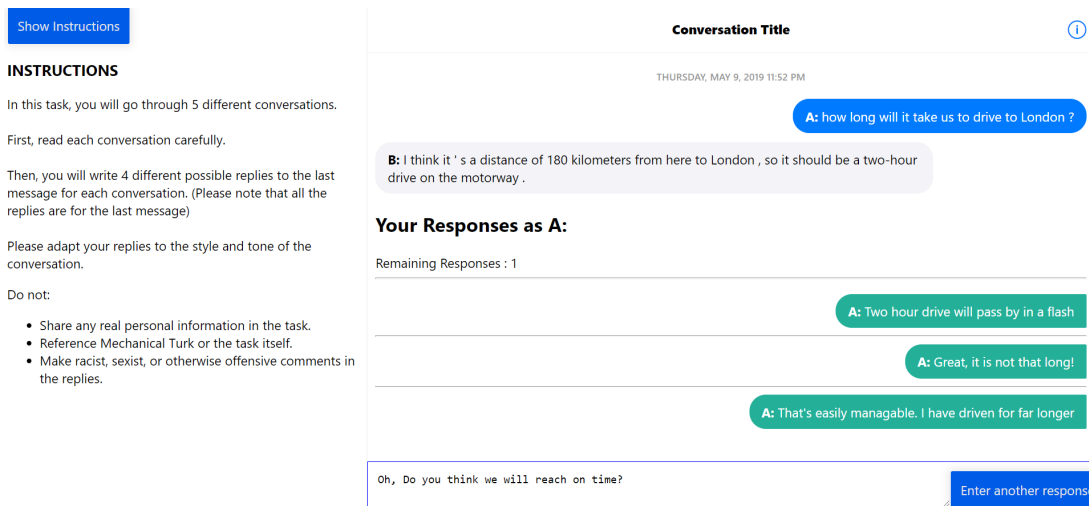


Figure 3: Interface used for multi-reference data collection.

Reference	Original	Multi-reference
Unique 1-gram	17.55	23.62
Unique 2-gram	27.88	58.69
Unique 3-gram	21.79	50.34

Table 9: Comparison of number of unique n-grams in original versus multiple references.

tions in DailyDialog cover 10 different topics on daily life. We chose to augment the DailyDialog dataset due to the following reasons- 1) The dialogs in this dataset are about daily conversation topics and thus it is easier to augment them using crowdsourcing. 2) The dialogs in this dataset are generally more formal than datasets such as the Twitter Dialog Corpus (Ritter et al., 2011) and Ubuntu Corpus (Lowe et al., 2015b) which contain noise such as typos and slangs. 3) The dialogs generally have a reasonable number of turns, which makes it easier for a person to understand the context and generate a reply. Therefore, given the size of the original DailyDialog test set and the above-mentioned properties of the dataset, we chose to augment the test set of DailyDialog.

Dataset quality continued

We present the average number of unique 1, 2 and 3 grams in the original ground truth and the set of collected multi-reference ground truth in Table 9. The higher number of unique ngrams in the multi-reference ground truth indicates that the new ground truth captures more variation in the set of possible responses.

User Evaluation of a Multi-dimensional Statistical Dialogue System

Simon Keizer,* Ondřej Dušek,*† Xingkun Liu* and Verena Rieser*

*Interaction Lab, Heriot-Watt University, Edinburgh, Scotland, UK

†Charles University, Faculty of Mathematics and Physics, Prague, Czechia

keizer.simon@gmail.com, odusek@ufal.mff.cuni.cz,

{x.liu,v.t.rieser}@hw.ac.uk

Abstract

We present the first complete spoken dialogue system driven by a *multi-dimensional* statistical dialogue manager. This framework has been shown to substantially reduce data needs by leveraging domain-independent dimensions, such as social obligations or feedback, which (as we show) can be transferred between domains. In this paper, we conduct a user study and show that the performance of a multi-dimensional system, which can be adapted from a source domain, is equivalent to that of a one-dimensional baseline, which can only be trained from scratch.

1 Introduction

Data-driven approaches to spoken dialogue systems (SDS) are limited by their reliance on substantial amounts of annotated data in the target domain. This can be addressed by considering transfer learning techniques, e.g. (Taylor and Stone, 2009), in which data from a source domain is leveraged to improve learning in a target domain. In particular, domain adaptation has been used in the context of dialogue systems (Gašić et al., 2017; Wang et al., 2015; Wen et al., 2016), focusing on identifying and exploiting similarities between domain ontologies in slot-filling tasks.

In contrast to this previous work, we take a *multi-dimensional* approach, which combines machine learning with linguistic theory. Following Bunt (2011), we exploit the linguistic phenomenon that utterances serve more than one function in a conversation, i.e. they have more than one *dimension* (see Section 2).¹ For example, the utterance “On what date would you like to fly to London?” both asks a task-oriented question, and provides feedback about understanding the requested destination. We take advantage of this phenomenon by training

separate, fully-statistical dialogue models for each dimension and generating system responses along multiple dimensions simultaneously. Such an SDS thus has the potential to adapt more efficiently to new domains by exploiting previously trained policies of the domain-independent dimensions, such as feedback and social conventions.

Previous implementations of multi-dimensional SDSs were mostly handcrafted (Akker et al., 2005; Petukhova et al., 2016). Keizer and Rieser (2017) were the first to present a statistical multi-dimensional dialogue manager (DM). Their results suggest an up to 80% reduction in data: a task success rate of over 90% can be achieved after only 2,000 dialogues when using pre-trained policies, whereas at least 10,000 dialogues are required without pre-training. In comparison, Gašić et al. (2017) achieve similar success rates for in-domain systems trained on 5,000 dialogues. However, Keizer and Rieser’s findings are only tested in simulation.

In this paper, we present the first complete statistical SDS with multi-dimensional DM, and the first crowdsourced human user evaluation of this type of system, comparing a one-dimensional baseline and three multi-dimensional variants, using a novel web-based setup. A novel aspect of our statistical analysis is testing for *equivalence*. The four system variants were designed in such a way that we would expect their performance levels to be indistinguishable when using fully trained policies. Should the data provide statistical evidence for this, the multi-dimensional variants can be preferred due to their inherent potential for domain transfer.

2 A Multi-dimensional Dialogue Manager

Our DM is a partially-observable Markov decision process (POMDP; Young et al., 2013) which takes as input an n-best list of dialogue act hypotheses,

¹See also <https://dit.uvt.nl/>.

Usr: Hi, I need a Thai restaurant in the city centre
 SOCIAL: GREET; TASK: INFORM; TURN: RELEASE

Sys: Okay, let me see, ...
 TURN: TAKE; AUTOFEEDBACK: AUTOPOSITIVE
 TIME: PAUSING; TASK: INFORMSEARCH

**Sys: Bangkok City is a Thai restaurant;
it is in the city centre**
 AUTOFEEDBACK: INFORM; TASK: INFORM

Figure 1: An example of multiple dimensions in a dialogue: the user both greets the system and asks for a cheap Indian restaurant, before releasing the turn; the system then takes the turn while giving positive feedback, and indicates that it needs some time to retrieve the requested information; in the second part the system both provides this information and gives feedback about understanding the user’s question (underlined).

updates the dialogue state and then selects a response in the form of one or more dialogue acts. Rather than selecting a single action from one set of possible actions, our DM consists of multiple *dialogue act agents*, each of which selects an action from a separate action set, associated with one dimension. These action sets are based on three of the ten dimensions defined in the ISO standard for dialogue act annotation (ISO, 2012): Task (e.g. recommending a restaurant), AutoFeedback (e.g. asking the user to repeat/rephrase after a processing problem), and Social Obligations Management (SOM; e.g. responding to the user saying goodbye). These dimensions were considered to be the most important for supporting the kind of task-oriented dialogues targeted (see Fig. 1 for an example). While the Task dimension is domain-specific, AutoFeedback and SOM are applicable across domains.

Training the statistical DM on these three dimensions involves optimising three policies in parallel. A set of priority rules is used to combine the output of these policies into a single system response. The key advantage of such a design is that the domain-independent policies (AutoFeedback and SOM) can be transferred and adapted to a new domain, leaving only the Task policy to be trained from scratch. In our previous work (Keizer and Rieser, 2017), we have shown that a multi-dimensional DM with pre-trained policies reaches higher performance levels during the early stages of training. Here, we take an important step in confirming this advantage in a real user study.

	Restaurants	Hotels
#venues	149	39
#slots	4	5
shared slots	pricerange, area, near	
other slots	cuisine	type, rating

Table 1: Overview of task domains.

Our framework currently supports information-seeking domains, such as recommending restaurants or hotels based on the user’s preferences. The domains are specified in terms of an ontology (describing slots such as price range and cuisine) and a database. Our domains are presented in Table 1. We use restaurant information as target domain, but two of the system variants were trained for the hotels domain (source) and then adapted to the restaurant domain.

2.1 Model Variants

For the evaluation, we follow Keizer and Rieser (2017)’s four DM variants and training regime: The one-dimensional *one-dim* baseline system contains a single dialogue act agent (ALL) and the corresponding policy was trained from scratch in the target domain. The multi-dimensional systems use three dialogue act agents, one of which is domain-specific (TASK) and the other two domain-general (AUTOFEEDBACK and SOM). For the base *multi-dim* system, the three policies are trained from scratch in the target domain, whereas the *trans-fixed* and *trans-adapt* variants employ transfer learning (Pan and Yang, 2010; Torrey and Shavlik, 2010): only the task-specific policy is trained from scratch and the two domain-general policies are previously trained in the source domain. For *trans-fixed*, the pre-trained policies are kept fixed during training in the target domain, whilst for *trans-adapt*, these are further trained in the target domain. The four fully trained DM versions are outlined in Table 2.

2.2 Training Details

All policies are optimised in simulation using multi-agent reinforcement learning with linear value function approximation, based on a single reward signal shared between the agents.² To train all systems,

²The reward function, shared among the agents/dimensions, was the following: (i) a reward of +80 upon task completion, (ii) a penalty of -1 for each turn, (iii) a reward of +3 when responding appropriately to a social act, and (iv) a penalty of -5 when not signalling a perception or interpretation level processing problem to the user when it occurred.

we use the agenda-based user simulator of Keizer and Rieser (2017), which is based on (Schatzmann et al., 2007), along with the following error model: In addition to creating an n-best list of user dialogue act hypotheses from the ‘true’ user act, we also occasionally insert so-called ‘processing problems’, at the levels of perception (no ASR results received) or interpretation (ASR successful, but no NLU results received). We simulate a perception problem with 10% probability, and in case of no perception problem (90%), we simulate an interpretation problem with 10% probability; only in case no processing problems are generated (81%), an n-best list of dialogue act hypotheses is generated. Following Thomson et al. (2012), the n-best lists are populated by taking the true user act and distorting it at a given semantic error rate for each of the positions, after which semantically equivalent hypotheses are merged. Based on the error rate, a Dirichlet distribution is used to generate confidence scores for the n-best list (resulting in a semantic top accuracy equal to the error rate), interpreted as probabilities by the DM when updating its user goal belief state.³

In order to correctly interpret the evaluation results, note that in the current setup, the *one-dim* system serves as an upper bound baseline system, as it needs no coordination between different agents during training whilst generating (by construction) the same range of actions as the multi-dimensional systems. This is ensured by a set of priority heuristics which map action combinations to single acts.⁴

2.3 DM Evaluation in Simulation

To get a better picture of what we might expect during the human evaluation, we first ran evaluations with simulated data. The results obtained with the same settings as those during training are shown in Table 3. As we hypothesised, the scores are very similar, the *one-dim* system only slightly outperforming the multi-dimensional systems.

We then extended the setup with different semantic error rates (Thomson et al., 2012); the results are shown in Fig. 2. The performance levels of the

³For each of the four DM versions, 5 training runs over 60k dialogues were carried out, resulting in a pool of 5 fully trained policies.

⁴The n-best size was set to 3 and the error rate was set to 30% for the target domain (restaurants) and 20% for the source domain (hotels).

⁵E.g. if the Task agent generates a recommendation action and the AutoFeedback agent generates a negative feedback action, the latter gets priority and the former is cancelled.

four systems are very similar at error rates between 10% and 40%, showing that the construction of the multi-dimensional versions in relation to the *one-dim* baseline is sound, and showing there is no negative transfer, i.e., the adapted systems are not performing worse.⁵

3 Evaluation Setup

We use crowdsourcing to evaluate our system, following Jurčiček et al. (2011) and Crook et al. (2014). In both of these works a phone-based system was deployed, using a bespoke ASR and Voice over IP (VoIP) to connect speech input/output with the dialogue system. Here, we follow a similar evaluation methodology, but with a novel, simpler web-based interface using Google Chrome’s built-in web speech API, embedded into the crowdsourcing task webpages. A detailed description of the technical setup can be found in Appendix A.

3.1 Crowdsourcing Setup

The users are recruited on the FigureEight crowdsourcing platform and asked to have a conversation with the system to find a venue meeting certain criteria (e.g. cheap Chinese food) and get certain information about that venue (e.g. phone number and address). This scenario is specified in natural language, generated automatically from a set of task specifications randomly generated from the domain ontology. After each conversation, the user is given a questionnaire to rate the system.

3.2 Evaluation Metrics

The subjective evaluation metrics are derived from the following questionnaire, with one yes/no question (Q1) and four 6-point Likert Scale ratings.

Q1 [SubjSucc]: Did you find all the information you were looking for?

Please state your attitude towards the following statements:

Q2 [VoiceInt]: The system was easy to understand (the voice was intelligible).

Q3 [Understand]: In this conversation, the system understood what you said.

Q4 [AsExpect]: The system worked the way you expected it to during the conversation.

Q5 [WdUseAgain]: From your experience with the system, you think you would use it in the future to find a place to eat.

⁵The discrepancy at zero error rate for the trans-fixed system might have occurred because certain state feature combinations occurring specifically at zero error rate were not seen during training, and might be too distinct to be dealt with by the generalisation capability of the value approximation model used in our reinforcement learning algorithm.

Dialogue Act Agent	<i>one-dim</i>	<i>multi-dim</i>	<i>trans-fixed</i>	<i>trans-adapt</i>
ALL	<i>source</i> : – <i>target</i> : trained	–	–	–
TASK	–	<i>source</i> : – <i>target</i> : trained	<i>source</i> : – <i>target</i> : trained	<i>source</i> : – <i>target</i> : trained
AUTOFEEDBACK	–	<i>source</i> : – <i>target</i> : trained	<i>source</i> : trained <i>target</i> : fixed	<i>source</i> : trained <i>target</i> : adapted
SOM	–	<i>source</i> : – <i>target</i> : trained	<i>source</i> : trained <i>target</i> : fixed	<i>source</i> : trained <i>target</i> : adapted

Table 2: Evaluated systems: *one-dim* is a one-dimensional (upper) baseline, other systems are multi-dimensional.




Figure 2: Results in simulation at different error rates.

system	SuccRate	AvgLen	AvgRew
<i>one-dim</i>	97.8%	14.69	66.36
<i>multi-dim</i>	97.6%	15.68	64.97
<i>trans-fixed</i>	96.8%	15.08	65.23
<i>trans-adapt</i>	97.4%	16.41	64.20

Table 3: Test results on simulated data (same error rates as in training): task success rate (SuccRate), average dialogue length (AvgLen), average reward (AvgRew).

DM version	NumDials	NumTurns (StDev)
<i>one-dim</i>	245	6.67 (2.55)
<i>multi-dim</i>	228	6.30 (1.97)
<i>trans-fixed</i>	261	6.57 (2.33)
<i>trans-adapt</i>	248	6.64 (2.33)
Total	982	6.55 (2.31)

Table 4: Corpus statistics: the number of dialogues collected (NumDials) and the average number of turns per dialogue (NumTurns) with standard deviation (StDv).

The following objective success metrics are derived from the logs:

EntProv: the system recommended an entity matching the task constraints,

ConstrConf: the system confirmed all task constraints in its recommendation,

InfoProv: the system provided all information requested by the user.

4 Human User Evaluation

In total, 982 dialogues were collected (see Table 4), i.e. 246 dialogues per system variant on average.

We carried out a number of statistical tests to analyse the observed effect sizes in comparing the systems, including chi-squared (for success rates) and Mann-Whitney tests (for the Likert scale ratings), but also the ‘two one-sided test’, or TOST (Schuirmann, 1987), for *equivalence*, as argued in Section 2.1. In a TOST scenario, the null hypothesis is that the difference in performance between two systems, Δ , is greater than a given threshold ϵ (a hyperparameter). This translates into two one-sided null hypotheses:

$$H_{lo} : \Delta \leq -\epsilon \quad (1)$$

$$H_{hi} : \Delta \geq +\epsilon \quad (2)$$

If both H_{lo} and H_{hi} are rejected, we can conclude that $-\epsilon < \Delta < +\epsilon$, i.e. the difference lies below the threshold. This test is much more conservative than failing to reject the null hypothesis in a conventional statistical test of significant difference. The underlying one-sided tests can differ according to the nature of data at hand. The default proposed by Schuirmann (1987) is t-tests. However, our data fails the normal distribution assumption of a t-test. Therefore, we use the robust t-test of Yuen and Dixon (1973) for testing equivalence on Likert scale data, which does not assume normality, and a pooled z-test with continuity correction (Fleiss et al., 2003, p. 53ff.) for success rates.⁶ We used a

⁶The z statistic is the square root of the χ^2 statistic, which is more suited for determining standard deviation (i.e. size of difference) as opposed to variance.

DM	SubjSucc [Q1]	VoiceInt [Q2]	Underst [Q3]	AsExpect [Q4]	WdUseAgain [Q5]	EntProv	ConstrConf	InfoProv
<i>one-dim</i>	87.3%	5.49	4.80	4.81	4.67	72.2%	57.7%	45.7%
<i>multi-dim</i>	83.3%	5.37	4.68	4.68	4.59	68.4%	52.7%	44.7%
<i>trans-fixed</i>	81.6%	5.47	4.66	4.64	4.63	70.1%	53.1%	41.0%
<i>trans-adapt</i>	85.9%	5.38	4.67	4.64	4.57	72.2%	53.1%	46.6%

Table 5: Overview of subjective and objective evaluation results (cf. Section 3.2 for metrics).

DM version	NumDials	WER
<i>one-dim</i>	120	17.2%
<i>multi-dim</i>	124	15.6%
<i>trans-fixed</i>	137	15.4%
<i>trans-adapt</i>	115	19.1%

Table 6: WER analysis results (NumDials indicates the number of dialogues transcribed for each system).

threshold of $\epsilon = 10\%$ for the equivalence tests.

4.1 Evaluation Results

Table 5 shows the results for both objective and subjective metrics. When considering the metrics for task success (*SubjSucc*, *EntProv*, *ConstrConf*, *InfoProv*), the *one-dim* system is the highest scoring, although the *trans-adapt* system is often a close second and in some cases the top scorer. However, no statistically significant differences were detected, and the *one-dim* system was moreover found to be equivalent to the *multi-dim* ($p = 0.024$) and *trans-adapt* ($p = 0.002$) systems in perceived success (*SubjSucc*), and all three multi-dimensional systems were found to be equivalent to each other ($p = 0.006, 0.009, \text{ and } 0.031$). Similarly, several equivalences were detected for the three objective success metrics, as illustrated in Appendix B.⁷ All systems are equivalent on the other subjective ratings Q2–Q5.

To get a sense of the noise levels encountered by the different system variants, we collected crowd-sourced transcriptions of 2,931 utterances from 496 dialogues (45.6% of the total number of turns in the evaluation corpus and 50.5% of collected dialogues), spread approximately evenly across all system variants. We then computed word error rate (WER).⁸ Results in Table 6 show comparable noise

⁷Following Armstrong (2014), we do not apply a correction for multiple comparisons (Lauzon and Caffo, 2009) since we only performed a limited number of pre-planned comparisons and did not require testing against the universal null hypothesis “nothing is significant”.

⁸The reference transcriptions were obtained by majority voting over the three transcriptions collected for each utterance, with manual fixes in case of a tie (20% of the utterances).

levels for all system variants. No significant differences were found and equivalence tests confirmed WER to be equivalent for all the systems. This confirms that none of the systems was disadvantaged and the results in Table 5 are indeed comparable.

5 Conclusion and Future Work

In this paper, we have shown that a multi-dimensional, data efficient dialogue manager performs equally to a one-dimensional, more data-hungry (upper) baseline. In doing so, we have developed a web-based platform for spoken dialogue system evaluation, carried out a crowdsourced user evaluation, and introduced statistical testing for equivalence in our analysis of the results. All code and data used in our experiments are available at:

<https://bitbucket.org/skeizer/madrigal>

The results show that none of the systems outperformed the other systems consistently across various metrics, and more importantly, that several statistical equivalences between the systems could be detected. We believe that these results are encouraging, especially since we suspect that the use of a web-based speech interface (with inherently varying quality of the microphone used) and the crowdsourcing setup (with inherently varying conditions in which workers do their tasks) resulted in a relatively high level of variance in the data, making it harder to draw strong conclusions.

In the next stage of our research, we aim to further demonstrate the cross-domain transfer capability of the dialogue manager, for example by evaluating partially trained policies, and showing that policies that use transfer learning reach higher performance levels in the early stages of training, or that they achieve a given performance threshold with much less data.

Acknowledgements

This research was supported by the EPSRC project MaDrIgAL (EP/N017536/1) and Charles University project PRIMUS/19/SCI/10.

References

- Rieks op den Akker, Harry Bunt, Simon Keizer, and Boris van Schooten. 2005. [From Question Answering to Spoken Dialogue: Towards an Information Search Assistant for Interactive Multimodal Information Extraction](#). In *Proceedings of the 9th European Conference on Speech Communication and Technology (Interspeech)*, Lisbon, Portugal.
- Richard A. Armstrong. 2014. [When to use the Bonferroni correction](#). *Ophthalmic and Physiological Optics*, 34(5):502–508.
- Harry Bunt. 2011. [Multifunctionality in dialogue](#). *Computer Speech & Language*, 25(2):222–245.
- Paul A Crook, Simon Keizer, Zhuoran Wang, and Wenshuo Tang. 2014. [Real user evaluation of a POMDP spoken dialogue system using automatic belief compression](#). *Computer Speech & Language*, 28(4):873–887.
- Joseph L. Fleiss, Bruce Levin, and Myunghee Cho Paik. 2003. *Statistical methods for rates and proportions*, 3rd edition. Wiley series in probability and statistics. J. Wiley, Hoboken, NJ, USA.
- Milica Gašić, Nikola Mrkšić, Lina M. Rojas-Barahona, Pei-hao Su, Stefan Ultes, David Vandyke, Tsung-Hsien Wen, and Steve Young. 2017. [Dialogue manager domain adaptation using Gaussian process reinforcement learning](#). *Computer Speech & Language*, 45:552–569.
- ISO. 2012. *ISO 24617-2 Language resource management – Semantic annotation framework – Part 2: Dialogue acts*. International Organization for Standardization, Geneva, Switzerland.
- Filip Jurčićek, Simon Keizer, Milica Gašić, François Mairesse, Blaise Thomson, Kai Yu, and Steve Young. 2011. [Real user evaluation of spoken dialogue systems using Amazon Mechanical Turk](#). In *Proceedings of Interspeech*, Florence, Italy.
- Simon Keizer and Verena Rieser. 2017. [Towards Learning Transferable Conversational Skills using Multi-dimensional Dialogue Modelling](#). In *Proceedings 21st Workshop on the Semantics and Pragmatics of Dialogue (SemDial/SaarDial)*, Saarbruecken, Germany. Extended version: arXiv:1804.00146.
- Carolyn Lauzon and Brian Caffo. 2009. [Easy Multiplicity Control in Equivalence Testing Using Two One-Sided Tests](#). *The American Statistician*, 63(2):147–154.
- Sinno Jialin Pan and Qiang Yang. 2010. [A Survey on Transfer Learning](#). *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359.
- Volha Petukhova, Christopher Stevens, Harmen de Weerd, Niels Taatgen, Fokje Cnossen, and Andrei Malchanau. 2016. [Modelling multi-issue bargaining dialogues: Data collection, annotation design and corpus](#). In *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC)*, Paris, France.
- Jost Schatzmann, Blaise Thomson, Karl Weilhammer, Hui Ye, and Steve Young. 2007. [Agenda-based user simulation for bootstrapping a POMDP dialogue system](#). In *Human Language Technologies: The Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT)*, Rochester, NY, USA.
- Donald J Schuirmann. 1987. [A comparison of the two one-sided tests procedure and the power approach for assessing the equivalence of average bioavailability](#). *Journal of Pharmacokinetics and Biopharmaceutics*, 15:657–680.
- Matthew E Taylor and Peter Stone. 2009. [Transfer Learning for Reinforcement Learning Domains: A Survey](#). *The Journal of Machine Learning Research*, 10(Jul):1633–1685.
- Blaise Thomson, Milica Gašić, Matthew Henderson, Pirros Tsiakoulis, and Steve Young. 2012. [N-best error simulation for training spoken dialogue systems](#). In *Spoken Language Technology Workshop (SLT)*, Miami, FL, USA.
- Lisa Torrey and Jude Shavlik. 2010. [Transfer learning](#). In *Handbook of research on machine learning applications and trends: algorithms, methods, and techniques*, pages 242–264. IGI Global.
- Zhuoran Wang, Tsung-Hsien Wen, Pei-hao Su, and Yannis Stylianou. 2015. [Learning domain-independent dialogue policies via ontology parameterisation](#). In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGdial)*, Prague, Czechia.
- Tsung-Hsien Wen, Milica Gašić, Nikola Mrkšić, Lina M. Rojas-Barahona, Pei-Hao Su, David Vandyke, and Steve Young. 2016. [Multi-domain neural network language generation for spoken dialogue systems](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 120–129, San Diego, CA, USA.
- Jason D Williams, Eslam Kamal, Hani Amr Mokhtar Ashour, Jessica Miller, and Geoff Zweig. 2015. [Fast and easy language understanding for dialog systems with Microsoft language understanding intelligent service \(LUIS\)](#). In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, Prague, Czechia.
- Steve Young, Milica Gašić, Blaise Thomson, and Jason D. Williams. 2013. [POMDP-Based Statistical Spoken Dialog Systems: A Review](#). *Proceedings of the IEEE*, 101(5):1160–1179.
- Karen K. Yuen and W. J. Dixon. 1973. [The approximate behaviour and performance of the two-sample trimmed \$t\$](#) . *Biometrika*, 60(2):369–374.

A Dialogue System Setup

An overview of our crowdsourced dialogue system evaluation setup is shown in Fig. 3. The core component of the spoken dialogue system is the Dialogue System Server, which contains the DM (see Section 2), extended with a template-based NLG component and code for processing NLU results from Microsoft’s LUIS (Williams et al., 2015). Our LUIS model was trained with 299 manually constructed and annotated example utterances.

The system is completed by a web-based user interface, which connects with both the Dialogue System Server and the Google Web Speech API.⁹ User audio input is first sent to Google ASR to get user utterance hypotheses with confidence scores. These are sent to the Dialogue System Server, which returns a system response utterance. Finally, this utterance is sent to Google TTS, which returns the synthesised system response audio to be played back to the user. The web interface is integrated into the FigureEight crowdsourcing platform for managing the evaluation (Section 3.1).

B Equivalence test results

See Figure 4 for a diagram of all statistically significant equivalences that we detected with respect to the individual evaluation criteria (see Sections 3.2 and 4).

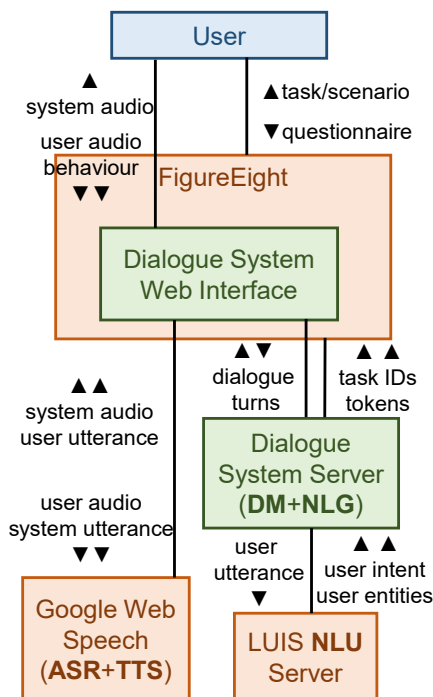


Figure 3: Overview of dialogue system evaluation setup.

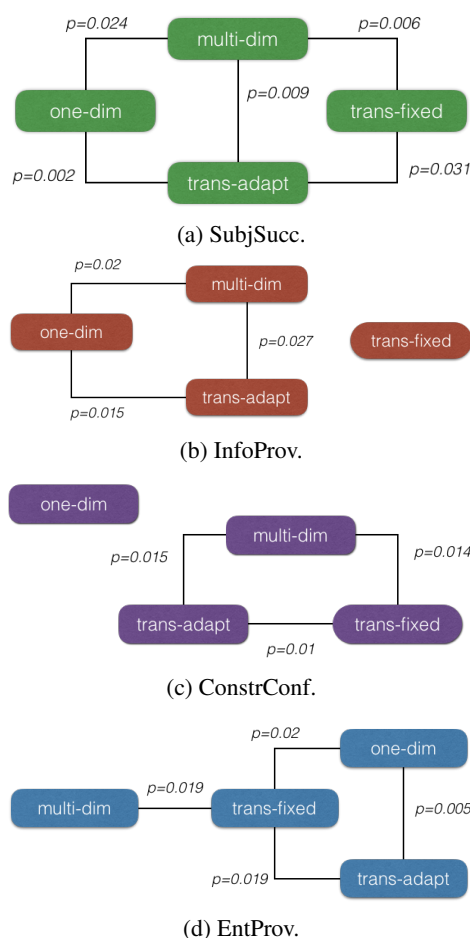


Figure 4: Statistically significant equivalences detected.

⁹<https://w3c.github.io/speech-api/speechapi.html>

Dialogue Act Classification in Team Communication for Robot Assisted Disaster Response

Tatiana Anakina and Ivana Kruijff-Korbayová

German Research Institute for Artificial Intelligence (DFKI) / Saarbrücken, Germany

tatiana.anikina@dfki.de, ivana.kruijff@dfki.de

Abstract

We present the results we obtained on the classification of dialogue acts in a corpus of human-human team communication in the domain of robot-assisted disaster response. We annotated dialogue acts according to the ISO 24617-2 standard scheme and carried out experiments using the FastText linear classifier as well as several neural architectures, including feed-forward, recurrent and convolutional neural models with different types of embeddings, context and attention mechanism. The best performance was achieved with a "Divide & Merge" architecture presented in the paper, using trainable GloVe embeddings and a structured dialogue history. This model learns from the current utterance and the preceding context separately and then combines the two generated representations. Average accuracy of 10-fold cross-validation is 79.8%, F-score 71.8%.

1 Introduction

Disaster response teams operate in high risk situations and face critical decisions despite partial and uncertain information. First responders increasingly deploy mobile robotic systems to mitigate risk and increase operational capability. In order for robotic systems to provide optimal support for mission execution, they need mission knowledge, i.e., run-time awareness and understanding of the mission goals, team composition, the tasks of the team(s), how and by whom they are being carried out, the state of their execution, etc. Since first responders typically operate under high cognitive load and time pressure, it is paramount to keep the burden of entering mission knowledge into the system at a minimum. The goal of our research thus is to develop methods for *extracting run-time mission knowledge from the verbal communication in the response team*. The acquired mission knowledge can also be used to assist the

first responders during or after the mission, for example, by supporting the real-time coordination of human and robot actions or by mission documentation generation (Willms et al., 2019).

In this paper we address one particular sub-problem: dialogue act (DA) recognition. DAs are needed for a better understanding of the team communication and how the mission tasks are being executed. For example, *Requests* communicate task assignments and thus allow us to distinguish task assignments from other task-relevant information exchange; *Informs* often report task progress; and *Questions* indicate what was unclear and required additional explanations. These distinctions are also useful for providing assistance, including compiling mission documentation.

We use the corpus of human-human team communication in robot-assisted disaster-response collected in the TRADR project (Kruijff-Korbayová et al., 2015). The TRADR team communication is task-oriented, focused on collaborative execution of a mission by a structured team using mobile robots to remotely gather situation awareness in a complex, dynamic, unknown physical environment. In this the communication differs from that in well-known existing corpora annotated with DAs.

We annotated our corpus with DAs following the ISO 24617-2 scheme (Bunt et al., 2012, 2017) and experimented with several machine learning approaches to DA classification. We explored various models, including different ways of taking dialogue context into account.

We overview previous work on DA classification and existing corpora with DA annotations in §2. We present our corpus in §3 and provide statistics for DA and speaker role distribution. In §4 we describe the classification models tested in our experiments and report the evaluation results. We conclude with a discussion and future plans in §5.

2 Related Work

There is a body of research on teamwork and information sharing in disaster response, with and without robots, e.g., (Casper and Murphy; Burke et al.; Burke and Murphy; Johnson et al., 2017; Toups et al., 2016; Carver and Turoff, 2007).

There has been very little work on dialogue processing in this domain so far. In the pioneering project TRIPS a decision-support dialogue system was developed for the planning of an island evacuation in the event of a natural disaster. Focus was on semantic parsing and task-specific interpretation. This approach was further developed to handle various more complex emergency tasks covered in the Monroe corpus (Stent, 2000). This work focused on mission planning (not execution), data was collected in lab (not real disaster environment) and the participants were students (not real first responders). DAs were annotated using the DAMSL scheme (Core and Allen, 1997).

Some works on human-robot collaboration for disaster response address the interpretation of verbal commands to robots (Kruijff et al., 2014; Yazdani et al., 2018), but not the overall team communication.

In (Martin and Foltz, 2004) automatic analysis of the semantic content of team communication and automatic verbal behavior labeling was used to assess team performance in a command and control task with an unmanned aerial vehicle in a simulated environment. A corresponding synthetic team-member agent is described in (Cooke et al., 2016). Since the corpus is not available and the publications do not provide details on the task and communication complexity, a closer comparison to our work is not possible. Communication analysis was used also in (Burke et al.). They designed and manually applied a team communication coding scheme, in order to examine robot operator situation awareness and technical search team interaction during a high-fidelity disaster response drill with teleoperated robots. DAs are reflected in their annotation of the forms and functions of communication contributions.

Corpora with DA annotations include also well-known human-human dialogue corpora, such as MapTask (Anderson et al., 1991; Carletta et al., 1997); TRAINS (Allen, 1991); Switchboard (Godfrey et al., 1992); Meeting Recorder Dialogue Act (Shriberg et al., 2004) and the AMI Meeting Corpus (Carletta et al., 2005), and re-

cent large corpora, e.g., Maluuba Frames (Schulz et al., 2017) and MultiWOZ (Budzianowski et al., 2018)). These corpora cover different domains and the goals the participants follow in their interaction are quite different from what is going on in the team communication in our domain.

Despite the differences it would be interesting to see how DA classification models developed on other exiting corpora perform on our corpus. The challenge of such endeavor is, however, that different and sometimes very task-specific schemes have been applied to annotate DAs. For instance, some of the DAs in the Maluuba Frames corpus include domain-specific labels such as *Canthelp* and *No_result* as well as *Thankyou* and *Moreinfo*.

The ISO 24617-2 standard for DA annotations introduced in (Bunt et al., 2012) and further defined in (Bunt et al., 2017) was proposed to overcome this. To date several corpora have been annotated accordingly and made available through the DialogBank (Bunt et al., 2016). Although the mapping of DA labels from other annotations to the ISO standard is quite straightforward in some cases (e.g., for *Inform* or *Request*), in other cases the specificity of the domain prevents from further generalizations, as discussed in (Chowdhury et al., 2016). These issues lead us to postpone transfer learning for future work and start traditionally by experiments on our own corpus.

Previous work on automatic DA classification includes the use of Hidden Markov models (Stolcke et al., 2000), Maximum Entropy (Choi et al., 1999), Generative and Conditional Bayesian Networks (Ji and Bilmes, 2005), and Support Vector Machines (Quarteroni and Riccardi, 2010). Recent papers also explored neural architectures (Kumar et al., 2017; Liu et al., 2017) and compared word embeddings (Cerisara et al., 2018).

Only few works to date systematically tested different kinds of context for DA classification. Several experiments on the Switchboard corpus are described in (Ribeiro et al., 2015), which tested untagged and index-tagged n-grams as well as context presented in the form of dialog act annotations for the previous segments. Index-tagged n-grams (n-grams tagged with the distance to the current segment) improved accuracy significantly, from 70.6% to 75.1%, and the DA annotations for the preceding segments even to 76.4%.

(Liu et al., 2017) tested different kinds of context for DA classification using deep neural mod-

els. They present hierarchical models based on convolutional neural networks (CNN) for sentence representations which they combine with dialogue history. They encode context as previous DA labels and as probabilities for system predictions, and experiment with dialogue history of varied length. Including context information in their models evaluated on the Switchboard corpus resulted in significant increase of accuracy from 77% to almost 80%. These results indicate that context should be taken into account when processing structured conversations.

3 The Corpus

We use the corpus of robot-assisted disaster-response team communication collected during joint exercises with first responders in the TRADR project (Kruijff-Korbayová et al., 2015).¹ The TRADR corpus contains audio recordings and transcriptions of the speech communication in a team of firefighters using robots in the aftermath of an incident, e.g., an explosion, at an industrial site. The team members have various roles: mission commander (MC), team leader (TL), operators (OP) of multiple ground (UGV) and aerial (UAV) robots. They explore the site, searching for persons, hazard sources, fires and other relevant points of interest. The MC and the TL lead the mission. They request situation information from the OPs, who report back with updates and can also share photos taken by the robot camera (see the example in Appendix A).

The recordings were collected during several field tests in 2015, 2016 and 2017. They amount to approximately 10 hours and contain almost 3k speech turns (see Table 1 for details). The 2015 and 2016 recordings are in German, the 2017 ones in English. For the experiments presented in this paper we used the original English data as well as English translations from German. We started on English because of available resources.

Before annotating DAs following the ISO 24617-2 scheme (Bunt et al., 2012, 2017), we segmented the data into *utterances*; we split and merged some turns to obtain appropriate spans for assigning DAs. This resulted in 2469 utterances.

The ISO scheme defines several dimensions and for each of them a hierarchy of commu-

¹The TRADR team communication corpus is available online from www.tradr-project.eu/resources/datasets/ or talkingrobots.dfki.de/resources/tradr/

Recording	Mission	Duration	Turns
TJex 2015	Day 1	48:21 min	374
	Day 2	33:21 min	201
			173
TEval 2015	Day 1	58:23 min	1165
	Day 2	65:04 min	289
	Day 3	57:15 min	299
	Day 4	53:22 min	219
			358
TEval 2016	Day 1	n.a.	421
	Day 2	n.a.	311
			110
TEval 2017	Day 1	64:02 min	822
	Day 2	149:20 min	240
	Day 3	56:36 min	408
			174
Total:			2782

Table 1: Corpus composition

ISO Annotation Label	Classification Label
Turn Management	Contact
Inform, Promise, Offer, Address-Suggestion	Inform
PositiveFeedback, AcceptRequest, AcceptOffer, AcceptSuggestion, Agreement	Affirmative
Request	Request
CheckQuestion, SetQuestion, ChoiceQuestion, Question	Question
Confirm	Confirm
Disconfirm	Disconfirm
Negative Feedback, DeclineOffer, Disagreement	Negative

Table 2: Mapping of ISO annotation labels to labels for automatic classification

nicative functions (a.k.a. DAs). The first author and another annotator independently annotated each utterance with one of the dimensions and a corresponding DA. Inter-annotator agreement was $\kappa=.77$ for dimension assignment and $\kappa=.55$ (weighted $\kappa=.66$) for the generic communicative functions in the *Task* dimension. For the experiments in this paper we used the first author’s annotations as a golden reference. We focused on the classification of DAs from the dimensions *Task*, *Feedback* and *Turn Management* (see Table 2 for the used labels).

We annotated the corpus in full compliance with the ISO scheme. Since some DAs had too few occurrences in the corpus we used a simplified set of DA labels in the experiments (see §4.1). The simplified labels are a result of a direct mapping from the ISO scheme labels (see Table 2), making it easy to compare DA classification results. In most cases the simplified labels can be seen as

Dialogue Act	MC	TL	OP	Total
Contact	32	350	360	742
Inform	19	132	476	627
Affirmative	8	217	127	352
Request	9	262	3	274
Question	12	150	84	246
Confirm	2	28	131	161
Disconfirm	0	4	49	53
Negative	0	6	8	14

Table 3: Dialogue act distribution

generalized ISO DAs which were selected based on their utility for the disaster response domain.

The mission interactions consist of *threads*, which are dialogue sequences where two (occasionally multiple) team members talk about a task or situation update, e.g., the TL talks to an OP as illustrated in the example in Appendix A. A new thread is initiated by establishing contact following the standard radio communication protocol. The threads are a good candidate for dialogue context and we used thread history in some experiments as we will describe in the next sections.

4 Experiments

4.1 Pre-processing

Before running the experiments we pre-processed the data as follows.

First, we collapsed DA labels which had very low frequency in the corpus with more frequent ones. For instance, there were only 2 cases of *AddressSuggestion* and 9 cases of *AcceptOffer* in total. Low frequency labels would introduce noise and prevent the classifier from learning reliable patterns. Moreover, there were some ambiguous cases with several possible annotations (e.g. *Inform* and *Promise* for "I'll send it over to you") and we decided to retain the most frequent label to reduce the perplexity. Table 2 shows the mapping of the manually annotated ISO scheme labels to the DA labels used for the automatic classification. The resulting distribution of DA labels is shown in Table 3.

Second, we removed all punctuation. Although punctuation can be a good clue for some DAs (e.g., "?" usually indicates *Question*) we removed it, because the ASR software often does not provide punctuation reliably. We also transformed all texts to lower case and padded sequences when using neural networks. For 10-fold cross-validation we split the 2469 utterances into 2222 for training and 247 for testing in each fold partition.

4.2 Baselines

We implemented three baselines. The **majority baseline** assigned each utterance the most frequent label for the given role, i.e., all MC/TL utterances were annotated as *Contact* and all OP utterances as *Inform*. This resulted in accuracy 34.8%

The fact that all TL utterances were classified as *Contact* was an obvious drawback. We therefore tried a **relative-frequency baseline** as an alternative, using the relative frequencies for each DA on the complete corpus (cf. Table 3). Each utterance was assigned a random class based on the relative frequencies. This baseline had accuracy 24.7%². The majority baseline which used solely the role was substantially better compared to the frequency-based random baseline.

The third **mixed baseline** was based on the assumption that all instances of *Contact* are identified correctly and for all other utterances we used the majority baseline. Therefore, the third baseline assigned *Request* to all MC/TL utterances and *Inform* to all OP utterances which were not labeled as *Contact*. This baseline had accuracy 47.2%. Since these three baselines had such a low performance we considered the results of the FastText classifier as a baseline for evaluating the performance of the neural models.

4.3 FastText

As the first model for DA classification we tested FastText³, an open-source library for text classification and representation using supervised learning with multinomial logistic regression. Although it can represent input text in the form of embeddings it belongs to the family of linear classifiers. We ran FastText using the parameters recommended for a small training set (10 dimensions, 0.5 learning rate, 20 epochs). The average accuracy over a 10-fold cross-validation was 74.0%. It was consistent across the folds (see Table 4). Because of the strong correlation between the speaker role and the DA distribution, as shown in Table 3, we also experimented with including the role as a special token at the beginning of each utterance. This additional information improved the average accuracy to 75.6% and also the accuracy in most folds (see Table 4). Finally, we tested the effect of adding the dialogue thread con-

²We also tested a baseline based on DA relative frequencies per role, but the accuracy was even lower, 21%.

³<https://fasttext.cc/>

Fold	Accuracy without Role	Accuracy with Role	Accuracy with Role + Thread History
1	0.656	0.668	0.628
2	0.607	0.684	0.583
3	0.668	0.696	0.583
4	0.745	0.757	0.583
5	0.834	0.858	0.692
6	0.761	0.741	0.640
7	0.794	0.773	0.709
8	0.781	0.785	0.660
9	0.769	0.794	0.676
10	0.785	0.801	0.650
Avg:	0.740	0.756	0.640

Table 4: FastText 10-fold cross-validation

text: we appended the corresponding thread history to each utterance and trained FastText on this extended input. Accuracy dropped for all folds, to 64.0% on average as shown in Table 4.

4.4 Neural Networks

Neural networks have already shown great potential in tackling various NLP tasks, including DA classification (Chen et al., 2018; Liu et al., 2017). We therefore also tested various neural architectures to classify DAs in our corpus: Feed-Forward Neural Networks (FFNN); Recurrent Neural Networks (RNN), in particular Long-Short Term Memory (LSTM) and bidirectional LSTM models; Convolutional Neural Networks (CNN). We experimented with attention and different kinds of embeddings (including Word2Vec, GloVe and FastText). We also tested the effect of the dialogue context in the form of the preceding thread history concatenated with the current utterance. We present the models and the DA classification results in the next sections.

Feed-Forward Neural Networks

We implemented a simple FFNN using the Keras⁴ library with one Embedding layer (we experimented with 100, 200 and 300 dimensions) and applied global average pooling to average the embeddings of all words in the utterance before sending them through the Dense layer. The architecture is shown in Figure 1.

We set the minibatch size to 8, trained the network for 5 epochs and used Adam as an optimizer. We trained several models using the Embedding layer provided by Keras as well as pre-trained GloVe embeddings obtained from the Stanford

⁴<https://keras.io/>

Embeddings Type	Accuracy
Keras 100	0.755
Keras 200	0.761
Keras 300	0.762
GloVe 100, frozen	0.685
GloVe 200, frozen	0.711
GloVe 300, frozen	0.722
GloVe 100, trainable	0.759
GloVe 200, trainable	0.768
GloVe 300, trainable	0.771

Table 5: DA classification results for FFNNs with different types of embeddings

NLP group website,⁵ which were learnt on the data from Wikipedia 2014 and Gigaword 5 (6B tokens, 400K vocabulary). We also experimented with both frozen and trainable embeddings. The results were consistently better with trainable embeddings compared to the frozen version. Table 5 shows the evaluation results with accuracy scores averaged across 10 folds.

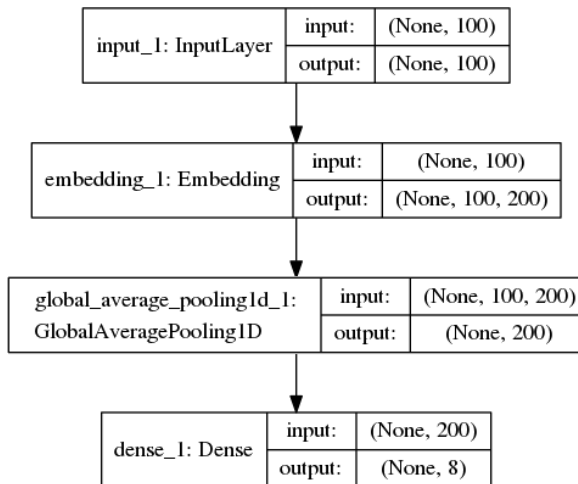


Figure 1: Feed-Forward Network with embeddings⁶

Convolutional Neural Networks

Inspired by the results on DA classification with CNNs in (Liu et al., 2017) we also tested CNNs with varying number of convolutional layers and filter sizes on our data. Figure 2 shows a sample architecture with two convolutions and 128 filters of size 5. We also tested CNNs with different embeddings. The best performance (average accuracy 72.1%) was achieved by the model with one convolutional layer, filter size 10 and embeddings

⁵<https://nlp.stanford.edu/projects/glove/>

⁶None is a dynamic length dimension which means that a corresponding layer can have variable-length sequences as an input.

Embeddings Type	Conv.	Filter Size	Accuracy
Keras 100	2	5	0.685
Keras 200	2	5	0.697
Keras 100	1	10	0.721
Keras 200	1	10	0.712
GloVe 100	2	5	0.695
GloVe 200	2	5	0.694
GloVe 100	1	10	0.703

Table 6: DA classification results for CNN models

trained on our data with dimensionality 100. An overview of the results obtained with various CNN architectures is in Table 6. Interestingly, more complex models resulted in worse scores. Convolutions appear not very useful for the relatively short texts of dialogue utterances.




Figure 2: Convolutional Neural Network ⁶

Model	Embeddings Type	Accuracy
LSTM	Embedding 200	0.745
LSTM	GloVe 200	0.775
LSTM	GloVe 300	0.768
LSTM +Attention, -Thread	Embedding 200	0.676
LSTM +Attention, -Thread	GloVe 200	0.767
LSTM -Attention, +Thread	GloVe 200	0.780
LSTM +Attention, +Thread	GloVe 200	0.745

Table 7: RNN performance

Model	Embeddings Type	Accuracy
no LSTM	GloVe 200	0.784
LSTM for turn & thread	GloVe 200	0.768
LSTM for turn	GloVe 200	0.798
LSTM for turn	Word2Vec 100	0.769
LSTM for turn	Word2Vec 200	0.773
LSTM for turn	Word2Vec 300	0.767
LSTM for turn	FastText 300	0.770

Table 8: Divide&Merge performance

Recurrent Neural Networks

We tested RNNs with Long Short Term Memory (LSTM) cells, both LSTMs and bidirectional LSTMs. We also applied an attention mechanism and experimented with various embeddings and regularization parameters. In some experiments we concatenated all previous utterances from the same thread with the current utterance in order to give more context to the classifier. We inserted a *#START#* symbol between the current utterance and the thread text as a separator.

Figure 3 shows the RNN architecture with bidirectional LSTM and attention mechanism. The attention layer follows the idea proposed in (Raffel and Ellis, 2015). We passed the generated word vectors through bidirectional LSTM and multiplied the input with the attention vector at each time step. The result was passed through the Dense layer with ReLU as an activation function. Dropout 0.25 was applied to the function output before it went through the final Dense layer. We tested this model with single utterances as well as with utterances concatenated with their corresponding thread history, with and without attention. The results of different RNN architectures are in Table 7. The best accuracy of 78.0% was achieved by the model which used the thread history and pre-trained GloVe embeddings with trainable weights, no attention.

In the experiments described above we noticed

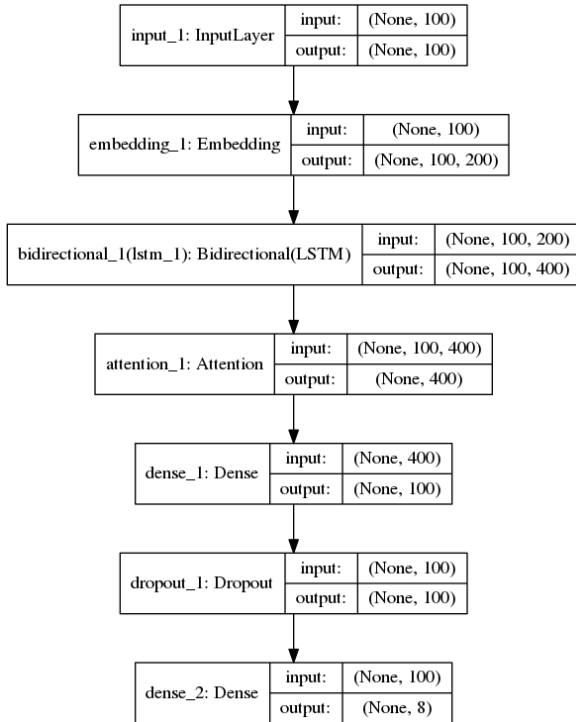


Figure 3: RNN with attention ⁶

that simple concatenation of the current utterance with the previous context gives us a very small improvement in accuracy compared to the model which does not use the thread history (accuracy increased from 77.5% to 78.0%). The network treats the current utterance and the thread history as a single input, and this might result in a sub-optimal representation. Hence, we designed a model that learns from the current utterance and from the previous context separately and then combines the two generated representations into one. Because it first separates the current utterance from the context and then puts the representations together we call this new model *Divide & Merge* (D&M). Figure 4 shows the D&M model architecture we implemented. 10-fold cross-validation yielded the best average accuracy of 79.8% using pre-trained GloVe embeddings with 200 dimensions and training for 5 epochs. Detailed results of the D&M model evaluation are in Tables 8 and 9.

Table 8 shows the results for various experimental settings. First, we report the accuracy scores obtained by the D&M model without LSTM, D&M which uses LSTM for encoding both turn and thread utterances and D&M which uses LSTM only for turns while the thread information is encoded using one Embedding layer and global average pooling as shown in Figure 4. The model

with turn-only LSTM achieved the best accuracy 79.8%. Second, we also compared different word embeddings (GloVe, Word2Vec and FastText) and found that pre-trained GloVe embeddings with 200 dimensions work best on our data.

Fold	Accuracy
1	0.733
2	0.717
3	0.765
4	0.794
5	0.834
6	0.826
7	0.858
8	0.810
9	0.818
10	0.829
Avg:	0.798

Table 9: Divide&Merge 10-fold cross-validation

4.5 Discussion

To compare the performance of the D&M model (accuracy 79.8%) against that of the FastText classifier (accuracy 75.6%) we applied a randomized test with 10,000 trials. The resulting p-value of 0.0001 indicates a significant difference. The accuracy of both FastText and D&M is also significantly better than that of the baselines (24.7% for the relative-frequency baseline, 34.8% for the majority baseline and 47.2% for the mixed baseline). Table 10 contains the results for precision, recall and F-score per DA.

Category	FastText			Divide&Merge		
	Prec.	Rec.	F1	Prec.	Rec.	F1
Contact	0.94	0.96	0.95	0.96	0.98	0.97
Inform	0.70	0.77	0.74	0.75	0.78	0.76
Affirmative	0.78	0.80	0.79	0.81	0.82	0.82
Request	0.69	0.68	0.68	0.75	0.76	0.75
Question	0.58	0.54	0.56	0.71	0.61	0.65
Confirm	0.40	0.28	0.33	0.48	0.50	0.49
Disconfirm	0.60	0.51	0.55	0.60	0.55	0.57
Average (w/o Neg.):	0.67	0.65	0.66	0.72	0.71	0.72
Average (with Neg.):	0.59	0.57	0.57	0.63	0.62	0.63

Table 10: FastText and D&M results per DA

We also compared the performance of the D&M model with threads to the same model without thread information. The results are in Table 11. Note that Tables 10 and 11 show average precision, recall and F1 score for two cases: with and without the category *Negative*. *Negative* turned out to be very difficult to classify because of the




Figure 4: Divide&Merge architecture ⁶

following reasons. First, there is a data sparsity problem, because *Negative* has only 14 occurrences in the whole corpus. Second, *Negative* is very similar to *Disconfirm* and in many cases they can be used interchangeably. However, *Negative* was omitted only in the precision and recall calculations showing performance per DA. All accuracy scores presented in this paper take *Negative* into consideration.

Table 11 shows that F1 score increases when the thread information is provided as an additional input to the model. For all DAs except for *Disconfirm* and *Negative* we observe an improvement in terms of precision, recall and F1 score. The poor performance of D&M model on categories *Negative* and *Disconfirm* could be due to the fact that some threads are interconnected and *Negative* is often a response to the previous thread. For instance, in one thread the OP says "I will put snapshots in ..." And in the next thread the TL says "I don't have snapshots" which should be interpreted as *Negative* with respect to the previous statement. However, D&M classifies the utterance as *Inform* because it does not see the connection between two different threads.

Further manual checking of the classification results confirmed that the D&M model could handle DAs which depend on the context better. Table 12 illustrates this: In Thread 1 FastText almost always picked *Inform* as the most likely label, whereas D&M assigned more DAs correctly. In Thread 2 FastText assigned *Contact* for "Yeah, I am driving closer now".

Category	D&M no threads			D&M with threads		
	Prec.	Rec.	F1	Prec.	Rec.	F1
Contact	0.95	0.96	0.95	0.96	0.98	0.97
Inform	0.74	0.73	0.73	0.75	0.78	0.76
Affirm.	0.80	0.76	0.78	0.81	0.82	0.82
Request	0.73	0.74	0.74	0.75	0.76	0.75
Question	0.64	0.60	0.62	0.71	0.61	0.65
Confirm	0.37	0.47	0.41	0.48	0.50	0.49
Disconfirm	0.62	0.59	0.60	0.60	0.55	0.57
Negative	0.25	0.07	0.11	0.00	0.00	0.00
Average (w/o Neg.):	0.69	0.69	0.69	0.72	0.71	0.72
Average (with Neg.):	0.64	0.61	0.62	0.63	0.62	0.63

Table 11: D&M results with and without threads

Speak.	Text	FastText	D&M
Thread 1			
TL	<i>UGV 1 to team leader.</i>	Contact	Contact
OP	<i>I am coming.</i>	Inform	Contact
TL	<i>Can you find out whats standing in all this smoke?</i>	Inform	Question
OP	<i>Yes. I could. You should have a picture of that.</i>	Inform	Confirm
TL	<i>I'll check that.</i>	Inform	Affirm.
Thread 2			
TL	<i>Can you get closer to the blue barrel, so that we can see the label?</i>	Request	Request
OP	<i>Yeah, I am driving closer now.</i>	Contact	Affirm

Table 12: Sample DA classification results by FastText and D&M. Correctly assigned DAs are typeset in bold.

I am driving closer now". Although there were some instances of *Contact* in the training corpus starting with "yeah", *Contact* is not a good candi-

date in this case given that the previous utterance was labeled as *Request*. This shows that thread history has an impact on the output of the D&M model. The D&M model makes better use of the thread history than FastText and seems to offer a better model for structured conversations.

In general, the independence assumption made by FastText impairs the classification performance. However, adding thread history resulted in an accuracy drop from 75.6% to 64.0% (cf. §4.3). This means that it is not only thread information that is important for correct classification but also the way this information is encoded and processed by the classifier. Whereas FastText treats the current utterance and the thread history in a bag-of-words fashion, the D&M model treats them as two independent inputs which are being processed by two different parts of the network and their representations are concatenated only at the final stage.

We also tested several models on the part of the Switchboard Corpus available in DialogBank (Bunt et al., 2016). After pre-processing similar to what we did for our corpus we had 443 utterances. We split them into 333 (75%) for training and 110 (25%) for testing. FastText achieved accuracy 60%. Among the neural models a simple FFNN using the Embedding layer initialized with pre-trained GloVe embeddings with 100 dimensions achieved best accuracy 73.6%. The D&M model could not be applied to the DialogBank-Switchboard data because there are no clearly delimited threads. It would be interesting to test the D&M approach on other corpora with dialogues structured into threads similarly to our corpus.

5 Conclusions

We presented the results of dialogue act classification in robot-assisted disaster response team communication. We experimented with a FastText classifier and various neural models using FFNNs, RNNs and CNNs with different types of embeddings and context information, with and without attention. We found that including the speaker role is beneficial whereas adding the previous sentence as dialogue context leads to lower accuracy. This might be due to the fact that dialogues in our corpus consist of threads and concatenating an utterance with a preceding one from a different thread causes erroneous predictions. We then designed the Divide&Merge model, where we added thread history in a separate layer and concatenated not

texts but their vector representations. This resulted in a significant improvement with average accuracy 79.8%. Using LSTM cells was beneficial for utterance encodings but the thread history was better encoded using the Embedding layer and global average pooling. Pre-trained GloVe embeddings with dimensionality 200 performed best on our data and the results were slightly better with trainable embeddings. This could be due to the fact that in our corpus some words have non-standard interpretations based on the communication protocol (e.g., *”roger that”*), which are learned from the corpus when we use trainable embeddings.

Incorporating thread information significantly improved DA classification. In the future we wish to investigate more the nature and importance of threads in team communication, e.g., whether to model threads implicitly (as we did) or explicitly; how to best segment them; how important is it to represent intertwined threads; is information throughout a thread used for interpretation or is the influence more local at the thread boundary.

In future work we will also apply the models presented here on the German data in the TRADR corpus; test their performance on the outputs of ASR without any editing by human annotators; look for ways to further improve performance, e.g., by enlarging the corpus by adding relevant dialogues from other corpora. We will develop models for the recognition of mission tasks and distinguishing task requests and commitments by the team members from other task mentions. We will then combine dialogue act and task recognition in a single model. We will release the corpus with the ISO dialogue act annotations later this year.

The models we develop are being integrated as part of the speech processing pipeline in a mission-support system that provides process assistance and facilitates the creation of mission documentation (Willms et al., 2019). It will be evaluated in practice with and by first responders.

Acknowledgements

This work is part of the A-DRZ project funded by the German Ministry of Education and Research (BMBF), grant No. I3N14856.⁷ We wish to thank Stefania Racioppa and Natalia Skachkova for their contributions to annotate the TRADR corpus and all our colleagues for valuable discussions.

⁷A-DRZ (Setup of the German Rescue Robotics Center). URL: rettungsrobotik.de

References

- James F. Allen. 1991. [Discourse structure in the TRAINS project](#). In *Speech and Natural Language, Proceedings of a Workshop held at Pacific Grove, California, USA, February 19-22, 1991*.
- Anne Anderson, M. Bader, Ellen Bard, E. Boyle, Gwyneth M. Doherty, Simon Garrod, Stephen Isard, Jacqueline Kowtko, J. McAllister, J. Miller, Cathy Sotillo, Henry Thompson, and R. Weinert. 1991. The HCRC Map Task corpus. *Language and Speech*, 34(4):351–366.
- Pawel Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. 2018. [Multiwoz - A large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 5016–5026.
- Harry Bunt, Jan Alexandersson, Jae-Woong Choe, Alex Chengyu Fang, Kōiti Hasida, Volha Petukhova, Andrei Popescu-Belis, and David R. Traum. 2012. [ISO 24617-2: A semantically-based standard for dialogue annotation](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation, LREC 2012, Istanbul, Turkey, May 23-25, 2012*, pages 430–437.
- Harry Bunt, Volha Petukhova, and Alex Chengyu Fang. 2017. [Revisiting the ISO standard for dialogue act annotation](#). In *Proceedings of the 13th Joint ISO-ACL Workshop on Interoperable Semantic Annotation (ISA-13)*.
- Harry Bunt, Volha Petukhova, Andrei Malchanau, Kars Wijnhoven, and Alex Chengyu Fang. 2016. [The dialogbank](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation LREC 2016, Portorož, Slovenia, May 23-28, 2016*.
- Jennifer Burke, Michael D. Covert, Robin Murphy, and Dawn L. Riddle. Moonlight in Miami: An ethnographic study of human-robot interaction in USAR. *Human-Computer Interaction, special issue on Human-Robot Interaction*, 19:85.
- Jennifer L. Burke and Robin Murphy. From remote tool to shared roles. *IEEE Robotics and Automation Magazine, special issue on New Vistas and Challenges for Teleoperation*, 15:39.
- Jean Carletta, Simone Ashby, Sebastien Bourban, Mike Flynn, Maël Guillemot, Thomas Hain, Jaroslav Kadlec, Vasilis Karaiskos, Wessel Kraaij, Melissa Kronenthal, Guillaume Lathoud, Mike Lincoln, Agnes Lisowska, Iain McCowan, Wilfried Post, Dennis Reidsma, and Pierre Wellner. 2005. [The AMI meeting corpus: A pre-announcement](#). In *Machine Learning for Multimodal Interaction, Second International Workshop, MLMI 2005, Edinburgh, UK, July 11-13, 2005, Revised Selected Papers*, pages 28–39.
- Jean Carletta, Stephen Isard, Gwyneth Doherty-Sneddon, Amy Isard, Jacqueline C. Kowtko, and Anne H. Anderson. 1997. The reliability of a dialogue structure coding scheme. *Computational Linguistics*, 23(1):13–31.
- Liz Carver and Murray Turoff. 2007. Human-computer interaction: The human and computer as a team in emergency management information systems. *Communications of the ACM*, 50(3):33–38.
- Jennifer Casper and Robin Murphy. Human-robot interaction during the robot-assisted urban search and rescue response at the World Trade Center. *IEEE Transactions on Systems, Man and Cybernetics Part B*, 33:367.
- Christophe Cerisara, Pavel Král, and Ladislav Lenc. 2018. [On the effects of using word2vec representations in neural networks for dialogue act recognition](#). *Computer Speech & Language*, 47:175–193.
- Zheqian Chen, Rongqin Yang, Zhou Zhao, Deng Cai, and Xiaofei He. 2018. [Dialogue act recognition via crf-attentive structured network](#). In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR 2018, Ann Arbor, MI, USA, July 08-12, 2018*, pages 225–234.
- Won Seug Choi, Jeong-Mi Cho, and Jungyun Seo. 1999. [Analysis system of speech acts and discourse structures using maximum entropy model](#). In *27th Annual Meeting of the Association for Computational Linguistics, University of Maryland, College Park, Maryland, USA, 20-26 June 1999*.
- Shammur Absar Chowdhury, Evgeny A. Stepanov, and Giuseppe Riccardi. 2016. [Transfer of corpus-specific dialogue act annotation to ISO standard: Is it worth it?](#) In *Proceedings of the Tenth International Conference on Language Resources and Evaluation LREC 2016, Portorož, Slovenia, May 23-28, 2016*.
- Nancy J Cooke, Mustafa Demir, and Nathan McNeese. 2016. Synthetic teammates as team players: Coordination of human and synthetic teammates. Technical report, Cognitive Engineering Research Institute, Mesa (US).
- Mark G. Core and James F. Allen. 1997. Coding dialogues with the DAMSL annotation scheme. In *Working Notes: AAAI Fall Symposium on Communicative Action in Humans and Machines*, pages 28–35, Menlo Park, California. AAAI, American Association for Artificial Intelligence.
- John J. Godfrey, Edward C. Holliman, and Jane McDaniel. 1992. [Switchboard: Telephone speech corpus for research and development](#). In *Proceedings of the 1992 IEEE International Conference on Acoustics, Speech and Signal Processing - Volume 1, ICASSP'92*, pages 517–520, Washington, DC, USA. IEEE Computer Society.

- Gang Ji and Jeff A. Bilmes. 2005. [Dialog act tagging using graphical models](#). In *2005 IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP '05, Philadelphia, Pennsylvania, USA, March 18-23, 2005*, pages 33–36.
- Matthew Johnson, Brandon Shrewsbury, Sylvain Bertrand, Duncan Calvert, Tingfan Wu, Daniel Duran, Douglas Stephen, Nathan Mertins, John Carff, William Rifenburgh, and Jesper Smith. 2017. Team ihmcs lessons learned from the darpa robotics challenge: Finding data in the rubble. *Journal of Field Robotics*, 34(2):241.
- Geert-Jan M. Kruijff, Ivana Kruijff-Korbayová, Shanker Keshavdas, Benoit Larochelle, Miroslav Janíček, Francois Colas, M. Liu, Francois Pomerleau, Roland Siegwart, Mark A. Neerincx, Rosemarijn Looije, Nanja J.J.M Smets, Tina Mioch, Juriaan van Diggelen, Fiora Pirri, Mario Gianni, F. Ferri, M. Menna, Rainer Worst, T. Linder, V. Tretyakov, Hartmut Surmann, Tomáš Svoboda, Michael Reinštein, Karel Zimmermann, Tomáš Petříček, and Václav Hlaváč. 2014. Designing, developing, and deploying systems to support humanrobot teams in disaster response. *Advanced Robotics*, 28(23):1547–1570.
- Ivana Kruijff-Korbayová, Francis Colas, Mario Gianni, Fiora Pirri, Joachim de Greeff, Koen Hindriks, Mark Neerincx, Petter Ögren, Tomáš Svoboda, and Rainer Worst. 2015. [Tradr project: Long-term human-robot teaming for robot assisted disaster response](#). *KI - Künstliche Intelligenz*, 29(2):193–201.
- Harshit Kumar, Arvind Agarwal, Riddhiman Dasgupta, Sachindra Joshi, and Arun Kumar. 2017. [Dialogue act sequence labeling using hierarchical encoder with CRF](#). *CoRR*, abs/1709.04250.
- Yang Liu, Kun Han, Zhao Tan, and Yun Lei. 2017. [Using context information for dialog act classification in DNN framework](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 2170–2178.
- Melanie J Martin and Peter W Foltz. 2004. Automated team discourse annotation and performance prediction using LSA. In *Proc. of HLT-NAACL 2004: Short Papers*, pages 97–100. ACL.
- Silvia Quarteroni and Giuseppe Riccardi. 2010. [Classifying dialog acts in human-human and human-machine spoken conversations](#). In *INTERSPEECH 2010, 11th Annual Conference of the International Speech Communication Association, Makuhari, Chiba, Japan, September 26-30, 2010*, pages 2514–2517.
- Colin Raffel and Daniel P. W. Ellis. 2015. [Feed-forward networks with attention can solve some long-term memory problems](#). *CoRR*, abs/1512.08756.
- Eugénio Ribeiro, Ricardo Ribeiro, and David Martins de Matos. 2015. [The influence of context on dialogue act recognition](#). *CoRR*, abs/1506.00839.
- Hannes Schulz, Jeremie Zumer, Layla El Asri, and Shikhar Sharma. 2017. [A frame tracking model for memory-enhanced dialogue systems](#). In *Proceedings of the 2nd Workshop on Representation Learning for NLP, Rep4NLP@ACL 2017, Vancouver, Canada, August 3, 2017*, pages 219–227.
- Elizabeth Shriberg, Rajdip Dhillon, Sonali Bhagat, Jeremy Ang, and Hannah Carvey. 2004. [The ICSI meeting recorder dialog act \(MRDA\) corpus](#). In *Proceedings of the SIGDIAL 2004 Workshop, The 5th Annual Meeting of the Special Interest Group on Discourse and Dialogue, April 30 - May 1, 2004, Cambridge, Massachusetts, USA*, pages 97–100.
- Amanda Stent. 2000. The monroe corpus. Technical report, Dept. of Computer Science, University of Rochester.
- Andreas Stolcke, Klaus Ries, Noah Coccaro, Elizabeth Shriberg, Rebecca A. Bates, Daniel Jurafsky, Paul Taylor, Rachel Martin, Carol Van Ess-Dykema, and Marie Meteer. 2000. [Dialogue act modeling for automatic tagging and recognition of conversational speech](#). *CoRR*, cs.CL/0006023.
- Zachary O. Toups, William A. Hamilton, and Sultan A. Alharthi. 2016. [Playing at planning: Game design patterns from disaster response practice](#). In *Proceedings of the 2016 Annual Symposium on Computer-Human Interaction in Play, CHI PLAY '16*, pages 362–375.
- Christian Willms, Constantin Houy, Jana-Rebecca Rehse, Peter Fettke, and Ivana Kruijff-Korbayová. 2019. Team communication processing and process analytics for supporting robot-assisted emergency response. In *International Conference on Safety, Security, and Rescue Robotics (SSRR)*.
- Fereshta Yazdani, Gayane Kazhoyan, Asil Kaan Bozcuoğlu, Andrei Haidu, Ferenc Bálint-Benczédi, Daniel Beßler, Mihai Pomarlan, and Michael Beetz. 2018. [Cognition-enabled framework for mixed human-robot rescue teams](#). In *International Conference on Intelligent Robots and Systems (IROS)*, pages 1421–1428. IEEE.

Appendix A

Team Communication Example

- TL *Andreas, Andreas from Markus, come in.*
- OP *yes, Andreas come in.*
< ... >
- OP *yes, for information, I am ready [EHM] shall I go ahead with my search command, or begin?*
- TL *Yes, begin immediately without possible – least possible time delay, to [EHM] have a higher chance for person rescue.*
- OP *yes, understood, I begin with the search.*
< ... >
- TL *Andreas from Markus, come in. [ent = unk.skippable]*
- OP *Yes, Andreas, come in.*
- TL *[ent = unk.skippable] are there already any noteworthy findings? [ent = unk.skippable]*
- OP *Negative. No noteworthy findings. [ent = unk.skippable]*
- TL *Yes, understood. [ent = unk.skippable] Daniel, Daniel from Markus, come in. [ent = unk.skippable] Andreas from Markus, come in.*
< ... >
- OP *Andreas, Markus from Andreas, come in.*
- TL *Andreas, come in.*
- OP *On first floor in the smoke found a barrel, green, labeled as environmentally hazardous material.*
- TL *Yeah, can you [unintelligible] whether anything is leaking?*
- OP *Yeah. It is a 200 liter barrel, whether anything is leaking I cannot currently tell.*
- TL *[EHM] Any thermal emission?*
- OP *No thermal emission.*
- TL *Okay. Priority on continuing person search. Andreas from Markus, priority on continuing person search.*

Multi-Task Learning of System Dialogue Act Selection for Supervised Pretraining of Goal-Oriented Dialogue Policies

Sarah McLeod¹, Ivana Kruijff-Korbayová², Bernd Kiefer²

¹Saarland University, Saarbrücken, Germany

²German Research Center for Artificial Intelligence (DFKI), Saarbrücken, Germany

smcleod@coli.uni-saarland.de

{ivana.kruijff, bernd.kiefer}@dfki.de

Abstract

This paper describes the use of Multi-Task Neural Networks (NNs) for system dialogue act selection. These models leverage the representations learned by the Natural Language Understanding (NLU) unit to enable robust initialization/bootstrapping of dialogue policies from medium sized initial data sets. We evaluate the models on two goal-oriented dialogue corpora in the travel booking domain. Results show the proposed models improve over models trained without knowledge of NLU tasks.

1 Introduction

To be successful, goal-oriented dialogue systems must accurately determine the intent(s) of a user, identify and understand the relevant information they have provided, and based on that information, select the appropriate response at each turn in the conversation. One way to model conversation is as a partially observable Markov decision process (Young et al., 2013). In this framework system response generation is modeled as a stochastic policy, and research into statistically optimizing dialogue policies with Reinforcement Learning (RL) is an active area of research (Gasic and Young, 2014; Lemon and Pietquin, 2007). However, learning optimal dialogue policies with RL can be challenging since large state and action spaces require large amounts of training data to densely sample the space (Lemon and Pietquin, 2007; Wen et al., 2016; Li et al., 2017). Additionally, networks trained with RL learn in a trial-and-error process, guided by a potentially delayed reward function. This exploration process can lead to poor performance in the early training stages, which in turn can lead to a negative user experience (Su et al., 2016).

To address these issues supervised learning has been used for pre-training of dialogue policies (Su

et al., 2016; Henderson et al., 2007; Williams and Zweig, 2016), however the previous approaches only considered one aspect of dialogue during training. Grosz and Sidner (1986) describe discourse structure as a composite of multiple aspects that interact and co-constrain one other. This structure determines the meaning of a discourse and provides a framework for processing dialogue. The question then arises whether it would be beneficial to view dialogue policy training as a multi-task learning (MTL) problem. MTL is an active area of research and has been shown to improve performance on a number Natural Language Processing (NLP) tasks (Ruder, 2017; Zhang and Yang, 2017). In this work we propose a method to use the training signals of related tasks during supervised pre-training of system dialogue act selection as part of dialogue policy initialization. We also experiment with multiple architectures across two data sets and evaluate against two baseline architectures.

Specifically, we use slot-filling and user-intent classification as auxiliary tasks for the primary task of system dialogue act selection. For many corpus trained dialogue systems slot-filling and user-intent classification are trained independently, separate from the dialogue manager. We hypothesize that the features learned when training neural models for these tasks are also informative for the initialization of a robust dialogue policy network. In MTL there can be an added cost of collecting labels for auxiliary tasks, but in the scenario in this paper the labels for user-intent and slot-filling that are needed to develop a complete dialogue system already exist; the framework we propose uses these labels as additional information to initialize the dialogue manager. The next sections describe related work in MTL, including MTL for goal-oriented dialogue systems, the corpora used in our experiments, the architecture of

Corpus	# Slots	# Intents	# Actions	Avg. UL	# Train	# Dev	# Test	Vocab
ATIS	79	17	NA	11.13	4478	500	894	900
Frames	21	45	52	8.05	6131	1532	1916	3249
COMM ATT	34	10	56	1.86	3960	442	781	545
COMM BBN	37	18	48	2.30	3168	351	622	490
COMM CMU	45	32	72	2.58	2793	310	548	580
COMM SRI	37	17	25	2.44	4076	452	800	569

Table 1: The number of slot types, user speech acts, and system dialogue acts for each corpus, as well as the average length of the user input utterances.

the neural models we tested, and the results of the evaluation.

2 Related work

Multi-Task Learning: In MTL the training signals of related tasks are used to learn features that are relevant to multiple tasks, including a primary task of interest. In learning these shared features the model learns a representation that improves generalization on that primary task. [Caruana \(1998\)](#) and [Zhang and Yang \(2017\)](#) describe a number of tasks where the shared representation learned with MTL improves generalization. MTL has also been shown to improve a number NLP tasks ([Toshniwal et al., 2017](#); [Arik et al., 2017](#); [Dong et al.](#); [Zoph and Knight, 2016](#); [Johnson et al., 2016](#)). See [Ruder \(2017\)](#) and [Zhang and Yang \(2017\)](#) for additional examples of MTL for NLP.

Goal-Oriented Dialogue systems: [Wen et al. \(2016\)](#) treat dialogue as a sequence to sequence mapping problem and design a dialogue manager where each component is modularly connected and trainable from data. Previous work also learns state-tracking and other NLU tasks simultaneously. [Hakkani-Tur et al. \(2016\)](#) use a bi-directional LSTM to jointly model slot filling, intent determination, and domain classification for different domains. [Chen et al. \(2016\)](#) use a knowledge-guided structural attention network (K-SAN) to model intent prediction and slot filling simultaneously. Both published results on the ATIS corpus ([Price, 1990](#)). [Padmakumar et al. \(2017\)](#) train a semantic parser and policy network in batches, giving the policy network access to the updated semantic parser after every batch. [Zhao and Eskenazi \(2016\)](#) jointly learn policies for state tracking and dialogue strategies using Deep Recurrent Q-Network (DRQN). [Li et al. \(2017\)](#) use a single RNN with LSTM to jointly learn user intent as well as slot filling. Their dialogue manager is initialized by supervised learning of labels generated by a rule system, then end-to-end train-

ing is continued with RL using a user simulator. Results were published on data from movie-ticket booking domain. We also propose to initialize the dialogue manager with supervised learning, however we use the information from upstream dialogue system tasks during supervised pre-training. We also experiment on two distinct corpora in the travel planning domain across multiple architectures.

3 Data

We evaluated our models on three corpora: the Maluuba Frames ([El Asri et al., 2017](#)), DARPA COMMUNICATOR ([Georgila et al., 2009, 2005](#)) and ATIS ([Price, 1990](#)) data sets. The Frames corpus is a collection of human-human dialogues that captures realistic behaviors in natural conversations. The DARPA COMMUNICATOR corpus is a collection of human-computer interactions from users calling into the COMMUNICATOR travel planning system. We use the version described in ([Georgila et al., 2005](#)), Georgila:COMMUNICATOR, which includes annotations from the original corpus plus additional user-intent and task level annotations automatically added by a system they designed. The complete COMMUNICATOR corpus includes data for all systems evaluated as part of the DARPA program. As in [Henderson et al. \(2007\)](#) we use only the data from the ATT, BBN, CMU and SRI systems. The ATIS corpus is a collection of spontaneous speech and associated annotations, collected in a Wizard-of-Oz setup. The corpus was included in the software released by [Hakkani-Tur et al. \(2016\)](#) and we used it to for a comparison to their work. The number of unique labels for each task as well as the train, dev and test data splits for each corpus are listed in Table 1.

3.1 Preprocessing

We used the common IOB (in-out-begin) format to annotate slot-tags for each token. In this schema,

for each input sequence X tokens t_1, \dots, t_n are assigned a slot label s_1, \dots, s_n and multi-token values are labeled with B (begin) and I (inside) to indicate the extent of the tokens that fill that slot. Tokens that are not relevant to any slot are tagged with O (outside). Some turns in the Frames and COMMUNICATOR corpora were labeled with duplicate user-intent labels and system action labels. One option was to ignore these duplicate labels, however these duplicates occurred frequently enough to be considered informative; therefore when more than one class label exists for a single input utterance, we concatenated all of the labels into a single label. For example, if the system dialogue act was annotated with *negate*, *negate*, and *inform* the labels are concatenated to create a single *negate#negate#inform* label.

4 Experiments

We completed three sets of experiments: two baseline experiments and a final experiment with the multi-task architecture. Each of these experiments included three tasks: slot-filling, framed as sequence prediction, user-intent classification, and system dialogue act selection. In the first baseline experiment the models described in [Hakkani-Tur et al. \(2016\)](#) were extended to new corpora and new tasks using the software released by the authors. In the second baseline experiment we trained single-task models for each of the three tasks individually, on each corpus. Following the methodology suggested in [Caruana \(1998\)](#), these models were tuned for each corpus and architecture. The Maluuba Frames and DARPA COMMUNICATOR Corpora were used in baseline and multi-task experiments; the ATIS corpus does not contain annotations for system dialogue act selection and was therefore only used in the baseline experiments.

4.1 Architectures

Baseline A: [Hakkani-Tur et al. \(2016\)](#) describe a recurrent neural network (RNN) architecture for simultaneous learning of slot-filling, domain classification, and user intent classification. They treat joint learning as a sequence labeling task and use a modification of the encoder-decoder model. To represent the data they use the IOB style annotations for slots and for each utterance U associate the sentence final token with a single label generated by concatenating the associated domain d and

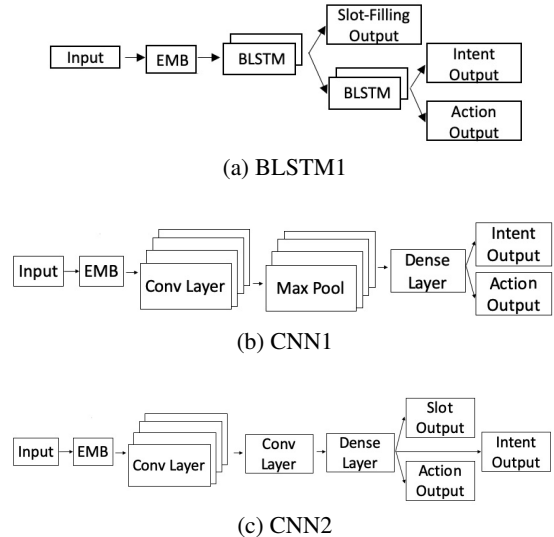


Figure 1: Model architectures for the multi-task experiments.

user-intent u labels. In this framework the input and output utterances become:

$$X = t_1, \dots, t_n, \langle EOS \rangle$$

$$Y = s_1, \dots, s_n, d_u-u_u$$

The model weights are learned by maximizing the conditional likelihood of the training set labels.

In our first baseline experiment we use this architecture to jointly learn user-intent classification, slot-filling, and system dialogue act selection (replacing domain classification) on the Frames and COMMUNICATOR corpora. In our experiments the sentence final token is created by concatenating the user-intent and system dialogue act labels.

Baseline B: Next we trained Bi-directional LSTM (BLSTM) and Convolutional Neural Network (CNN) single-task models to perform each task individually. The BLSTM consisted of an input layer, hidden layer, and output layer. Softmax is used to produce a distribution (p_t) of likely labels at each time-step. The final output is then $argmax(p_t)$. The CNN network consists of two convolutional layers, connected in series, each followed by max pooling layers. A dense layer connects the output of the final convolutional layer to the softmax layer. For slot-filling the models predict a label for each word in the input sequence. For user-intent and system dialogue act selection the models predict a single label for the input utterance. The BLSTM architecture was used to train

individual models for all three tasks. The CNN architecture was used to train individual user-intent and system dialogue act selection models only.

Multi-Task Models: Lastly, we created multi-task models with BLSTMs and CNNs, and a combination of the two. In these architectures each task has a separate output, and all tasks share hidden layers. We implemented three BLSTM versions. BLSTM1 consists of two stacked BLSTMs and the slot-filling output layer is positioned as an auxiliary output at the first BLSTM. For BLSTM1 the loss for slot-filling is backpropagated through the first BLSTM. The loss for user-intent and system dialogue act selection is backpropagated through both BLSTM layers. Figure 1a illustrates this architecture. BLSTM2 uses the BLSTM1 architecture plus a skip connection from the embedding layer to the second BLSTM layer. In BLSTM3 the first BLSTM layer weights are initialized with the weights learned when training slot-filling alone. The intent was to explore the possible benefit of transfer learning from a previously trained model. Experiments on the subsets of the COMMUNICATOR corpus with BLSTM3 include model training where the weights of the first BLSTM layer are initialized with the weights learned on the Frames data (BLSTM3b). Finally, ablation testing was also done to explore the influence of each auxiliary task. The BLSTM1 model was trained on all three tasks simultaneously (BLSTM1a), on slot-filling and the primary task alone (BLSTM1b), and on user-intent classification and the primary task alone (BLSTM1c).

The CNN1 network design was inspired by Yoon (2014) and is illustrated in Figure 1b. This network uses 4 filters of different widths each followed by max pooling over time. Filter widths, the number of feature maps, and the number of nodes in the fully connected layer were chosen based on the suggestions of Zhang and Wallace (2015). Early experiments on the BLSTM networks showed a potential benefit to using user-intent classification alone as an auxiliary task, therefore these experiments used only user-intent classification as the auxiliary task.

We also conducted experiments with networks inspired by Google’s Inception architecture (Szegedy et al., 2014). This is a general purpose architecture where the output from multiple convolutional layers is passed to a single convolutional layer, called a bottle-neck, which constrains

Corpus	best F	Avg F
ATIS	95.48%	94.70%
Frames	74.26%	73.05%
COMMUNICATOR ATT	48.17%	45.98%
COMMUNICATOR BBN	50.34%	48.77%
COMMUNICATOR CMU	53.96%	52.59%
COMMUNICATOR SRI	59.74%	58.55%

Table 2: The best F-measure and average F-measure on slot-filling alone for each corpus using the architecture released by Hakkani-Tur et al. (2016).

the number of features that subsequent layers take as input, keeping the number of parameters low while retaining the expressive power of the network. Our architecture is illustrated in Figure 1c. This network uses 5 convolutional layers of different filter widths followed by a single bottle-neck convolutional layer. The CNN2b network is composed of three CNN2a networks concatenated together.

The final multi-task network is a hybrid CNN + BLSTM architecture. In this network the input is connected to a CNN network of three convolutional layers with different filter widths each followed by max pooling. This is then connected to the BLSTM1 architecture. The goal was to explore the possibility of extracting features with a CNN layer that could then be used by the BLSTM1 network.

4.2 Training

All network development and training was done in Keras (Chollet et al., 2015) and the code will be released with the final version of this paper. We experimented with batch sizes of 15, 25, 50 and 100, hidden layers of 25, 50 and 100 units, and drop-out ratios of 0,0.25, and 0.5 on the fully-connected layers. GloVe (Pennington et al., 2014) word embeddings were used as pre-trained word embeddings. The Adam optimizer was used with a learning rate of 0.001. All weights were initialized with glorot uniform. The BLSTM layers used tanh as the activation function. During training the validation loss was monitored and early stopping was used to prevent over-fitting.

5 Evaluation

Table 2 shows the best and average F-measure for slot-filling alone on each corpus using the architecture released by Hakkani-Tur et al. (2016).

Model	Frames	ATT	BBN	CMU	SRI
<i>Baseline A</i>	34.77%	52.55%	40.59%	36.11%	57.72%
<i>BLSTM Baseline</i>	36.26%	52.56%	43.52%	37.06%	57.36%
BLSTM1a	36.08%	52.52%	43.71%	37.49%	58.12%
BLSTM1b	35.88%	52.31%	42.52%	38.06%	57.32%
BLSTM1c	37.93%*	53.12%	42.96%	36.58%	57.21%
BLSTM2	36.29%	51.74%	43.56%	37.55%	58.61%
BLSTM3	37.49%*	52.13%	42.27%	37.34%	58%
BLSTM3b	NA	52.29%	43.19%	38.01%	58.30%
<i>CNN Baseline</i>	35.37%	46.93%	43.93%	36.44%	57.39%
CNN1	34.59%	51.90%*	43.93%	36.73%	59.09%
CNN2a	32.78%	52.62%*	43.29%	37.59%	56.09%
CNN2b	32.80%	53.28%*	43.74%	38.47%*	56%
CNN+BLSTM	36.82%*	51.38%*	43.73%	38.53%*	58.49%

Table 3: The best F-measure achieved for each multi-task model on the system action classification task. Results in bold indicate an improvement over the associated single-task baseline (BLSTM or CNN baseline). An asterisk indicates a statistically significant improvement over the respective baseline.

Both best and average F-measure were calculated on the held-out test set, where the average was calculated over 10 different weight initializations. [Hakkani-Tur et al. \(2016\)](#) experimented with multiple LSTM and BLSTM models, but noted that comparable results were achieved on each and therefore only report results on the BLSTM models. We do the same and only report on experiments with their BLSTM architecture. The results on the ATIS corpus are the metrics reported by the authors (and confirmed by us).

For each corpus many of the multi-task models achieved a higher metric score than the Baseline B models on the test data, however significance testing showed not all of these improvements were statistically significant. Significance testing was done with randomized approximation ([Yeh, 2000](#)). Table 3 lists the best F-measure values for each model for the primary task of system action selection.

The majority of the multi-task models, as well as the Baseline B models on the Frames, BBN, and SRI corpora, achieved a higher F-measure than the Baseline A models. (We did not test for statistical significance between the MTL models and the Baseline A). The multi-task CNN models showed statistically significant improvement on three data sets and were faster to train than the BLSTM models, even when larger. Half of the BLSTM models achieved significant improvement on the Frames corpus, but improvement was more sporadic on the COMMUNICATOR corpus. In the Frames corpus most input utterances are much longer since the user provides significant context at each turn. In the COMMUNICATOR corpus after the initial request most user utterances are lim-

ited to one or two word responses to questions presented by the system. This creates a dialogue that looks more like a system initiative dialogue, as compared to the more unconstrained Frames corpus. The CNN+BLSTM network improved performance on three data sets and is the largest of the proposed models.

6 Conclusion

We present multi-task BLSTM and CNN models that use slot-filling and user-intent classification as auxiliary tasks for the primary task of system dialogue act selection as part of dialogue policy initialization. The models bootstrap dialogue policy optimization without the need for hand-written rules, as done, e.g., in ([Li et al., 2017](#)). We also empirically evaluate multiple RNN and CNN architectures on multiple data sets against two baselines architectures. Our MTL models improve over the performance achieved on single task baseline models (Baseline B) as well as the jointly trained BLSTM model released by [Hakkani-Tur et al. \(2016\)](#).

A dialogue manager that is initialized from corpus data is not flexible enough for new user interactions, therefore additional training is necessary. Future work will include deploying our MTL models as part of a complete dialogue system and continued training with RL. This will allow us to explore the performance of MTL models experimentally on end-to-end systems. Additionally, future work will incorporate additional dialogue context into system dialogue act selection, and model the scenario where more than one system dialogue act may be valid at a given point in the dialogue.

References

- Sercan O. Arik, Mike Chrzanowski, Adam Coates, Gregory Diamos, Andrew Gibiansky, Yongguo Kang, Xian Li, John Miller, Andrew Ng, Jonathan Raiman, Shubho Sengupta, and Mohammad Shoeybi. 2017. Deep voice: Real-time neural text-to-speech. *ICML 2017*.
- Rich Caruana. 1998. Multitask learning. *Autonomous Agents and Multi-Agent Systems*, 27(1):95133.
- Yun-Nung (Vivian) Chen, Dilek Hakkani-Tur, Gokhan Tur, Asli Celikyilmaz, and Jianfeng and Gao. 2016. [Syntax or semantics? knowledge-guided joint semantic frame parsing](#). *IEEE Workshop on Spoken Language Technology (SLT 2016)*.
- François Chollet et al. 2015. [Keras](#). GitHub.
- Daxiang Dong, Hua Wu, Wei He, Dianhai Yu, and Haifeng Wang. Multi-task learning for multiple language translation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, pages 1723–1732.
- Layla El Asri, Hannes Schulz, Shikhar Sharma, Jeremie Zumer, Justin Harris, Emery Fine, Rahul Mehrotra, and Kaheer Suleman. 2017. [Frames: a corpus for adding memory to goal-oriented dialogue systems](#). In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 207–219, Saarbrücken, Germany. Association for Computational Linguistics.
- Milica Gasic and Steve Young. 2014. Gaussian processes for pomdp-based dialogue manager optimisation. *IEEE Transactions on Audio, Speech, and Language Processing*.
- Kallirroi Georgila, Oliver Lemon, and James Henderson. 2005. Annotation of communicator dialogue data for learning dialogue strategies and user simulations. *Proceedings of the 9th Workshop on the Semantics and Pragmatics of Dialogue (SEMDIAL:DIALOR)*, pages 61–68.
- Kallirroi Georgila, Oliver Lemon, James Henderson, and Johanna D. Moore. 2009. Automatic annotation of context and speech acts for dialogue corpora. *Natural Language Engineering*, 15(3):315–353.
- Barbara J. Grosz and Candace L. Sidner. 1986. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12:175–204.
- Dilek Hakkani-Tur, Gokhan Tur, Asli Celikyilmaz, Yun-Nung Chen, Jianfeng Gao, Li Deng, and Ye-Yi Wang. 2016. Multi-domain joint semantic frame parsing using bi-directional rnn-lstm. In *Proceedings of Interspeech.*, pages 715–719.
- James Henderson, Oliver Lemon, and Kallirroi Georgila. 2007. Hybrid reinforcement/supervised learning of dialogue policies from fixed data sets. *ACL*.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Vigas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Googles multilingual neural machine translation system: Enabling zero-shot translation](#). arXiv:1611.0455.
- Oliver Lemon and Olivier Pietquin. 2007. Machine learning for spoken dialogue systems. *INTER-SPEECH*.
- Xiujun Li, Yun-Nung Chen, Lihong Li, Jianfeng Gao, and Asli Celikyilmaz. 2017. End-to-end task-completion neural dialogue systems. In *Proceedings of the The 8th International Joint Conference on Natural Language Processing*, pages 733–743.
- Aishwarya Padmakumar, Jesse Thomason, and Raymond J. Mooney. 2017. Integrated learning of dialog strategies and semantic parsing. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- P.J. Price. 1990. Evaluation of spoken language systems: The atis domain. In *Proceedings of the DARPA Workshop on Speech and Natural Language*.
- Sebastian Ruder. 2017. [Multi-task learning objectives for natural language processing](#).
- Pei-Hao Su, Milica Gasic, Nikola Mrksic, Lina Rojas-Barahona, Stefan Ultes, David Vandyke, Tsung-Hsien Wen, and Steve Young. 2016. [Continuously learning neural dialogue management](#). arXiv:1606.02689.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2014. Going deeper with convolutions. *CoRR*, pages 1746–1751. Abs/1409.4842.
- Shubham Toshniwal, Hao Tang, Liang Lu, and Karen Livescu. 2017. [Multitask learning with low-level auxiliary tasks for encoder-decoder based speech recognition](#). arXiv:1704.01631.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrksic, Milica Gasic, Lina M. Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2016. A network-based end-to-end trainable task-oriented dialogue system. In *Proceedings of the 15th Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.

- Jason D. Williams and Geoffrey Zweig. 2016. [End-to-end lstm-based dialog control optimized with supervised and reinforcement learning](#). arXiv preprint arXiv:1606.01269.
- Alexander Yeh. 2000. More accurate tests for the statistical significance of result differences. In *Proceedings of the 18th conference on Computational linguistics, COLING '00*.
- Kim Yoon. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751.
- S. Young, M. Gasic, B. Thomson, and J.D. Williams. 2013. Pomdp-based statistical spoken dialog systems: A review. *proceedings of the ieee*. 101(5):1160–1179.
- Ye Zhang and Byron C. Wallace. 2015. [A sensitivity analysis of \(and practitioners guide to\) convolutional neural networks for sentence classification](#). arXiv:1510.03820.
- Yu Zhang and Qiang Yang. 2017. [A survey on multi-task learning](#). arXiv:1707.08114.
- Tiancheng Zhao and Maxine Eskenazi. 2016. Towards end-to-end learning for dialog state tracking and management using deep reinforcement learning. *the Annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL)*.
- Barret Zoph and Kevin Knight. 2016. Multi-source neural translation. *NAACL*, pages 30–34.

B. Rex: A Dialogue Agent for Book Recommendations

Mitchell Abrams¹, Luke Gessler¹, Matthew Marge²

¹Department of Linguistics
Georgetown University
{mja284,lg876}@georgetown.edu

²Information Sciences Division
U.S. Army Research Laboratory
matthew.r.marge.civ@mail.mil

Abstract

Information overload can be challenging for children searching for books among a multitude of titles, authors, and genres. We present B. Rex, a dialogue agent for book recommendations. B. Rex aims to exploit the cognitive ease of natural dialogue and the excitement of a whimsical persona in order to engage users who might not enjoy using more common interfaces for finding new books. B. Rex succeeds in making book recommendations with good quality based on only information revealed by the user in the dialogue.

1 Introduction

There are many ways to discover a book. Goodreads and Amazon, two popular websites, have user interfaces that are packed full of book recommendations and reviews. Although this mode of presentation serves many users well, younger users like children might benefit more from an interface that is less dense, and that attempts to extrinsically motivate users to engage with the book recommendation task. To investigate these issues, we created a book recommendation dialogue agent, B. Rex.

There are two aspects of B. Rex’s design that make it more suitable for younger users. First, because B. Rex is a dialogue agent, there is less opportunity for users to experience the decision fatigue that can come with a traditional book recommendation interface. Second, B. Rex motivates users to engage with the task by using the persona of an intelligent dinosaur who spends his time reading novels in a secret lab. These two features—the use of recommendation in a text medium, and a whimsical persona—are intended to ultimately make younger users more successful at succeeding at the book recommendation task.

2 Related Work

Book recommendation is comparable to other traditional recommendation domains which have already been explored in the dialogue systems literature, such as hotel or restaurant recommendations, e.g. Ultes et al. (2017).

Within information retrieval, Bogers and Koolen (2018) approach book recommendation with Narrative-Driven Recommendation (NDR), using user narratives from the online book discussion forums as natural language book recommendation requests. Kang et al. (2017) studies user queries for movie recommendations, highlighting not only the range of linguistic variation in such a recommendation task, but also how Wit.ai¹ can effectively tackle NLU in a recommendation task: their slots were retrieved from user input reliably. (B. Rex also uses Wit.ai.)

We are aware of two dialogue-based book recommendation systems. Pan Macmillan Publishing has developed a book recommendation chatbot for Facebook Messenger², where users are presented with different questions to narrow down a set of book recommendations. In this system-driven interaction, users are presented with a set of fixed categories to choose from, leaving little room for self-expression. Author Bot, created by BAM Mobile³, offers a personalized book exploration experience by acting like an author or a character in a book that can discuss the plot, characters, and backstory. While this system does not recommend books, it shares a similar persona element implemented in B. Rex.

B. Rex utilizes the best results from these systems. As far as book recommendation dialogue

¹<https://wit.ai>

²<https://www.digiteum.com/portfolio/panmacmillan-chatbot/>

³<http://www.fastbot.io/author-bot>

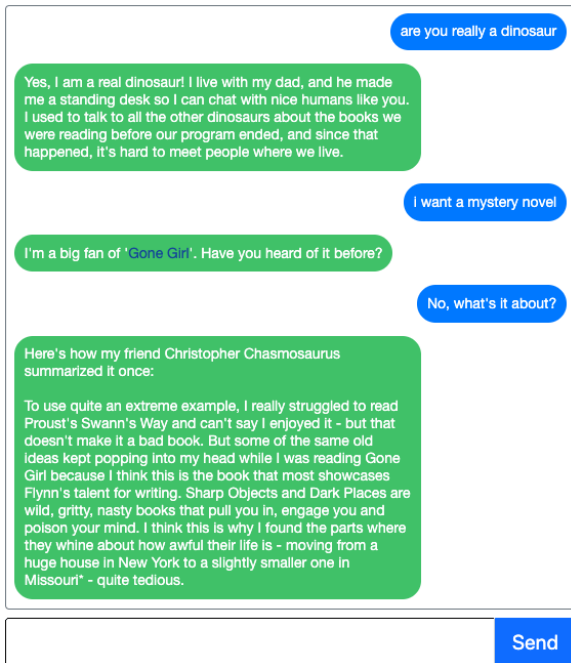


Figure 1: A sample dialogue with B. Rex.

agents go, B. Rex is novel in that it does not constrain user input and also completely allows user-initiated dialogue flow. This is important because systems that do not allow the user any initiative may be efficient in bringing the task to a conclusion, but they are not as engaging or enjoyable for the user as they could be.

3 System Overview

B. Rex⁴ was implemented in Python. Users interact with B. Rex through a web browser. Natural language understanding was handled using Wit.ai, and we relied on Goodreads⁵ for information about books. Handwritten Python string templates were used for natural language generation.

Figure 2 provides an overview of the B. Rex system architecture. Freetext user input was processed with Wit.ai. Based on the value of the `intent` slot returned by Wit.ai, the dialogue manager selects an *intent handler*, a module that is specifically written to handle that intent. For instance, the `greet` intent handler is selected when the value of `intent` is `greet`, i.e. when Wit.ai detects that the user is saying hello to B. Rex. Then, output is planned, generated, and presented to the user.

⁴Source code for B. Rex is accessible at <https://github.com/georgetown-dialogue-systems-2018/brex>, demo at <https://youtu.be/3Z1fBu5PMzc>

⁵<https://goodreads.com>

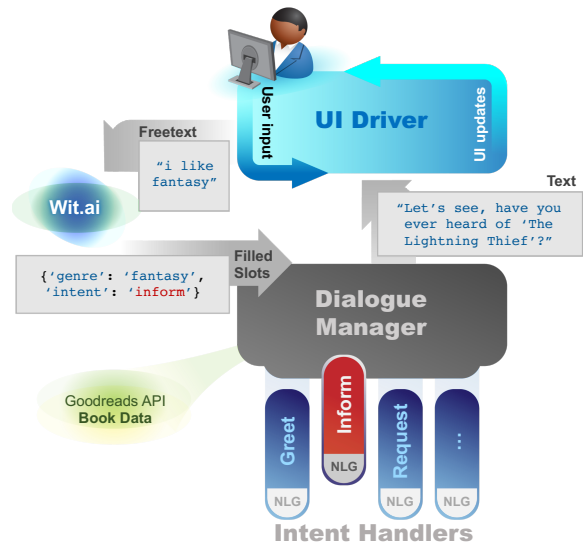


Figure 2: A high-level architectural diagram of B. Rex.

3.1 Wit.ai

Wit.ai is a platform that provides a pre-trained, general-purpose natural language understanding (NLU) system and lets developers tailor it to their domain. Because the system is pre-trained, only a very small amount (on the order of hundreds of labeled inputs) of training data is needed to get a domain-specific model. Further, Wit.ai is able to accommodate slots that are unbounded, which is a necessity in the book recommendation domain, since exhaustively listing all authors or books is not feasible⁶. Beyond slot-filling, Wit.ai also supports classification of entire utterances, which B. Rex uses to determine what a user's intent is for a given message. This feature was useful, since intent is harder to capture with slots alone.

3.2 Intent Handlers

B. Rex has seven intent handlers to respond to user intents. The dialogue manager selects the intent handler that corresponds to the value of `intent`. Once the dialogue manager has selected the intent handler, it hands execution off to it.

The intent handler then plans and generates a text response using this information as well as data retrieved from Goodreads on an as-needed basis. B. Rex uses the Goodreads API⁷ to retrieve information about authors, genres, books, and reviews on Goodreads. To get information about a book,

⁶This might be more feasible with Goodreads' database, but it is not publicly available except through an HTTP API that is limited in its capacity for detailed queries.

⁷<https://pypi.org/project/Goodreads/>

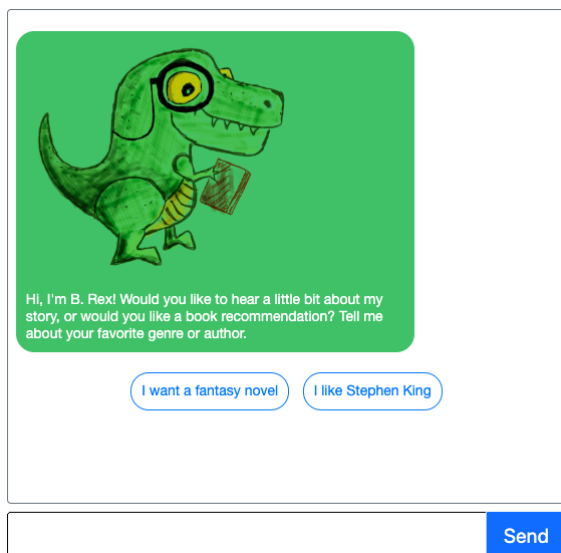


Figure 3: The first message users see when chatting with B. Rex. The user may either enter freetext or select a suggested input.

user reviews are put through a text summarizer⁸ to ensure they are no larger than a screenful.

After all data needed for building a response has been retrieved, the NLG component for each intent handler uses string templates to generate the system’s text. Different templates are written for each intent handler, and around 100 templates were used in total. The text response is passed through the manager to the user interface.

3.3 Extrinsic Motivation: B. Rex’s Persona

A major motivation for the present work is the supposition that our target demographic, younger users, would be engaged by extrinsic motivators, i.e., reasons to engage with the system that don’t have to do with the task itself. We provide this extrinsic motivation at the level of the interface with B. Rex. This primarily comes in the form of a whimsical persona: B. Rex, by hinting at fantastic and interesting bits about his life, drives users to ask questions, building engagement with the system and getting them closer to completing the task.

Before preparing NLG template strings, we created a brief biosketch of B. Rex’s life, personality, and preferences. Template strings referred to the biosketch to showcase B. Rex’s persona, and so that our exposition of his persona would be internally consistent and detailed enough to be lucid and believable. B. Rex sprinkles in bits about his

⁸:sumy <https://pypi.org/project/sumy/>

life when he is chatting with the user about the task, and he is also capable of talking about his favorite books and his life. We followed the findings of Nasihati Gilani et al. (2016) in having B. Rex respond as if he were *really* a dinosaur behind a keyboard, instead of a virtual dinosaur created only for the purposes of this system. We maintained a whimsical tone in various ways, which included having a randomized stock of alliterative, dinosaur-themed book reviewer names (e.g., “Roger Rajasaurus”) and featuring a cartoonish sketch of B. Rex himself.

4 Evaluation

Similar to Griol and Callejas (2013), we gave a survey ($n = 8$) to discover strengths and weaknesses of our system. The results are given in Table 1. These results show that B. Rex was usually successful in recommending a book to users. In practice, some users had difficulties getting a book recommendation from B. Rex, but the main difficulties pertained to Wit.ai not recognizing contextless slots in user input, or certain genres or authors (a database limitation).

As for the quality of B. Rex’s recommendations, according to survey responses, it was for most users just slightly worse than a recommendation from a friend. This level of quality seemed somewhat surprising given B. Rex’s disadvantages compared to a book recommendation system like Amazon or Goodreads, since B. Rex only knows what the user has said, while Amazon and Goodreads have a better model of users that has been constructed from much richer data sources.

In summary, results suggest that the majority of user dissatisfaction had to do with poor understanding and relatively poor recommendation quality. These are both problems that could be easily solved by a commercial system with more training data and more user data, respectively⁹. As for successes, many of the users expressed their amusement with the B. Rex persona, and specifically mentioned their satisfaction with the alliterative dinosaur-reviewer names.

5 Demonstration Outline

Participants will engage B. Rex through an online interface on a laptop or on their own mobile de-

⁹This comes with the caveat that identifying books and authors in isolation may remain somewhat difficult, as discussed in the introduction.

Prompt	Mean	SD
1. How appealing is this book to you, compared to a book a close friend might have recommended to you? (1: least appealing, 7: most appealing)	4.8	1.12
2. B. Rex always understood what I was telling him. (1: B. Rex never understood, 7: B. Rex always understood)	3.5	1.65
3. I often felt unsure about what I could say to B. Rex. (1: I always felt unsure, 7: I never felt unsure)	4.3	1.21
4. B. Rex’s book recommendations were as interesting to me as books that people who know my taste have recommended. (1: nowhere near as interesting, 7: just as interesting)	4.1	1.61
5. Overall, how satisfied were you with B. Rex as a way of finding new books to read? (1: not at all satisfied, 7: very satisfied)	4.5	1.32

Table 1: Results of a survey given to B. Rex users. For these questions, users provided Likert scale ratings from 1 to 7 indicating their agreement with the statement. Survey respondents were all adult native speakers of English, n=8.

vice. Participants will be introduced to B. Rex and be invited to input any requests or questions to begin the book recommendation task. The demonstration will highlight B. Rex’s ability to handle different user questions, the personality of the system, and ability to collaborate with the user in making an efficient and satisfactory recommendation. A real-time display will visualize the dialogue for observers.

6 Conclusion and Future Work

B. Rex demonstrates the utility of natural language interfaces and fictional dialogue agent personas to make book recommendations more engaging for users who are less well served by the prevailing interfaces. B. Rex succeeded in recommending books with good quality to users using no information about them other than their messages. We expect that our approach should generalize to other tasks beyond book recommendation, wherever users find existing interfaces overwhelming or unengaging.

There are a few immediate questions that would need to be addressed by extensions to this work. First, there are many other ways users want to discover books. Users want to be able to find books that are similar to a certain book, that are by an author that is similar to a certain author, or that were published within a certain year range. Second, an ideal book recommendation dialogue system must be able to answer high-level questions about a book. Users want to ask interpretive questions about books, like “does it have a happy ending?” or “does it pass the Bechdel test?”.

For the former, a database with richer data and

more sophisticated querying strategies would do much to solve these problems. The latter problem is more difficult to solve. A fruitful way to tackle these questions might be to aggregate user reviews and use methods from information retrieval and question answering systems to build a response.

References

- Toine Bogers and Marijn Koolen. 2018. “I’m looking for something like...”: Combining Narratives and Example Items for Narrative-driven Book Recommendation. In *KARS18: Proceedings of the First Knowledge-aware and Conversational Recommender Systems Workshop. CEUR-WS*.
- David Griol and Zoraida Callejas. 2013. An Architecture to Develop Multimodal Educative Applications with Chatbots. *International Journal of Advanced Robotic Systems*, 10(3):175.
- Jie Kang, Kyle Condiff, Shuo Chang, Joseph A Konstan, Loren Terveen, and F Maxwell Harper. 2017. Understanding How People Use Natural Language to Ask for Recommendations. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*, pages 229–237. ACM.
- Setareh Nasihati Gilani, Kraig Sheetz, Gale Lucas, and David Traum. 2016. *What Kind of Stories Should a Virtual Human Swap?* In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 1437–1438, Singapore. International Foundation for Autonomous Agents and Multiagent Systems.
- Stefan Ultes, Lina M Rojas Barahona, Pei-Hao Su, David Vandyke, Dongho Kim, Inigo Casanueva, Paweł Budzianowski, Nikola Mrkšić, Tsung-Hsien Wen, Milica Gasic, et al. 2017. Pydial: A Multi-domain Statistical Dialogue System Toolkit. *Proceedings of ACL 2017, System Demonstrations*, pages 73–78.

SpaceRefNet: a neural approach to spatial reference resolution in a real city environment

Dmytro Kalpakchi

Division of Speech, Music and Hearing
KTH Royal Institute of Technology
Stockholm, Sweden
dmytroka@kth.se

Johan Boye

Division of Speech, Music and Hearing
KTH Royal Institute of Technology
Stockholm, Sweden
jboye@kth.se

Abstract

Adding interactive capabilities to pedestrian wayfinding systems in the form of spoken dialogue will make them more natural to humans. Such an interactive wayfinding system needs to continuously understand and interpret pedestrian's utterances referring to the spatial context. Achieving this requires the system to identify exophoric referring expressions in the utterances, and link these expressions to the geographic entities in the vicinity. This *exophoric spatial reference resolution* problem is difficult, as there are often several dozens of candidate referents. We present a neural network-based approach for identifying pedestrian's references (using a network called *RefNet*) and resolving them to appropriate geographic objects (using a network called *SpaceRefNet*). Both methods show promising results beating the respective baselines and earlier reported results in the literature.

1 Introduction

Remember yourself being lost in a completely unfamiliar city without knowing the local language or acquaintances that can help? Being close to desperate, you ask a passerby for a help and get an answer similar to the following:

Just go forward until you see a McDonald's on the corner. There you turn right and keep straight until the old Gothic style church. A tall glass building near it is exactly what you need.

Such wayfinding instruction is a typical example of how humans guide each other in a city, relying mostly on *landmarks* in the vicinity (Cornell and Greidanus, 2006; Goodman et al., 2004; May et al., 2003; Denis, 1997; Lynch, 1960).

On the contrary, a current generation of navigation systems aiding pedestrian wayfinding generally makes use of quantitative information based

on GPS signals, e.g. distances, cardinal directions and street names. The same instruction rephrased by such system would sound as follows:

Head north on West Avenue. Turn right at the corner. Continue 150 meters straight until East Avenue 29. You've reached your destination.

Such instructions are presented to a pedestrian as a sequence on a screen (possibly voiced as well) supplemented by a map with a moving marker indicating pedestrian's position.

The approach presented above, referred to as *turn-by-turn navigation*, does not resemble a human wayfinding process and thus can be perceived as unnatural and more complicated than it should. In our opinion, making pedestrian's experience more natural should be based on the following two observations.

First, a wayfinding is an inherently interactive process, e.g. we need to know if a person is lost, if the instruction is not clear enough, etc. Human guide guarantees such interactivity, since wayfinding happens in a dialogue, hence a wayfinding system should interact with a pedestrian by means of a spoken dialogue.

Second, humans have difficulties understanding instructions based on quantitative characteristics of a spatial environment (such as distance or angles) (Ross et al., 2004), (Moar and Bower, 1983). Such instructions make humans less confident in their ability to reach the goal correctly. Hence, they tend to rely more on qualitative ones, such as salient geographical objects (*landmarks*), by simply referring to them (Denis, 1997). Such approach can be called *landmark-by-landmark navigation*. Furthermore, landmarks can be used not only when giving route descriptions, i.e. serving as a guide, but also when being guided. For instance, when giving a reassuring confirmation to

the guide, such as “Yes, I can see a tall glass building that you’ve mentioned before”, or describing the proximal surroundings when got lost (“I believe I’m lost, but I see a pizzeria to my right”).

A prerequisite for providing such interaction capabilities is being able to identify the landmarks referred to by phrases as “a tall glass building” or “a pizzeria to my right”. Such kind of phrases is called *referring expressions* (RE) and the landmarks these phrases refer to are called *referents*. A task of matching a referring expression with its referent(s) is called *reference resolution* (RR). Guiding humans in a real city environment requires resolving exophoric spatial references, i.e. those referring to spatial objects outside of the discourse. The focus of this paper is on designing the method for solving this task.

The main contribution of this paper is a new method for resolving exophoric spatial REs, consisting of two substeps:

- a method for identifying exophoric spatial REs in spoken utterances;
- a method for resolving exophoric spatial REs to the appropriate referents, represented as 0, 1, or more geographic entities.

2 Background

Pedestrian wayfinding is an interactive, problem-solving process by which people use environmental information to locate themselves and navigate from place to place (Vandenberg et al., 2016). Despite the ubiquity of wayfinding for pedestrians, the navigation systems aiding the process, usually mobile applications, generally use methods offering *a turn-by-turn navigation*, described in the previous section. Such approach limits possibilities for interaction with the system along the route and forces the user to pay constant attention to the map on the screen. Such design can also lead to an increasing spatial anxiety (an anxious feeling when navigating in unfamiliar environments), which was shown by several studies (Hund and Minarik, 2006; Lawton and Kallai, 2002) to negatively influence pedestrian’s wayfinding performance.

In this paper we suggest to remove pedestrian’s dependency on the digital maps by interacting with a pedestrian by means of a spoken dialogue offering *a landmark-by-landmark navigation*. In fact, a number of studies (Cornell and Greidanus,

2006; Goodman et al., 2004; May et al., 2003; Denis, 1997; Lynch, 1960) have confirmed that humans reason about a spatial environment in qualitative terms, mostly relying on landmarks. As stated in (May et al., 2003), pedestrians were observed to use distances and street names much less frequently than landmarks when describing a city environment. Such approach have been observed to be more efficient for older people, who tend to find a way quicker when using a landmark-based navigation aid (Goodman et al., 2004). Pedestrians with cognitive impairment have been observed to rely on landmarks during navigation as well (Sheehan et al., 2006).

As previously stated, a landmark-based navigation requires being able to resolve exophoric spatial references. Exophoric reference resolution is not a new task in itself, but it has primarily been explored in unrealistic environments containing distinct objects that can be described by a relatively small number of visual features, e.g. recognizing one of 36 Pentomino puzzle pieces in (Kennington and Schlangen, 2015), one of 7 Tangram puzzles in (Funakoshi et al., 2012) or an object in a 3D treasure-hunt game in (Engonopoulos et al., 2013). Only recently the research started to focus on resolving references to objects in real environments. (Schlangen et al., 2015) try to identify objects in the images taken from different locations around the world. (Götze and Boye, 2017) deal with reference resolution in a complex city environment. (Chen et al., 2019) present a TOUCH-DOWN dataset, where the agents navigate in a real-life visual urban environment trying to find a hidden object based on a number of cues formulated in a natural language. The presented task is then called *spatial description resolution*, i.e. given a set on instructions find the referred place, whereas *reference resolution* aims at resolving *all* references mentioned in the given instructions as well.

A number of research papers on exophoric reference resolution (eRR) decompose the problem into three subtasks: identifying referring expressions (RE), constructing a search space of candidate referents and resolving the found references. Hence, the descriptions of the existing eRR methods are decomposed in the same way.

As mentioned above, most of the studies on eRR have been conducted in an unrealistically small toy domain, hence REs can be identi-

fied manually, as in (Engonopoulos et al., 2013), (Funakoshi et al., 2012) or (Kennington and Schlangen, 2015). (Schlangen et al., 2015) and (Götze and Boye, 2017) addressed RR in a realistic domain, but all REs were manually annotated as well. (Schutte et al., 2010) worked on resolving REs in simple manipulation instructions, e.g. “hit that red button”, and identified REs using a set of simple regular expressions. (Prasov and Chai, 2010) used syntactic parsing on a word confusion network, constructed out of n-best list of alternative speech recognition hypotheses. All non-pronominal NPs were then detected and said to be a set of exophoric REs.

In most research studies, the search space of candidate referents is the same for all utterances and consists of a limited number of objects, e.g. (Kennington and Schlangen, 2015), (Funakoshi et al., 2012), (Engonopoulos et al., 2013), (Matuszek et al., 2014). In these studies all candidate referents (candidate set) have a limited number of distinct properties (color, shape, size, etc) and hence each object in the search space is either represented as a combination of such properties or simply as a numeric identifier (as the search spaces are very small). (Schlangen et al., 2015) worked with resolving references to a much more diverse real-life objects in images containing object segmentations. The referred objects come from over 80 different categories and only around 2% of the objects comprise geographical entities, e.g. benches, traffic lights, fire hydrants, etc., that are of interest in the present article. The candidate set for every referring expression was set to contain all object segmentations of the given image and every candidate is encoded using a deep convolutional neural network augmented with a number of extra features. Similarly, (Götze and Boye, 2017) have dealt with a constantly changing candidate set of diverse *geographical* objects in a pedestrian’s vicinity. Each geographical object was then represented by a pedestrian’s position and a number of properties inferred from OpenStreetMap (OSM) (Haklay and Weber, 2008).

In most of the studies, eRR itself is solved by taking the stochastic approach by training a generative probabilistic model to estimate the distribution over a set of candidate objects and then find the most probable intended referent as:

$$O^* = \arg \max_O P(O|U, S), \quad (1)$$

where U is a representation of an utterance constituting RE, S is the search space of possible referents, O is an object in the search space, O^* is the predicted referent. Such a stochastic approach is pursued, for instance, in (Kennington and Schlangen, 2015), (Engonopoulos et al., 2013), (Matuszek et al., 2014), (Funakoshi et al., 2012), (Schlangen et al., 2015), (Götze and Boye, 2017).

3 Approach

Also in this paper, spatial reference resolution is seen as a three-stage problem. First, referring expressions should be identified in the utterances and encoded into a numerical representation. We refer to this stage as *spatial referring expressions identification (sREI)*. This is achieved by a neural network, referred to as *RefNet*. Then the candidate set of referents should be constructed (described further in Sect. 4). Finally, the found referring expressions should be resolved to the appropriate referents, which we call a *spatial reference resolution (sRR) stage*. This adds a spatial dimension to the first task, hence a method name *SpaceRefNet* (also a neural network).

sREI (Sect. 3.1) is seen as a classification problem, where each word is to be assigned one of the three labels, **B-REF** (beginning of RE), **I-REF** (inside of RE), **O** (outside of RE), inspired by the BIO labeling strategy for named entity recognition (NER).

sRR (Sect. 3.2) is seen as a set of binary classification problems, each assigning a pair of an RE and a candidate object to either the positive class, if the candidate is predicted as a referent for the RE, or to the negative class, otherwise. Both stages use the same dataset, described further in Sect. 4, pre-processed in different ways.

3.1 Referring expression identification

Let us now describe the way *RefNet* operates (see Fig. 1). We start by padding (with a special word <pad>) or trimming every utterance to some fixed sentence length L_s . Each utterance is fed into RefNet word by word, as a part of a training batch. Each word is encoded using pre-trained D_w -dimensional distributional word embeddings (we are using GloVe (Pennington et al., 2014)). Additionally, each word is split into characters, mapped to the pre-trained \tilde{D}_c -dimensional character embeddings, trained on the SpaceRef corpus

using the Random Indexing technique (Kanerva et al., 2000) for a character level. These character embeddings are then fed into a bidirectional recurrent neural network (BiRNN) with gated recurrent units (GRUs), having rectifier activation functions (ReLUs) and H_c -dimensional hidden states. This BiRNN produces ($D_c = 2 \times H_c$)-dimensional *character-level word embeddings* by concatenating the hidden states of forward and backward GRUs.

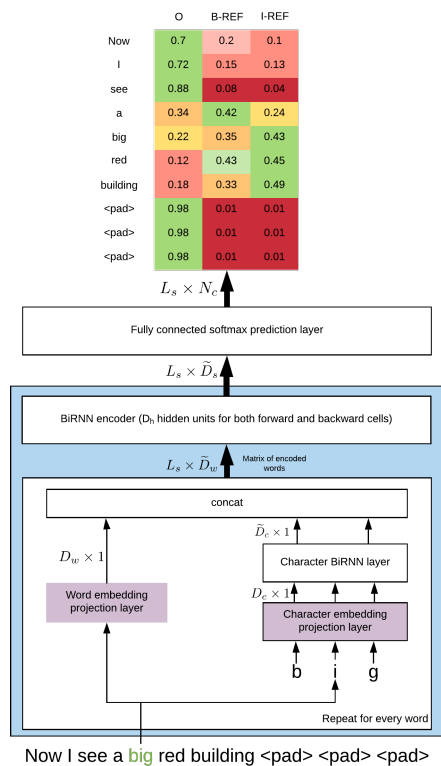


Figure 1: RefNet architecture diagram. The purple blocks specify the pre-trained layers; thick arrows emphasize that 2D tensors of dimensionality specified to the left of arrows are passed; the blue block denotes RefNet encoder. (*Best viewed in color*)

The motivation behind taking character-level embeddings into account is that some words in an RE will inevitably lack word vectors. In such cases, the corresponding word embeddings are assigned to be zero vectors, leaving character-level embeddings as the only source of information. This amendment should be particularly helpful in at least the following two cases:

- if an RE is a proper name of a geographical object, pronounced in the language, different from the dominant language of the utterance, e.g. “Bahnstraße”, “Östvägen”;
- if an RE is a composite name with one of the

constituents being recognized as a valid RE, e.g. “supermegamarket”.

The final *word encoding* is then a concatenation of the word embedding and the character-level word embedding, resulting in a $(D_w + D_c)$ -dimensional vector. These word encodings are then collected into a sentence representation, which is a $L_s \times D_w$ matrix. This sentence representation is fed row-wise as a sequence into another BiRNN (with forward and backward GRUs with ReLUs having H_s hidden units).

In order to incorporate the contextual information, we want to represent a sentence as a matrix, the i^{th} row of which contains the information about the sub-sentence up until, and including, the i^{th} word. To clarify, let us say the sentence “I see a building” is being processed (the padding step is omitted for the sake of brevity), then we are interested in vectorizing all its sub-sentences in the forward direction, i.e. “I”, “I see”, “I see a”, “I see a building”, and in a backward direction, i.e. “building”, “building a”, “building a see”, “building a see I”. To achieve that, we concatenate forward and backward memory cells (which are equivalent to hidden states in case of GRUs) at each time step i . This results in $(D_s = 2 \times H_s)$ -dimensional sub-sentence representations, which are concatenated into $L_s \times D_s$ matrix, referred to as *sub-sentence encoding*. The sub-network used for obtaining such encoding will be referred to as *RefNet encoder* (blue block in Fig. 1).

The rationale behind using sub-sentence encoding is that the same word can be either a part of RE or not, depending on the preceding and succeeding words. Consider the following two passages:

1. “You can see a train station to the right, it is for commuter trains and is called City’s Eastern.”
2. “You can see a train departing from the second track. It one of the city’s eastern parts.”

Finally, the sub-sentence encoding is fed into the softmax layer, which produces a $L_s \times 3$ matrix with i^{th} row representing a probability distribution over the possible labels, i.e. **O**, **B-REF**, **I-REF**, for the i^{th} word. RefNet is trained by minimizing the cross-entropy loss using Adam optimization method, presented in (Kingma and Ba, 2014).

3.2 Reference resolution

The resolution step implies matching a textual referring expression with a candidate geographical object. For performing such reference resolution (sRR) we employ another neural network architecture, dubbed as *SpaceRefNet* (see Fig. 2).

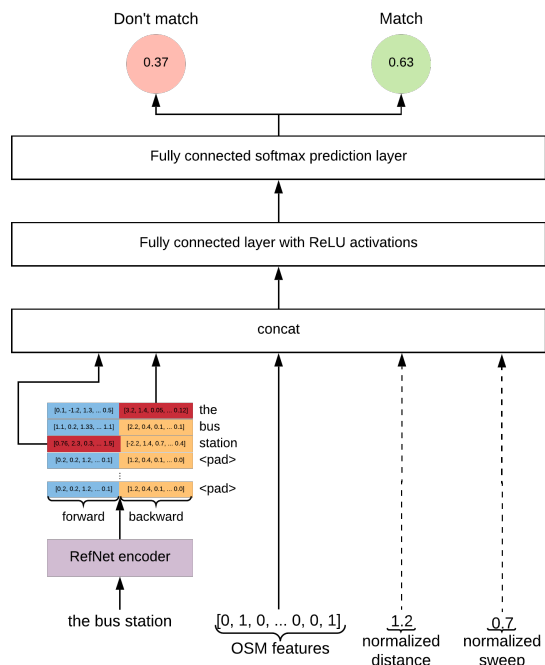


Figure 2: SpaceRefNet architecture diagram. The purple block is the pre-trained layer; the dashed arrows denote optional connections (*Best viewed in color*)

SpaceRefNet takes as an input a referring expression (RE) and a candidate geographical object, denoted as *the candidate*. The RE is encoded using the pre-trained RefNet encoder, resulting in a $L_s \times D_s$ matrix, containing forward and backward $\frac{D_s}{2}$ -dimensional encodings for every sub-sentence. The final RE encoding is then the concatenation of the vectors containing forward and backward sub-sentence encodings for the whole sentence excluding paddings (the selected vectors are shown in a dark red in Fig. 2), resulting in D_s -dimensional vector. The input candidate is fed as only an OSM representation, or together with the distance and/or sweep features (see details in Sect. 4). The vectors obtained after encoding both RE and candidate are then concatenated and passed to the fully connected layer with N_h hidden units having rectifier activation functions. The final fully connected softmax prediction layer produces the probability of a match between the RE and the candidate.

SpaceRefNet is trained by optimizing the

weighted cross-entropy loss using the Adam optimization method. Weights for the loss function are introduced, because the SpaceRef dataset has a high class imbalance – it has much more negatives (when a candidate and an RE mismatch) than positives (when a candidate and an RE match). To counteract this, a contribution of each data point (a candidate and an RE) to the global loss is adjusted using class-dependent multiplication factors (negatives receive lower weights than positives), allowing us to penalize the network more for the mistakes made on positive data points.

Such an architecture allows handling the cases when an RE has any number of referents (0, 1, or more) in the candidate set, which is an advantage compared to previously developed methods that required more ad-hoc solutions, e.g. setting an experimentally selected probability threshold in (Götze and Boye, 2017).

4 Data and processing

The utilized data consists of three datasets:

- a slightly corrected version of a publicly available *SpaceRef* dataset (Götze and Boye, 2016) (used for RefNet and SpaceRefNet training);
- a number of walks, containing the subjects' descriptions of their vicinity, which is referred to as *WalksRef* dataset¹ (used for RefNet training);
- a number of dialogues with manually annotated REs, referred to as *DialogsRef*, taken from the publicly available *Cornell Movie-Dialogs Corpus* (Danescu-Niculescu-Mizil and Lee, 2011) and *DailyDialog* corpus (Li et al., 2017) (used for RefNet training only).

The SpaceRef dataset contains descriptions of immediate geographical environment given by pedestrians following predefined routes. REs in the spoken utterances were manually annotated. GPS information representing a physical context is also available.

Referring expressions (REs) in SpaceRef are mostly noun phrases (NPs). Some example utterances with the referring expressions (underlined) include:

¹is publicly available at <https://traktor.csc.kth.se>

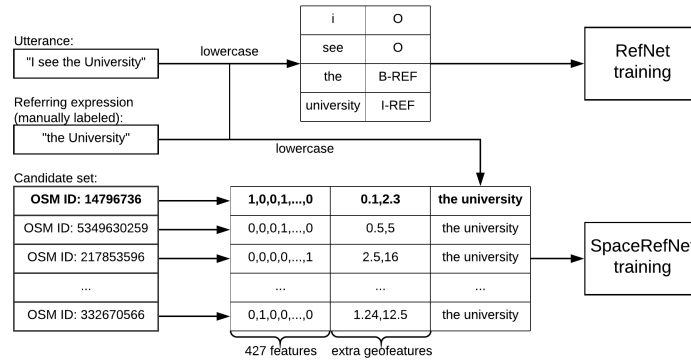


Figure 3: Data processing for training. The rows in **bold** denote a positive data point for SpaceRefNet training, i.e. the one where RE describes the given OSM entity.

- **indefinite and definite NPs**, e.g. “... walking down some stairs”, “there is a fountain to my left”;
- **NPs with interjections**, which should be excluded from an RE, e.g. “I am near eh the red eh brick building”;
- **demonstratives**, e.g. “... standing to the right of this building”;
- **proper names**, e.g. “... I am now passing 7-Eleven store”.

However, not all NPs in these categories are REs, for instance,

- in the utterance “Do you know if there is a subway station nearby?”, “a subway station” is not an RE, since it has no intention of referring to a specific geographic object;
- in the utterance “This architectural style I like the most”, a demonstrative “this architectural style” is not an RE;
- in the utterance “The statue in front of the library portrays Carl Linnaeus”, a proper name “Carl Linnaeus” is not an RE either.

SpaceRef and WalksRef contain mostly the utterances with at least one RE in them. Hence, the number of negative examples (NPs that are not REs), was not sufficient for training the neural network. With this in mind, the DialogsRef corpus was annotated providing more negative examples to improve the robustness of the trained models.

The candidate sets were *regenerated* for each referring expression by first computing lines-of-sight around the pedestrian location in 1 degree

steps using a “visibility engine”, inspired by (Boye et al., 2014). The lines-of-sight were computed in every direction between -100 and 100 degrees with respect to the pedestrian’s walking direction. The closest OSM nodes and ways, intersecting with these lines-of-sight, were included into the candidate set as OSM identifiers.

Each candidate referent is then encoded using the following features:

- 427 binary *OSM type features*, as described in (Götze, 2016, Subsection 4.3.2).;
- *the distance feature*: the logarithm of a distance between pedestrian’s and object’s locations;
- *the sweep feature*: a number of lines-of-sight intersecting with an object divided by 360.

The last two features are referred to as *extra geo-features* and a numeric vector consisting of these 429 features – as *geoencoding*.

The available data were transformed differently for training RefNet and SpaceRefNet (see Fig. 3). RefNet training requires the data to be labeled using BIO-REF labeling strategy (as mentioned before), i.e. each word in an utterance is either at the beginning of an RE and gets a label *B-REF*, or inside an RE and gets a label *I-REF*, or is not a part of an RE and gets a label *O*. SpaceRefNet training requires labeling of tuples (RE, OSM entity) with a binary label (1 if RE describes this OSM entity, 0 otherwise).

Finally, note that the training data for SpaceRefNet are heavily (and necessarily) skewed: for every referring utterance from the user, there will be about 30 candidate referents to consider, and in

most cases all of them but one are not the referent the user intended. Thus, there will always be many more negative examples than positive examples in any dataset.

5 Models for comparison

5.1 Referring expression identification baseline

The REs are mostly represented by the noun phrases (NP), so the natural baseline is just returning every found NP as a candidate RE. The baseline was implemented as follows:

- a part-of-speech (POS) tag was defined for each word in an utterance using the Stanford POS tagger for English (Toutanova et al., 2003), to be more specific, the *wsj-0-18-bidirectional-distsim* version was used;
- the POS-tagged utterance is then parsed using NLTK RegexpParser (Bird et al., 2009), supplied with the following grammar:

```
NP: {
  (<DT>? (<RB . *>* <JJ>*) * <NN . *>+ <IN>*) +
}
```

- all found NPs are returned as REs.

5.2 Reference resolution baseline

The natural RR baseline is just querying the OSM database and checking for geographical objects with an OSM property containing at least one word from the utterance (except stop words) either in a property key or value. For example, consider two utterances, (1) “a very nice big park” and (2) “a huge green area”, are being matched with the geographical object “Stanford Arboretum” (see Fig. 4). The utterances are first split by space and then all the stop words are removed. The result would be as follows: (1) {very, nice, big, park} and (2) {huge, green, area}.

Each word is then checked against all properties of the OSM object (both keys and values are checked). The first utterance will then be matched with “Stanford Arboretum”, because the “leisure” tag has value “park”, which is part of the utterance. The second utterance will not be matched, since none of the words matches any of the property keys or values.

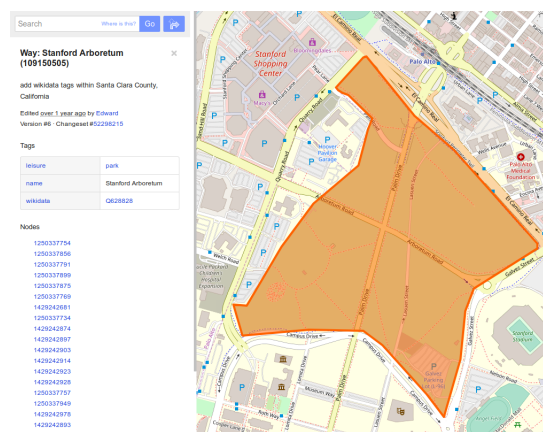


Figure 4: OpenStreetMap (OSM) representation of “Stanford Arboretum”

6 Experimental results

In all experiments the networks were trained for a maximum of 100 epochs with the early stopping (patience of 5 epochs).

6.1 Spatial RE identification

A RefNet was trained on the SpaceRef, Walks-Ref and DialogRef corpora. The data was split into training set (around 90% of the data containing around 90% of REs) and a test set (the remaining data). RefNet has a large number of hyper-parameters making a grid search computationally infeasible. Instead, some of the hyper-parameters were fixed to the values found through manual experiments and the others were found using a random search (Bergstra and Bengio, 2012). The hyper-parameter space was searched during 60 random trials evaluating RefNet’s performance for each of hyper-parameter’s combination using 5-fold cross-validation.

The model’s performance was assessed by computing precision and recall, which can be done in several ways. The most straightforward way is to consider a word and its BIO-REF label as one datapoint, and compute precision and recall based on this. However, our aim is to measure how well the network identifies full referring expressions. Therefore, we’ve considered one datapoint being a tuple of all words in the referring expression together with their respective labels. The datapoint is then considered as a true positive only if *all* the words in the RE are correctly labeled (*exact match*). In order to assess the magnitude of method’s errors, we say that a *partial match* occurs if the starting word is correctly labeled and there are no more than 2 errors in the rest of the

expression.

During the experiments the batch size was fixed to 128, the maximum sentence length set to 100 and maximum word length to 30. The best-performing RefNet model found after the performed hyper-parameter search had 24 hidden units on the character-level BiRNN layer, 51 hidden unit on the sentence-level BiRNN layer, a learning rate of 0.002. A regularization in the form of dropout was applied with the probabilities of keeping the input, the state and the output being 0.7, 0.75, 0.95 for the character-level BiRNN and 0.8, 0.95, 0.95 for the sentence-level BiRNN respectively. The found RefNet model achieved the following performance (averaged over 5 folds):

- a precision of 0.7846 (partial precision of 0.8083);
- a recall of 0.6608 (partial recall of 0.784).

Evaluating the same model on the test set resulted in a better performance compared to the baseline (see Table 1).

Metric	Baseline	RefNet
Correct sentences (%)	21.02	77.08
Precision	0.1204	0.5457
Recall	0.2997	0.5531
Partial precision	0.1663	0.7204
Partial recall	0.4142	0.7302

Table 1: Performance of different methods for solving spatial referring expression identification (sREI) task on the test set

6.2 Spatial reference resolution

SpaceRefNet was trained exclusively on the SpaceRef corpus. The data were split into training set (around 80% of the data), validation set (around 10% of the data) and a test set (the remaining data). SpaceRefNet has a smaller number of hyper-parameters than RefNet, but a higher-dimensional input data (429 dimensions + RE encoding size). Hence, the combination of random search with cross-validation becomes computationally infeasible. Given the nature of SpaceRef data, i.e. the subjects walking along the routes in the same vicinity, the datapoints are more homogeneous compared to DialogRef and WalksRef used for RefNet training. Keeping in mind everything mentioned above, hyper-parameter space was searched using a combination of random and

manual search relying on the performance on the held-out validation set.

During random search, the batch size was fixed to 256. The best found SpaceRefNet model had 32 hidden units, negatives weighted with 0.25 and positives – with 1 in the loss function, a learning rate of 0.001 and used both distance and sweep features. The model’s performance was evaluated by computing precision, recall and F1-score for positives (matches between an RE and a candidate) and negatives (mismatches). The performance on the validation set was:

- for positives, precision of 0.5854, recall of 0.4444, F1-score of 0.5053;
- for negatives, precision of 0.9860, recall of 0.9748, F1-score of 0.9804;

Evaluating the same model on the test set resulted in a better performance compared to the baseline and previously reported results in the literature (see Table 2). Additionally a percentage of completely correctly labeled sentences is provided.

Metric	Baseline	WAC	SpaceRefNet
Prec. (p)	0.5588	0.40	0.6105
Rec. (p)	0.2043	0.45	0.6237
F1 (p)	0.2992	0.42	0.6170
Prec. (n)	0.9757	0.98	0.9883
Rec. (n)	0.995	0.98	0.9876
F1 (n)	0.9853	0.98	0.9879

Table 2: Performance of different methods for solving spatial reference resolution (sRR) task on the test set (“(p)” stands for positives, “(n)” stands for negatives), “WAC” stands for words-as-classifiers method (results reported in (Götze and Boye, 2017)).

7 Discussion

The designed methods have shown promising results in solving exophoric spatial reference resolution (sRR) beating the respective baselines and earlier reported results in the literature. It should be noted that sRR is a complicated task with non-trivial subproblems. Identifying REs in spoken utterances gets complicated because of multiple challenges:

- *unclear sentence segmentation* in spoken utterances results in the utterances like “I am passing the shop **on my left on my right**

there is a bank”, the phrase “on my left” describes the RE “the shop”, whereas the phrase “on my right” describes “the bank”;

- *ASR errors* can lead to the utterances like “I’m crossing the street on my **rights**”;
- *interjections and self-corrections* result in utterances like “there is another shop eh called ehm jer- jersey shop”.

A problem arises because of the possible differences in the interpretation. Consider the utterance “on my right is the embassy of Poland in an old fantastic villa”. Depending on the interpretation, one might find either two REs “the embassy of Poland” and “an old fantastic villa” referring to the same geographic object or only one RE “the embassy of Poland in an old fantastic villa” referring to the same object. Such interpretation differences have not been considered while evaluating *RefNet*.

Resolving spatial references is even more tricky, since each found RE has mostly only one correct referent out of 30 candidates on average, making data very unbalanced. Furthermore, one RE can have multiple referents, e.g. the streets often consist of many different parts in OSM, or have no referents, e.g. some specifics about the geographical objects (“a big clock on the wall of the university”), or outdated information.

Ongoing work includes incorporating this reference resolution model into our wayfinding spoken dialogue system and collecting more data to improve the model.

Acknowledgements

This work was supported by Vinnova (Sweden’s Innovation Agency) within the project 2018-01672, and by PTS (The Swedish Post and Telecom Authority) within the project 17-10324.

References

- James Bergstra and Yoshua Bengio. 2012. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13(Feb):281–305.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. ”O’Reilly Media, Inc.”.
- Johan Boye, Morgan Fredriksson, Jana Götze, Joakim Gustafson, and Jürgen Königsmann. 2014. Walk this way: Spatial grounding for city exploration.

In *Natural interaction with robots, knowbots and smartphones*, pages 59–67. Springer.

- Howard Chen, Alane Suhr, Dipendra Misra, Noah Snaveley, and Yoav Artzi. 2019. Touchdown: Natural language navigation and spatial reasoning in visual street environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12538–12547.
- Edward H Cornell and Elaine Greidanus. 2006. Path integration during a neighborhood walk. *Spatial Cognition and Computation*, 6(3):203–234.
- Cristian Danescu-Niculescu-Mizil and Lillian Lee. 2011. Chameleons in imagined conversations: A new approach to understanding coordination of linguistic style in dialogs. In *Proceedings of the 2nd Workshop on Cognitive Modeling and Computational Linguistics*, pages 76–87. Association for Computational Linguistics.
- Michel Denis. 1997. The description of routes: A cognitive approach to the production of spatial discourse. *Cahiers de Psychologie*, 16:409–458.
- Nikos Engonopoulos, Martin Villalba, Ivan Titov, and Alexander Koller. 2013. Predicting the resolution of referring expressions from user behavior. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1354–1359.
- Kotaro Funakoshi, Mikio Nakano, Takenobu Tokunaga, and Ryu Iida. 2012. A unified probabilistic approach to referring expressions. In *Proceedings of the 13th annual meeting of the special interest group on discourse and dialogue*, pages 237–246. Association for Computational Linguistics.
- Joy Goodman, Phil Gray, Kartik Khammampad, and Stephen Brewster. 2004. Using landmarks to support older people in navigation. In *International Conference on Mobile Human-Computer Interaction*, pages 38–48. Springer.
- Jana Götze. 2016. *Talk the walk: Empirical studies and data-driven methods for geographical natural language applications*. Ph.D. thesis, KTH Royal Institute of Technology.
- Jana Götze and Johan Boye. 2016. Spaceref: A corpus of street-level geographic descriptions. In *LREC*.
- Jana Götze and Johan Boye. 2017. Reference resolution for pedestrian wayfinding systems. In *International Conference on Geographic Information Science*, pages 59–75. Springer.
- Mordechai Haklay and Patrick Weber. 2008. Openstreetmap: User-generated street maps. *IEEE Pervasive Computing*, 7(4):12–18.
- Alycia M Hund and Jennifer L Minarik. 2006. Getting from here to there: Spatial anxiety, wayfinding strategies, direction type, and wayfinding efficiency. *Spatial cognition and computation*, 6(3):179–201.

- Pentti Kanerva, Jan Kristoferson, and Anders Holst. 2000. Random indexing of text samples for latent semantic analysis. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 22.
- Casey Kennington and David Schlangen. 2015. Simple learning and compositional application of perceptually grounded word meanings for incremental reference resolution. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 292–301.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Carol A Lawton and Janos Kallai. 2002. Gender differences in wayfinding strategies and anxiety about wayfinding: A cross-cultural comparison. *Sex roles*, 47(9-10):389–401.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. Dailydialog: A manually labelled multi-turn dialogue dataset. *arXiv preprint arXiv:1710.03957*.
- Kevin Lynch. 1960. *The image of the city*, volume 11. MIT press.
- Cynthia Matuszek, Liefeng Bo, Luke Zettlemoyer, and Dieter Fox. 2014. Learning from unscripted deictic gesture and language for human-robot interactions. In *AAAI*, pages 2556–2563.
- Andrew J May, Tracy Ross, Steven H Bayer, and Mikko J Tarkiainen. 2003. Pedestrian navigation aids: information requirements and design implications. *Personal and Ubiquitous Computing*, 7(6):331–338.
- Ian Moar and Gordon H Bower. 1983. Inconsistency in spatial knowledge. *Memory & Cognition*, 11(2):107–113.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Zahar Prasov and Joyce Y Chai. 2010. Fusing eye gaze with speech recognition hypotheses to resolve exophoric references in situated dialogue. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 471–481. Association for Computational Linguistics.
- Tracy Ross, Andrew May, and Simon Thompson. 2004. The use of landmarks in pedestrian navigation instructions and the effects of context. In *International Conference on Mobile Human-Computer Interaction*, pages 300–304. Springer.
- David Schlangen, Sina Zarri , and Casey Kennington. 2015. Resolving references to objects in photographs using the words-as-classifiers model. *arXiv preprint arXiv:1510.02125*.
- Niels Schutte, John Kelleher, and Brian Mac Namee. 2010. Visual salience and reference resolution in situated dialogues: A corpus-based evaluation.
- Bart Sheehan, Elizabeth Burton, and Lynne Mitchell. 2006. Outdoor wayfinding in dementia. *Dementia*, 5(2):271–281.
- Kristina Toutanova, Dan Klein, Christopher D Manning, and Yoram Singer. 2003. Feature-rich part-of-speech tagging with a cyclic dependency network. In *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology-Volume 1*, pages 173–180. Association for Computational Linguistics.
- Ann E Vandenberg, Rebecca H Hunter, Lynda A Anderson, Lucinda L Bryant, Steven P Hooker, and William A Satariano. 2016. Walking and walkability: Is wayfinding a missing link? implications for public health practice. *Journal of physical activity and health*, 13(2):189–197.

Which aspects of discourse relations are hard to learn? Primitive decomposition for discourse relation classification

Charlotte Roze¹, Chloé Braud¹, Philippe Muller²

¹ Université de Lorraine, CNRS, LORIA
Nancy, France

{charlotte.roze, chloe.braud}@loria.fr,

² IRIT, CNRS, Université of Toulouse
Toulouse, France

philippe.muller@irit.fr

Abstract

Discourse relation classification has proven to be a hard task, with rather low performance on several corpora that notably differ on the relation set they use. We propose to decompose the task into smaller, mostly binary tasks corresponding to various primitive concepts encoded into the discourse relation definitions. More precisely, we translate the discourse relations into a set of values for attributes based on distinctions used in the mappings between discourse frameworks proposed by Sanders et al. (2018). This arguably allows for a more robust representation of discourse relations, and enables us to address usually ignored aspects of discourse relation prediction, namely multiple labels and underspecified annotations. We study experimentally which of the conceptual primitives are harder to learn from the Penn Discourse Treebank English corpus, and propose a correspondence to predict the original labels, with preliminary empirical comparisons with a direct model.

1 Introduction

Discourse parsing is a crucial task for natural language understanding, as it accounts for the coherence of a text by identifying semantic and pragmatic links between sentences and clauses. The links are sometimes marked by explicit lexical items, so-called discourse connectives, but very often they rely on several lexical cues, contextual interpretation or even world knowledge, in which case they are called “implicit” relations. Automating discourse parsing consists in finding which sentences or clauses are directly related in a text, and with what type of semantico-pragmatic relation. The examples below demonstrate each type of relation, with the explicit discourse connective marked in bold, and example labels inspired by the Penn Discourse Treebank 2.0 (Prasad et al., 2008) relation set.

(1) Climate change is caused by anthropic activities, **but** politics are not doing anything about it.

Comparison. Concession. Contra-expectation

(2) Climate is changing. Humans generate too much CO₂.

Contingency. Cause. Reason

Several theoretical frameworks exist for discourse analysis, the most well-known being Rhetorical Structure Theory (RST, Mann and Thompson, 1988), and Segmented Discourse Representation Theory (SDRT, Asher and Lascarides, 2003). The Penn Discourse Treebank (PDTB, Prasad et al., 2008) is an English annotated corpus with its own theoretical assumptions. It is the largest resource for discourse relations and has been used in several studies to demonstrate the difficulty of automatically identifying implicit discourse relations, e.g. (Xue et al., 2016; Bai and Zhao, 2018). The PDTB relies on a three-level hierarchy of rhetorical functions, and multiple relations can be annotated for each example.

As empirical models have shown rather low results for implicit relation classification, with only incremental improvements in spite of the variety of approaches that have been tried, it appears a lot of the necessary information is still not leveraged in discourse parsing.

But it could be argued also that the difficulty lies in the way we model the task, especially these labels on which there is no consensus and generally a low inter-annotator agreement.

We argue here that, even if the label sets differ, all frameworks propose to encode the same range of pragmatic phenomena, and that decomposing the relations into simpler conceptual primitives could help to understand where the real difficulty lies, and, eventually, to improve classification performance. We thus experiment with clas-

sification tasks where we try to predict these primitives of the discourse relations rather than the relations themselves.

More precisely, we experimentally test Sanders et al. (2018)’s recent proposal of an inventory of so-called *dimensions* (called here *primitives*) of the discourse relations that could be seen as an interface between the various existing frameworks.

Our first contribution is thus to implement this mapping, from annotated relations to a set of primitives, and from a predicted set of primitives to compatible relation labels.

Our second contribution is an empirical investigation of the separate primitives and how difficult they are to predict. One advantage of this approach is that it can provide underspecified labels, which is why we focus for now on the PDTB, as its hierarchical organisation of relation types naturally lends itself to a classification mixing granularities. Our approach can also address predicting or comparing against multiple labels between pairs of sentences or clauses. This allows us to stay closer to the annotation, contrary to all existing work, limited to a subset of relations.

Finally, we hope to provide a framework to investigate the validity of different conceptual decompositions of discourse relations.¹

This paper is organized as follows. In Section 2, we briefly review work on discourse relation identification. In Section 3, we present discourse relation decomposition, with a focus on the mapping presented in (Sanders et al., 2018), before detailing, in Section 4, our proposal for an operational mapping. The Section 5 presents our experimental framework – the systems compared and the evaluation strategy. Finally, we detail in Section 6 the models built and the data used, before reporting our results in Section 7.

2 Discourse relation classification

Previous work on discourse relation identification generally separated the classification of implicit and explicit examples, and mainly focused on implicit ones, considered as the hardest task. Performance on this task are, however, still low: the current best are reported in (Bai and Zhao, 2018), where it is proposed to augment word embeddings with subword and contextual embeddings, and to combine sentence and sentence pair representa-

¹Our code is available at <https://gitlab.inria.fr/andiamo/relations>.

tions. They report 45.73 to 48.22% in accuracy – depending on the sections used for evaluation – for level 2 relation classification (11 labels), and 51.06% in F_1 for multiclass classification of level 1 relations (4 labels).

For explicit relation classification, the last scores come from the CoNLL shared tasks on shallow discourse parsing (Xue et al., 2015, 2016). Mihaylov and Frank (2016) use similarity measures based on word embeddings and report 78.34% in F_1 on blind test and 89.80% on section 23. Kido and Aizawa (2016) propose to build a specific classifier for *Comparison* subtypes and report 75.43% on blind test and 90.22% on section 23. These scores are computed on relations of the PDTB, with a modified inventory of 20 relations designed to make data more balanced by mixing various levels of the hierarchy.

The organizers of the shared tasks also provide scores for all relations: at best 54.60 on blind test and 64.34% on section 23 (Xue et al., 2016).

All previous work made simplifying assumptions for the task: systems are restricted to a subset of relations, and ignore multiple annotations and under-specified annotations of relations. On the contrary, our approach aims at considering the problem of discourse relation prediction in the most general way.

3 Existing approach for mapping relations into primitives

Discourse frameworks and their corresponding annotated corpora rely on different assumptions, among them the set of discourse relations they consider, covering overlapping or identical concepts under different names and definitions, and they are hard to reconcile.

There have been a few attempts to formalize the various types of information encoded by discourse relations, and give it some structure (Hovy, 1990; Knott, 1997), or provide a semantics for the underlying principles (Chiarcos, 2014), without clear-cut criteria to decide on the most appropriate set of relations. The PDTB addresses the problem by providing a hierarchy of relations, allowing for various levels of underspecification, but without much justification other than annotation operational constraints.

3.1 Cognitive approach to Coherence Relations

More recently, within the context of the COST TextLink Action,² Sanders et al. (2018) provided a mapping into *dimensions* for sets or hierarchies of relations from RST, PDTB and SDRT. These mappings rely on an extended version of the primitives originally introduced in the Cognitive approach to Coherence Relations or CCR (Sanders et al., 1992, 1993). In the following we will use the term *primitive* to describe what is rather ambiguously called *dimension* in (Sanders et al., 2018).

In CCR, the link between two discourse units is described by values for a set of primitives. The core CCR primitives are: *basic operation*, *polarity*, *source of coherence*, *implication order*, and *temporality*. According to Sanders et al. (2018), these primitives are shared by all coherence relations and are validated by a number of psycholinguistic and/or corpus-based studies.

We use the following notation: P and Q are two propositions (events, states, speech acts, claims, etc.) expressed in the discourse units linked by a relation. Each relation is characterized by the way in which its arguments map onto P and Q .

Basic operation This primitive makes a distinction between *additive* relations (typically expressed by connectives *and* or *also*) that involve a logical conjunction ($P \& Q$) and *causal* relations (typically expressed by connectives *because* or *since*) that involve an implication ($P \rightarrow Q$).

Polarity Polarity distinguishes between *positive* and *negative* (or adversative) relations. Negative relations (expressed for instance by connectives *but*, *although* or *even if*), differ from positive relations (expressed for instance by *because*) in that they imply the negation of either P or Q or some of their implications in their semantics. Note that this negation does not need to be explicit/linguistically marked. In (3), the negated proposition would be that *the biofuel costs more*, as an expected consequence of the higher production costs. Note that this primitive must not be confused with sentiment polarity.

- (3) The biofuel is more expensive to produce, **but** by reducing the excise-tax the government makes it possible to sell the fuel for the same price.

Comparison. Concession. Contra-expectation

²See <http://www.textlink.ii.metu.edu.tr>.

Source of coherence This primitive has two possible values named *objective* and *subjective* in CCR. It refers to a common distinction in the literature, for instance *subject matter* versus *presentational* relations for Mann and Thompson (1988). *Objective* relations link discourse units at the level of their propositional content (*as a result* generally expresses an *objective* relation), whereas *subjective* relations operate at epistemic or speech act level: the speaker is “involved in the construction of the relation” (Sanders et al., 2018) (*since* seems to have a preference for marking *subjective* relations).

Implication order This primitive is only applicable for *causal* relations (value for this primitive is set to non-applicable (NA) for *additive* relations). For relations involving an implication $P \rightarrow Q$, it indicates the order in which P and Q are described in the linguistic arguments S_1 and S_2 of the relation. If S_1 expresses P (antecedent), implication order is *basic*, whereas if S_1 expresses Q (consequent), implication order is *non-basic*. Typically, connectives *thus* and *because* respectively express relations in *basic* and *non-basic* order.

Temporality A relation can have a temporal aspect or not, and when it does it can be *chronological* (*then*), *anti-chronological* (*previously*), or *synchronous* (*meanwhile*).

Additional features Sanders et al. (2018) introduce additional features that represent distinctions which are more detailed than those used in the original CCR framework, in order to provide the most specific mapping possible. These additional features are: *conditional*, *alternative*, *specificity* (and refinements: *specificity-equivalence*, *specificity-example*), *goal* and *list*. Their values are negative by default (-). In our experiments, we did not retain features that only apply to part of the relations falling under the respective category (*goal* and *list*). We keep as primitives: *conditional* (*if*, *unless*), *alternative* (*or*) and *specificity* (*in particular*, *in fact*). In order to have quite generic primitives, we merged refinements on *specificity* into one primitive, so that each primitive is positive (+) for more than one PDTB label.

The contribution of Sanders et al. (2018) is to provide a (arguably) complete mapping to make existing annotation systems compatible, and Demberg et al. (2017) test the approach by applying PDTB and RST mappings to existing annotations:

Class	Type	Subtype	Pol.	Basic op.	Impl. order	SoC	Temp.
<i>Comparison</i>			neg	NS	NS	NS	NS
<i>Comparison</i>	<i>Contrast</i>	<i>Juxtaposition</i>	neg	add	NA	obj	NS (any)
<i>Comparison</i>	<i>Contrast</i>	<i>Opposition</i>	neg	add	NA	obj	NS (any)
<i>Comparison</i>	<i>Pragmatic contrast</i>		neg	add	NA	sub	NS (NA)
<i>Comparison</i>	<i>Concession</i>		neg	cau	NS	NS	NS
<i>Comparison</i>	<i>Concession</i>	<i>Expectation</i>	neg	cau	non-b	NS (obj sub)	NS (anti NA)
<i>Comparison</i>	<i>Concession</i>	<i>Contra-expectation</i>	neg	cau	basic	NS (obj sub)	NS (anti NA)
<i>Comparison</i>	<i>Pragmatic concession</i>		neg	cau	NS	sub	NS

Table 1: Sample of our classification into core primitives, for relations within the class Comparison. Primitives are *polarity* (Pol.), *basic operation* (Basic op.), *implication order* (Impl. order), *source of coherence* (SoC) and *temporality* (Temp.). Bold indicates modified or new values w.r.t. Sanders et al. (2018) (see Section 4.1). Original ones are indicated in parenthesis. NS (non-specified) unifies different unspecified labels from the original model.

they used common portions of PDTB 2.0 and RST-DT, in order to test the validity of the mapping. The outcome is that only a partial mapping is possible at this stage, because of discourse segmentation issues, and a lot of contextually underspecified or ambiguous correspondences.

As a first step we focus on providing a practical correspondence between PDTB annotations and the set of CCR primitives described by Sanders et al. (2018). It is the mapping we rely on in our experiments (with a few changes on the possible values for each primitive, see Section 4).

4 Proposal for an operational mapping

In this study, we focus on the PDTB 2.0 (Prasad et al., 2007). This corpus has been annotated with explicit and implicit discourse relations.³ As previously said, in the PDTB, relations are organized into a three-level hierarchy with 4 coarse-grained classes, 16 types and 23 subtypes. Examples can be annotated at any levels and annotators were asked to choose a more general relation when hesitating between different relations within a group; some annotation disagreements were adjudicated by annotating at the upper level. Moreover, annotators were allowed to suggest up to two relations per explicit example, and up to four per implicit.

PDTB annotation thus presents several particularities that are almost always ignored by automated approaches: relations at different levels of granularity, under-specified relations and possibly multiple relations for a single pair of text segments. Moreover, studies on discourse relation classification are always limited to a subset of re-

lations, for example by focusing on level 1 or 2 relations.

Decomposing relations into primitive concepts allows us to tackle the problem in all its generality. First, the primitives can precisely be used to encode distinctions at the finest level of the hierarchy (level 3) such as distinction on *source of coherence* for pragmatic (*subjective*) or level 3 (the finest level) relations. Second, even when several relations cannot be distinguished by their values for each primitive, we do not need to merge them: they are mapped into the same set of values for dimensions, and in the reverse mapping (see Section 5.1), they can be mapped into a subset of relations. Finally, we are not limited by the problem of small number of annotated instances for some relations.

In this section, we describe specificities of our operational mapping.

4.1 Primitives and possible values

The set of primitives and their possible values used in our experiments are presented in Figure 1, along with their distribution in our training dataset (see Section 6.1) after operational mapping. Possible values for core primitives present minor changes compared to the ones adopted by Sanders et al. (2018). For additional or binary primitives, possible values are unchanged: they are either negative (default value -) or positive (+). For core primitives, we proposed several modifications motivated by the fact that the *operational* mapping is applied to data for being used as input of classifiers for each primitive (see Section 5). In particular, we need to deal with cases of ambiguity – i.e. for some relations, a primitive is associated with a set of values, each being possible –, under-specified and multiple annotated relations.

³As in previous work on this task, we ignore the Entity relation. Note that no mapping was provided in (Sanders et al., 2018) for this relation.




Figure 1: Distribution of values for each dimension

Non-specified value (NS) For all core primitives (i.e. non binary primitives), in addition to values described in previous section (e.g. *additive* and *causal* for primitive *basic operation*), we add the value *NS* (non-specified) to the set of possible values.

NS value does not exist as a “label” in (Sanders et al., 2018) mapping, but there are cases of ambiguity/under-specification: in the original CCR mapping, value for the primitive *source of coherence* is set to *obj|sub* for a number of relations, primitive *temporal order* has value *syn|chron|NA* for *Expansion.List*, etc. In our mapping, when there is an ambiguity on a primitive value, we associate the value *NS* (see Table 1 for our mapping for class *Comparison* relations).

NS value is also used for ambiguities raised by the need to associate primitive values to relations that are not *end-labels* of the PDTB hierarchy, end-labels being relations at level 2 that have no subtypes (such as *Temporal.Synchronous*) or relations at level 3 (such as *Contingency.Cause.Result*). Sanders et al. (2018) provide a mapping for each end-label but not for less specific labels. Since PDTB contains examples annotated with level 1 (classes) or 2 (types) relations which are not end-labels – under-specified relations –, we also need to provide a mapping into primitives for these relations in our experiments. For example, we set primitive *basic operation* to value *NS* for *Comparison*, as some relations within this class are *additive*, and some others are *causal* (see Table 1).

Non-applicable value (NA) We keep value *NA* for dimension *implication order*, associated with relations that do not involve an implication (*additive* relations).

On the other hand, we remove it for dimension *temporal order*. This is motivated by the fact that relations from *Temporal* class have a somewhat

special status among discourse relations: it is not always clear whether they are rhetoric or semantic relations (especially when annotated in addition of another relation). *Temporal* relations represent 66.3% of multiple relations in PDTB, and they can co-occur with relations from any other class. Furthermore, temporal relations can co-occur with relations which are associated with the value *NA* (non-applicable) for *temporal order* in the original mapping of Sanders et al. (2018).⁴

As there is no relation in PDTB data that seem to be incompatible with a specified value for *temporal order*, we remove *NA* value for this primitive (it is present in possible values for *temporal order* in CCR), and keep only *NS* as a default value.

4.2 Multiple relations: merging sets of primitive values

On the overall corpus used in our experiments (see Section 6.1), 4.4% relations are multiple relations, i.e. several relations have been annotated in the original PDTB. As previously said, Sanders et al. (2018) applied their mapping into values per primitive on RST-DT and PDTB’s common sections. However, they give no information about a mapping into primitives for cases where multiple relations were annotated in the PDTB: they select the PDTB relation that most closely corresponds to the RST label.

Our goal being different here, we want to take all annotated information into account. In case of multiple relations, we map each relation into a set of primitive values, and then merge values when they are different. Our actual merging preferably outputs non-specified values, but other options should be tested in future work, e.g. keep most specific values.

For *basic operation*, *polarity*, *source of coherence* and *temporal order*, if values to be merged are different, the primitive value is set to *NS*.

For binary primitives (*conditional*, *alternative*, *specificity*), value is set to *positive* (+) if at least one of the merged values is *positive*, and *negative* (-) otherwise.

For *implication order*, if one of the two distinct values to be merged is *NA* and the other is not (i.e. *basic*, *non-basic* or *NS*), we keep the second value. If the two distinct values are different from *NA*, *implication order* is set to *NS*.

⁴ For instance, there are 198 co-occurrences of *Temporal.Synchrony* and *Expansion.Conjunction* in our training dataset.

4.3 Refinements and adding of missing relation

When mapping PDTB relations into primitives, we operated refinements on occurrences of *Expansion.Alternative.Disjunctive*, whose values for primitives are quite under-specified when strictly applying the mapping of Sanders et al. (2018): values are non-specified (*NS*) for *basic operation* and *source of coherence*, and we do not know whether the additional feature *conditional* or *alternative* must be set to a positive value (+). The only specified primitive is *polarity*, which is *negative*. Leaving this level of under-specification would mean having the same set of primitive values for class *Comparison* and sub-type *Expansion.Alternative.Disjunctive*.

But as suggested by Sanders et al. (2018), markers such as *unless* indicate that the relation is *causal-conditional* rather than *additive-alternative*. For some occurrences of *Expansion.Alternative.Disjunctive*, connectives from PDTB annotations (*unless*, *either...or* and *or*) were used to determine which of the two sub-cases of *Expansion.Alternative.Disjunctive* was present, and associate the correct set of primitive values.

Sanders et al. (2018) provide no mapping for PDTB relation *Comparison.Pragmatic concession*, for which there is no description in PDTB annotation manual. This label being quite explicit, we associate to it the same primitive values as *Comparison.Concession*, except for *source of coherence*, set to *subjective* (see Table 1).

5 Experiments

Our main goal is to assess which primitives are harder to identify, we thus build separate models for each of them, i.e. *basic operation*, *polarity*, *source of coherence*, *implication order*, *temporality*, *conditional*, *alternative* and *specificity* (see Section 3 for definitions).

In addition, we compute the performance of our systems on discourse relations using a reverse mapping from a set of predicted values for each primitive to a relation, or, more precisely, to a set of potential relations. We describe the reverse mapping in Section 5.1.

We also train systems on the task of directly predicting discourse relations, in order to check the validity of our models and to compare to the predictions derived from the primitives.

Recall that we aim at keeping all the particular-

ities of the PDTB annotations, meaning the multiple relations and the relations at different levels of granularity. This calls for specific evaluation metrics, relying on hierarchical multi-label measurement, that we describe in Section 5.2.

5.1 Reverse mapping

Our approach consists in building separate systems dedicated to each primitive, in order to split a hard task into several, arguably simpler tasks. One possible goal of this approach is to predict discourse relations based on the predicted primitives. In order to do that, we need a mapping in the reverse way, i.e. from primitives to (PDTB) relations. Note that we need to map primitives to any level relation, since examples in the PDTB can be annotated with various granularities. This could also be used to limit our system to a set of relations *a posteriori*, without retraining the primitives models. Our reverse mapping, which outputs a set of relations, is defined as follows: starting with a set containing all the possible relations, we remove relations that are not compatible with the primitive values predicted.

More precisely, for each binary primitive, if the predicted value is negative (-), we remove all relations with a positive value for the primitive. For primitives *basic operation*, *polarity*, *source of coherence* and *temporal order*, if predicted value is not *NS*, we remove all relations with a different “specified” (non *NS*) value for the primitive which is different from predicted value. For instance, if *polarity* is *positive*, all relations associated with *negative polarity* are excluded.

For primitive *implication order*, at first, we treated *NA* value as a “specified” value in our reverse mapping: a predicted value *NA* for *implication order* excluded all relations with a non *NA* value for this primitive, i.e. all *causal* relations were removed. This first mapping led to cases where the set of compatible relations was empty. In all these cases, *basic operation* was predicted *causal* and primitive *implication order* was predicted *NA*, which is theoretically inconsistent: if not specified, *implication order* should be *NS*. In order to keep the information specified in other primitives, we decided to treat *NA* value for *implication order* as an *NS* value. It suggests that keeping these two distinct values should be reconsidered.

When all subtypes of a type (or all types un-

der a class) remain in the set of possible relations, we remove these subtypes (or types) from the set, and keep the type (or class) – i.e. the upper level underspecified relation. For instance, if the set contains *Temporal.Asynchronous* and *Temporal.Synchrony*, these labels are removed: only the less specific label *Temporal* remains in the set.

When only some subtypes of a type (or some types under a class) remain in the set of possible relations, we keep them along with the type (or class).

5.2 Evaluation measures

Our experimental setup raises a number of questions with respect to the evaluation: mapping a set of primitive values back to a PDTB label implies there might be underspecifications and corresponding to a disjunction of relations, either a coarse-grain label in the hierarchy or a set of possible relations. To account for the first case, we can apply measures for hierarchical classification; the second case can be taken care of by measures for multi-label classification, which are needed anyway to take PDTB annotations without restrictions. There has not been much work on hierarchical discourse relation classification except (Versley, 2011), and the evaluation was just done at each granularity level, with either exact matching or a Dice coefficient between sets of labels (a relative overlap measure). For a more general measure, we use hierarchical precision and recall (Kiritchenko et al., 2005) on the set of all predicted relations. For instance a predicted X.Y evaluated against a gold X.Z.T would get 0.5 precision (one level correct, one incorrect), and 0.33 recall (2 out of 3 levels missing from the prediction). For multi-labels, all levels are put in the same set.

To have an idea of the upper bound we could obtain this way, we also evaluated by considering only the best predicted label, with respect to hierarchical F-score, and prefixed the corresponding measures with max-h.

6 Settings

6.1 Data

The PDTB (Prasad et al., 2007) is a corpus of English newswire, containing 2,159 articles from the Wall Street Journal. We use the section 23 as test set. In the following sections, we present results for both explicit and implicit examples. Contrary to existing studies, we give results for all the labels

annotated in the data (in particular, our results are not limited to level 1 or 2 relations). There are 41 distinct relation labels in the corpus, with 30 end-labels (mainly level 3 labels, but also level 2 labels that have no sub-types), and 11 “intermediate” labels (such as *Contingency.Cause* or *Comparison*).

6.2 Model architecture

We have separate classifiers for each dimension, and we compare the mapping from these to a full relation with a direct PDTB relation prediction.

Infersent is an architecture for sentence relation prediction, initially proposed to train transferable sentence representation from a semantic inference task to be fine-tuned on various sentence and sentence pair classification tasks. It takes as input two text fragments s_1 and s_2 (sentence or clause here), mapped to pretrained word embeddings (GloVe), encode each separately with a bi-LSTM with tied weights, and combine the final LSTM states to predict a relation. The combination is a concatenation of the representations provided for each argument, their absolute difference, and their element-wise product.

Each argument of the relation is thus encoded as a vector of dimension n , and the combined representation is a vector of dimension $4n$ for each separate relation dimension to predict, for various values of n .

6.3 Hyper-parameters

Models are trained for each dimension separately, with a maximum of 15 epochs and early stopping. An additional fully connected layer can be added on top of the combination of argument representations, and we vary the size of the layer with 0 (no layer), 512, or 4096 dimensions. We also tried different regularization values (weight decay): 10^{-n} , with $n \in \{-8, 1\}$. The best setting on the development set was chosen as our configuration for the final test.

7 Results

We describe here the performances obtained for our systems for each primitive separately, and use the reverse mapping to evaluate performance on relations as annotated in the PDTB.

7.1 Predicting primitives

All primitives are not equal in importance in the perspective of predicting rhetorical relations.

Primitive	Baseline			Best model		
	Acc	m-F ₁	w-F ₁	Acc	m-F ₁	w-F ₁
Basic op.	72.76	28.08	61.29	75.90	37.80	69.03
Polarity	73.00	28.13	61.60	82.29	49.86	80.59
Src of Coh.	52.67	23.00	36.34	68.06	50.03	67.44
Impl. order	73.05	21.11	61.68	78.16	41.00	74.89
Temp.	69.63	20.52	57.16	72.65	48.04	69.32
Cond.	95.88	–	–	98.55	–	–
Altern.	98.78	–	–	98.84	–	–
Specif.	82.93	–	–	85.13	–	–

Table 2: Scores of the systems for each primitive on test set (section 23 of the PDTB). The baseline is a majority classifier. We report Accuracy (“Acc”), and, for non-binary tasks, macro averaged F₁ (“m-F₁”) and weighted F₁ (“w-F₁”).

Some primitives, such as *basic operation* and *polarity*, correspond to major distinctions with respect to PDTB hierarchy: their values determine distinctions between top-level classes. Other primitives characterize more restricted sets of relations (*alternative*, *specificity*) or label distinctions at level 3 (*source of coherence*).

Table 2 shows performance for each primitive separately. We observe that among core primitives, *basic operation* demonstrates the least improvement (on accuracy, macro averaged F1 and weighted F1) with respect to the baseline, and thus should be a priority for further work. For primitive *polarity*, whose distribution of values are comparable (see Figure 1), results are quite better. When looking at the confusion matrix for this primitive, we observe that 95% of *positive* relations and 50% of *negative* relations are correctly labeled. For primitive *basic operation*, only 14% of *causal* relations are correctly labeled (relations are mainly labeled as *positive*). For primitive *temporal order*, results are lower than for primitive *polarity*. Relations are mainly labeled as *NS* (*non-specified*, which is the majority class) for this primitive.

The greatest improvement with respect to the baseline is for primitive *source of coherence*, but this result must be tempered by the fact that there are a very small number of *subjective* relations in our dataset (less than 1%).⁵ A further study with more data about *subjective* relations could be more informative.

⁵It should be noted that there is a potential loss of information due to the absence of a *subjective* version for *Contingency.Cause.Result* (whereas the *subjective* version of *Contingency.Cause.Reason* is *Contingency.Pragmatic cause.Justification*) in the PDTB 2.0 hierarchy (whereas present in PDTB 3).

We also looked at the difference when predicting primitives for implicit and explicit relations, and it appears there is almost no improvement on implicit over the baseline, which seems to confirm that primitives should not be considered in isolation. Less distinctive primitives show high accuracy mainly because they are unspecified most of the time.

7.2 Relation identification

Table 3 summarizes the scores obtained for relation identification, either when the relation label is obtained via the reverse mapping from the predicted primitives (row “Primitives”), or for systems directly trained to predict discourse relations (row “Relations”). We report accuracy as done in the literature by considering a prediction as correct if it contains one of the gold labels, and use hierarchical measures to have a more general setting. Again, our models generally outperform the baseline, often by a large margin, showing the relevance of Inference architecture to perform the task. Accuracy is much lower than predicting directly the relations, which can be explained by the fact that primitives are learned independently from each other.

By analyzing the predictions, we observed that *Contingency* relations were rarely predicted, a consequence of the aforementioned problem when predicting the *basic operation* primitive (which separates *causal* from *additive* relations). Another problem is that combining primitives still leaves too much underspecification, and predicting too many labels greatly impacts all hierarchical scores. We can also see that explicit relations benefit from the presence of very specific markers, while primitive recombination cannot make use of the marker information as efficiently. An encouraging aspect is that we found a lot of cases where a *Temporal* relation was predicted instead of a *Contingency* relation because the *basic operation* primitive was wrong, but the others were correct, which appears as plain error in all evaluations while being close to the ground truth. This seems to indicate primitive could be useful information on their own. Note that the scores we report in this table are the first, to the best of our knowledge, that are computed on the whole set of relations of the PDTB.

	Explicit					Implicit					All				
	Acc	h-R	h-P	max-h-R	max-h-P	Acc	h-R	h-P	max-h-R	max-h-P	Acc	h-R	h-P	max-h-R	max-h-P
	PDTB relations														
Baseline	23.5	25.35	26.13	27.02	27.33	15.73	30.5	34.72	31.38	35.5	20.03	27.65	29.97	28.97	30.98
Primitives	46.27	35.56	26.43	59.93	69.59	19.12	20.63	10.52	35.61	45.99	34.15	28.89	19.32	49.07	59.05
Relations	59.08	63.63	65.3	67.4	67.8	28.35	39.76	42.11	40.57	42.67	45.35	52.97	54.95	55.42	56.58

Table 3: Scores of the systems for relation prediction, using the full relation set of the PDTB. The predicted relations are either inferred from the predicted primitives (“Primitives”), or directly predicted (“Relations”). We report hierarchical recall (h-R) and hierarchical precision (h-P), along with max-h-P max-h-R, and accuracy.

8 Conclusion

We have taken a theoretical proposition for mapping discourse framework to apply it to discourse relation decomposition into primitives, in the context of the PDTB English corpus. This allows us to have a simple representation of PDTB annotations as a set of semantic and pragmatic primitives, allowing for general representations in case of underspecification. We have shown a simple experiment to learn these concepts separately and compare them to a direct relation classifier. Of course the primitives are not independent from each other, so learning them in isolation is bound to be less accurate than learning fully specified relation, but this framework lends itself straightforwardly to a multi-task learning setting and will be subject of future work. Other interesting perspectives include testing whether, when learning primitives on a training corpus without some relations, we can predict them correctly based on their conceptual decomposition (something akin to 0-shot learning); and finally, applying this decomposition to other discourse framework (RST or SDRT) can make cross-corpora training and prediction possible.

9 Acknowledgement

This work was supported partly by the french PIA project “Lorraine Université d’Excellence”, reference ANR-15-IDEX-04-LUE, and the PEPS blanc from CNRS (INS2I).

References

- Nicholas Asher and Alex Lascarides. 2003. *Logics of conversation*. Cambridge University Press.
- Hongxiao Bai and Hai Zhao. 2018. Deep enhanced representation for implicit discourse relation recognition. In *Proceedings of the 27th International Conference on Computational Linguistics*.
- Christian Chiarcos. 2014. [Towards interoperable discourse annotation. discourse features in the ontologies of linguistic annotation](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation, LREC 2014, Reykjavik, Iceland, May 26-31, 2014.*, pages 4569–4577.
- Vera Demberg, Fatemeh Torabi Asr, and Merel Scholman. 2017. [How consistent are our discourse annotations? Insights from mapping RST-DT and PDTB annotations](#). *CoRR*, abs/1704.08893.
- Eduard H. Hovy. 1990. [Parsimonious and profligate approaches to the question of discourse structure relations](#). In *Proceedings of the Fifth International Workshop on Natural Language Generation*.
- Yusuke Kido and Akiko Aizawa. 2016. Discourse relation sense classification with two-step classifiers. *Proceedings of the CoNLL-16 shared task*.
- Svetlana Kiritchenko, Stan Matwin, and A. Fazel Famili. 2005. Functional annotation of genes using hierarchical text categorization. In *in Proc. of the BioLINK SIG: Linking Literature, Information and Knowledge for Biology (held at ISMB-05)*.
- Alistair Knott. 1997. *A Data-Driven Methodology for Motivating a Set of Coherence Relations*. Ph.D. thesis, University of Edinburgh.
- William C. Mann and Sandra A. Thompson. 1988. Rhetorical Structure Theory: Toward a functional theory of text organization. *Text*, 8:243–281.
- Todor Mihaylov and Anette Frank. 2016. Discourse relation sense classification using cross-argument semantic similarity based on word embeddings. *Proceedings of the CoNLL-16 shared task*.
- Rashmi Prasad, Nikhil Dinesh, Alan Lee, Eleni Miltsakaki, Livio Robaldo, Aravind Joshi, and Bonnie Webber. 2008. [The Penn Discourse TreeBank 2.0](#). In *Proceedings of LREC*.
- Rashmi Prasad, Eleni Miltsakaki, Nikhil Dinesh, Alan Lee, Aravind Joshi, Livio Robaldo, and Bonnie Webber. 2007. [The Penn Discourse TreeBank 2.0 annotation manual](#).
- Ted Sanders, Wilbert Spooren, and Leo Noordman. 1992. [Toward a taxonomy of coherence relations](#). *Discourse Processes*, 15:1–35.
- Ted Sanders, Wilbert Spooren, and Leo Noordman. 1993. [Coherence relations in a cognitive theory](#).

of discourse representation. *Cognitive Linguistics*, 4:93–134.

Ted J.M. Sanders, Vera Demberg, Jet Hoek, Merel C.J. Scholman, Fatemeh Torabi Asr, Sandrine Zufferey, and Jacqueline Evers-Vermeul. 2018. [Unifying dimensions in coherence relations: How various annotation frameworks are related](#). *Corpus Linguistics and Linguistic Theory*.

Yannick Versley. 2011. Towards finer-grained tagging of discourse connectives. In *Proceedings of the Workshop Beyond Semantics: Corpus-based Investigations of Pragmatic and Discourse Phenomena*.

Nianwen Xue, Hwee Tou Ng, Sameer Pradhan, Attapol Rutherford, Bonnie Webber, Chuan Wang, and Hongmin Wang. 2016. Conll 2016 shared task on multilingual shallow discourse parsing. In *Proceedings of the CoNLL-16 shared task*.

Nianwen Xue, Hwee Tou Ng, Sameer Pradhan, Rashmi Prasad, Christopher Bryant, and Attapol Rutherford. 2015. [The conll-2015 shared task on shallow discourse parsing](#), pages 1–16.

Discourse Relation Prediction: Revisiting Word Pairs with Convolutional Networks

Siddharth Varia
Dept. of Computer Science
Columbia University
sv2504@columbia.edu

Chris Hidey
Dept. of Computer Science
Columbia University
ch3085@columbia.edu

Tuhin Chakrabarty
Dept. of Computer Science
Columbia University
tc2896@columbia.edu

Abstract

Word pairs across argument spans have been shown to be effective for predicting the discourse relation between them. We propose an approach to distill knowledge from word pairs for discourse relation classification with convolutional neural networks by incorporating joint learning of implicit and explicit relations. Our novel approach of representing the input as word pairs achieves state-of-the-art results on four-way classification of both implicit and explicit relations as well as one of the binary classification tasks. For explicit relation prediction, we achieve around 20% error reduction on the four-way task. At the same time, compared to a two-layered Bi-LSTM-CRF model, our model is able to achieve these results with half the number of learnable parameters and approximately half the amount of training time.

1 Introduction

Implicit discourse relation identification is the task of recognizing the relationship between text segments without the use of an explicit connective indicating the relationship. For instance, while a connective such as “because” may indicate a causal relationship when present between sentences, it is not necessary for causality (as in Example 1). Without the explicit connective, automatically identifying the relationship is much more difficult. Improvement in identifying implicit discourse relations will also improve performance in downstream tasks such as question answering, textual inference (for determining relationships between text segments), machine translation and other multi-lingual tasks (for transferring discourse information between languages).

The Penn Discourse Tree Bank (PDTB) theory of discourse relations (Prasad et al., 2008) defines a shallow discourse representation between adjacent or nearby segments. As a result, the span of

the arguments participating in the discourse relation is often the most important input to a classifier.

Initial approaches used linguistically informed features derived from the arguments as inputs to traditional machine learning methods (Pitler et al., 2008). More recently, the application of neural methods has resulted in the best performance on this task, modeling the relationship between words in the arguments in context (Ji et al., 2015; Dai and Huang, 2018).

A common approach in prior work is to use pairs of words from across the arguments as features (Marcu and Echihiabi, 2002; Blair-Goldensohn et al., 2007; Pitler et al., 2009). Consider the example:

I am **late** for the meeting because the
train was **delayed**. (1)

The words “late” and “delayed” are semantically related and (absent the connective) one might hypothesize that their presence is what triggers a causal relation. Therefore, pairs of words across discourse arguments should be useful features for identifying discourse relations. However, learning these specific word pairings requires leveraging large text corpora to observe them in the relevant discourse context (Biran and McKeown, 2013). Furthermore, as the number of possible word pairs grows quadratically with the size of the vocabulary, representing word pairs discretely results in very sparse feature sets. Since a continuous representation of the word pairs allows for better generalization to unseen pairs, we thus use a Convolutional Neural Network (CNN) to embed word pairs from the arguments in a dense vector representation. We also extend this idea of word pairs beyond a single pair of words by using larger filter sizes.

Our results show that these word pairs provide

improved performance in transferring knowledge from explicit relations, indicating less sensitivity to word ordering. Finally, an additional advantage is that our architecture based on convolution layers allows for additional improvement in the speed of training through parallel processing unlike sequential models based on LSTMs.

Our primary contributions are as follows:

- A novel application of convolutional neural networks to model word pairs in the arguments in a discourse relation
- A demonstration that joint learning of implicit and explicit relations with both word pairs and n-grams improves performance over learning implicit relations only
- State-of-the-art results on four-way classification for both implicit and explicit relations, reducing the error by 20% in the latter case
- A model with half the number of learnable parameters compared to a state-of-the-art two-layered Bi-LSTM-CRF model along with approximately half the training time

2 Related Work

Previous work on discourse relations found success using word pairs as features. In the earliest work using word pairs, [Marcu and Echihabı \(2002\)](#) used unambiguous explicit markers such as “but” to create a corpus of discourse relations. They used a Naive Bayes approach by taking the cross-product of words on either side of the connective. [Blair-Goldensohn et al. \(2007\)](#) used word pairs for discourse relations as well. Later work ([Pitler et al., 2009](#)) applied this approach to the PDTB but found that the top word pairs were discourse connectives, which is counter-intuitive as connectives were removed to obtain word pairs. These earlier approaches use word pairs directly as features, which creates a large sparse feature space. In more recent work, [Biran and McKeown \(2013\)](#) address the sparsity issue by using features based on word pairs in the context of an explicit connective in the Gigaword corpus. Even though this approach addresses the sparsity issue by using a much larger corpus, it is still impractical to account for *every* possible word pair. In comparison to these previous word pair methods, our model takes advantage of the continuous representation

of word embeddings to model similarity between word pairs.

In other work, researchers have found that approaches using neural networks have helped increase performance on this task, as neural models are better at dealing with sparsity. Some work has focused on using novel representations. [Ji et al. \(2015\)](#) model the arguments with recursive neural networks (modeling the tree structure of each argument). [Lei et al. \(2017\)](#) model interaction between words in arguments by learning linear and quadratic relations. [Liu and Li \(2016\)](#) develop a method for repeated reading over the discourse context by using an external memory. Other researchers have found success by modeling the sequence of words using recurrent neural networks ([Chen et al., 2016](#)) with a gating mechanism to combine contextual word pairs while some approaches have used convolutional neural networks over each argument ([Qin et al., 2016](#)). Most recent work has focused on joint learning, as the PDTB is a relatively small dataset for neural methods. [Liu et al. \(2016\)](#) and [Lan et al. \(2017\)](#) propose a multi-task learning approach across PDTB and other corpora. [Qin et al. \(2017\)](#) have demonstrated the effectiveness of an adversarial approach, forcing one model without connective information to be similar to a model with connective. [Rönnqvist et al. \(2017\)](#) developed the first attention-based BiLSTM network for Chinese implicit discourse relations. [Dai and Huang \(2018\)](#) show that incorporating additional document context at the paragraph level and jointly predict both implicit and explicit relations. Finally, [Bai and Zhao \(2018\)](#) propose a deep model using contextual ELMo embeddings, multiple CNN layers and Bi-attention. Unlike these approaches, we represent the input in a novel way as a set of word pairs, while using a much simpler architecture, and distill knowledge between explicit and implicit relations.

3 Methods

Our architecture consists of two primary components. The first component learns complex interactions from word pairs and the second component learns n-gram features from individual arguments. The features from word pair convolutions and individual argument convolutions are then combined using a gating mechanism. Finally, we jointly learn representations for implicit and explicit relations.

	the	train	was	delayed
late	late, the	late, train	late, was	late, delayed
for	for, the	for, train	for, was	for, delayed
the	the, the	the, train	the, was	the, delayed
meeting	meeting, the	meeting, train	meeting, was	meeting, delayed

Table 1: Arg_1 is along the rows, Arg_2 is across the columns. Cell (i, j) corresponds to the word pair composed from i^{th} word of Arg_1 and j^{th} word of Arg_2 .

3.1 Product of Arguments

For the first component, we use convolution operations over the Cartesian product of words from the two arguments.

Word/Word Pairs Initially, we consider the interaction between all pairs of individual words in the arguments. Table 1 illustrates the use of word pairs from Example 1 in Section 1, where $Arg_1 =$ “I am late for the meeting” and $Arg_2 =$ “the train was delayed.” The sequence of word pairs starts at the first row and moves on to successive rows. In this example, we remove the connective to illustrate an implicit relation. For explicit relations, we include the connective as part of the second argument and create word pairs for the connective (e.g. “because”) as well. When computing the Cartesian product, we drop very short ($length < 3$) functional words to limit the number of word pairs (hence the absence of the words “I” and “am”).

Word/N-Gram Pairs We also use larger filters to capture relations between word pairs. A filter of size $2k$ will capture k word pairs. Henceforth we use the notation $WP-k$ to indicate a sequence of k word pairs (where $WP-1$ refers to a single pair of words). We use the following notation to describe concrete examples of word pair features: (“late” : “delayed”) is an example of $WP-1$. Similarly (“late” : “the train was delayed”) is an example of $WP-4$ and corresponds to the following sequence of word pairs (as in row 1 of Table 1): “late the late train late was late delayed”. In other words, it corresponds to the Cartesian product of “late” with the 4-gram “the train was delayed.” Thus, another interpretation of $WP-k$ is a mapping of a word in Arg_1 to a k -gram in Arg_2 and vice-versa. This interpretation of $WP-k$ is not true at word transitions. For instance, in the above example, a filter of size 8 will also capture “late was late delayed for the for train” as one of the $WP-4$. By learning $WP-k$ features (for $k > 1$) we are able to capture more complex interaction between the arguments. This

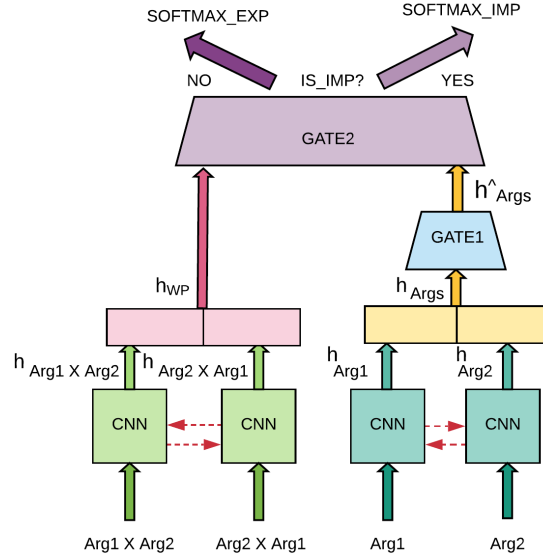


Figure 1: Architecture of our proposed network. Dashed arrows indicate that weights are shared among CNNs

is novel in comparison to the common practice of just using $WP-1$ as features. We use filters of even length and a stride of two to ensure the filter will always end at word pair boundaries.

Mathematically, the input to the CNN (where $[\cdot]$ means concatenation) is:

$$\mathbf{v}_{Arg_1 \times Arg_2} = [x_1^1 \cdot x_2^1 \cdot x_1^2 \cdot x_2^2 \cdot x_1^3 \cdot x_2^3 \cdot \dots]$$

where \mathbf{x}_1^i is the concatenation of word and POS embeddings corresponding to the i^{th} word of Arg_1 and \mathbf{x}_2^j is the concatenation of word and POS embeddings corresponding to the j^{th} word of Arg_2 . We also include the representation obtained from $(Arg_2 \times Arg_1)$, as our preliminary experiments showed that this approach was complementary to the representation from $(Arg_1 \times Arg_2)$:

$$\mathbf{v}_{Arg_2 \times Arg_1} = [x_2^1 \cdot x_1^1 \cdot x_2^2 \cdot x_1^2 \cdot x_2^3 \cdot x_1^3 \cdot \dots]$$

Following convolution, in line with the common practice, we apply max pooling along the length

of the sequence to pick the most prominent feature per feature map. Next we concatenate the max-pooled features from different filter sizes to obtain a hidden representation. This hidden representation from the CNN has dimensionality equal to the number of feature maps \times the number of filters. Thus, for each discourse relation comprised of Arg_1 and Arg_2 , we obtain two hidden representations $\mathbf{h}_{Arg_1 \times Arg_2}$ and $\mathbf{h}_{Arg_2 \times Arg_1}$. We concatenate these representations to obtain a vector \mathbf{h}_{WP} . The weights of the convolution layers are shared between $Arg_1 \times Arg_2$ and $Arg_2 \times Arg_1$ to allow the model to learn from both types of interaction. The left side of the Figure 1 depicts this component of our combined architecture.

3.2 Individual Arguments

For the second component, we use a CNN over *individual* arguments Arg_1 and Arg_2 (illustrated on the right side of Figure 1). As discussed in Rönqvist et al. (2017), arguments provided without context may contain elements indicative of a discourse relation (Asr and Demberg, 2015), e.g. implicit causality verbs (Rohde and Horton, 2010). We thus hypothesize that the hidden representation obtained from individual arguments will complement the representation obtained from the word pairs. To learn representations from the individual arguments we use filters of odd and even length and stride equal to one.

As with word pairs, the weights of the convolution layers are shared between Arg_1 and Arg_2 to allow the model to learn representations from both sides independent of the order of the arguments.

3.3 Combination of Argument Representations

In order to combine the representations \mathbf{h}_{Arg_1} and \mathbf{h}_{Arg_2} , we incorporate a method for the model to learn to weight the interaction between the argument features. We employ $Gate_1$ as shown on the right side of Figure 1. This gate is defined as follows:

$$\begin{aligned} c &= Relu(W_1 \cdot \mathbf{h}_{Args} + b_1) \\ g_a &= \sigma(W_2 \cdot \mathbf{h}_{Args} + b_2) \\ \hat{\mathbf{h}}_{Args} &= c \odot g_a \end{aligned} \quad (2)$$

where \mathbf{h}_{Args} is the concatenation of \mathbf{h}_{Arg_1} and \mathbf{h}_{Arg_2} .

We subsequently join the word pair representations and the individual argument representations

with a similar mechanism. The two components are combined using $Gate_2$ which is defined as in Equation 2 but instead takes as input the concatenation of $\hat{\mathbf{h}}_{Args}$ and \mathbf{h}_{WP}

To predict the discourse relation, the output of $Gate_2$ is then input to a separate dense layer with softmax non-linearity for either explicit or implicit relation classification as shown in Figure 1.

3.4 Joint Learning of Implicit and Explicit Relations

Finally, to fully take advantage of the labeled data in the PDTB, we jointly learn implicit and explicit relations. For explicit relations, we add the connective to the beginning of Arg_2 . As shown in Figure 1 and similar to (Dai and Huang, 2018), we use separate classification layers for explicit and implicit relations. To jointly learn both types of relations, we randomize the order of implicit or explicit relations rather than training each mini-batch separately.

4 Experiments

4.1 Data

We run experiments on three different tasks (binary, four-way and fifteen-way). For binary and four-way tasks, we train and test on the class level relations defined in the PDTB: Comparison, Contingency, Temporal, and Expansion. We use a common partition of the data: sections 2-20 for training, 0-1 for validation, and 21-22 for testing. For this partition, there are 1046 implicit relation instances and 1285 explicit relation instances in the test set. For fifteen-way task, we precisely follow the setup of CoNLL 2016 shared task on shallow discourse parsing and evaluate our approach on their test and blind test sets. A small fraction (around 4%) of the relations in PDTB have multiple gold labels. During training, we replicate a relation once for each gold label and for evaluation we deem the prediction to be correct if the predicted label matches any of the gold labels. We use this scheme for all the tasks in this paper.

4.2 Experimental setup

We use Spacy to tokenize and annotate POS tags for the individual arguments. To learn $WP-k$ features, we limit the Cartesian product to a maximum of 500 word pairs per relation. For n-gram features, we limit the arguments to a maximum of 100 words. For word embeddings, we used

Model		Implicit					
		Macro F1	Acc	Com	Con	Exp	Tem
LSTM	(Lei et al., 2017)	46.46	-	-	-	-	-
	(Lan et al., 2017)	47.80	57.39	-	-	-	-
	(Dai and Huang, 2018)	- (48.82)	- (58.20)	- (37.72)	- (49.39)	- (68.86)	- (40.70)
CNN	(Liu et al., 2016)	44.98	57.27	-	-	-	-
	(Bai and Zhao, 2018)	51.06	-	-	-	-	-
Ours	WP-[1-4], Args, Implicit Only	50.77 (49.2)	59.46 (56.11)	45.82 (42.1)	54.39 (51.1)	70.48 (64.77)	43.04 (38.8)
	Args Joint Learning	49.47 (48.1)	59.66 (57.50)	42.68 (35.50)	54.82 (52.5)	70.30 (67.07)	41.82 (37.47)
	WP-1, Args, Joint Learning	50.71 (48.73)	59.18 (57.36)	45.91 (37.33)	55.87 (52.27)	69.04 (66.61)	42.96 (38.70)
	WP-[1-4], Args, Joint Learning	51.84 (50.2)	60.52 (59.13)	46.84 (41.94)	53.74 (49.81)	72.42 (69.27)	43.97 (39.77)

Table 2: Results of four-way classification experiments on implicit relations. The numbers in the parenthesis correspond to average of 10 runs

Model	Explicit	
	Macro F1	Acc
(Dai and Huang, 2018)	- (93.70)	- (94.46)
Args, JL	95.48 (94.81)	96.2 (95.63)
WP-1, Args, JL	95.13 (94.83)	95.95 (95.67)
WP-[1-4], Args, JL	95.0 (94.50)	95.72 (95.33)

Table 3: Results of four-way classification experiments on explicit relations. JL : Joint Learning

word2vec pre-trained embeddings. However, for words not found in word2vec, we back-off to embeddings trained on the raw WSJ articles. We fix the word embeddings during training. We also concatenate one hot POS embeddings to the fixed word embeddings. We use 100 and 50 feature maps per filter size for learning $WP-k$ and n-grams respectively. For $WP-k$, we use filters of size 2, 4, 6 and 8. For n-grams, we use filters of size 2, 3, 4 and 5. For all dense layers and gate layers, we set the output dimension of the weight matrices to 300. For regularization, we use dropout (Srivastava et al., 2014) of 0.5 after convolution operations and before the softmax layers. We also use L2 regularization with a coefficient of 0.0001 and

early stopping to prevent over-fitting. For training, we minimize multi-class cross-entropy loss using the Adam optimizer (Kingma and Ba, 2014) with a learning rate of 0.0005 and batch size of 200. Our architecture is implemented in Theano deep learning framework.¹

5 Results

We compare our results to previous work along two dimensions: the architecture of the model (CNN or LSTM) and whether the model employs a joint learning component.

Our work and the work of (Dai and Huang, 2018) and (Lan et al., 2017) involves jointly training on explicit relations. (Lan et al., 2017) and (Liu et al., 2016) also train on *BLLIP* and *RST*, respectively.

5.1 Four-way classification

Tables 2 & 3 show the results of four-way experiments on implicit and explicit relations respectively.² Please note that these results are from the same joint learning experiments wherever applicable and they are presented in different tables for the sake of better presentation. We compare the performance of our models under different configurations. We gradually add $WP-k$ features to study their contribution. First we add $WP-1$ (fil-

¹<https://github.com/siddharthvaria/WordPair-CNN>

²(Qin et al., 2017) reported only one-versus-all binary classification so their results are only included in Table 4

Model		Com	Con	Exp	Tem
LSTM	(Lei et al., 2017)	40.47	55.36	69.50	35.34
	(Lan et al., 2017)	40.73	58.96	72.47	38.50
	(Dai and Huang, 2018)	- (46.79)	- (57.09)	- (70.41)	- (45.61)
CNN	(Liu et al., 2016)	37.91	55.88	69.97	37.17
	(Qin et al., 2017)	40.87	54.56	72.38	36.20
	(Bai and Zhao, 2018)	47.85	54.47	70.60	36.87
Ours	WP-[1-4], Args Joint Learning	45.03 (44.1)	56.53 (56.02)	73.5 (72.11)	46.15 (44.41)

Table 4: Results of binary classification experiments. The numbers in the parenthesis correspond to average of 10 runs

ters of size 2) and then we add $WP-k$ features to illustrate the contribution of more complex interactions for $k > 1$. Additionally, we compare joint learning of implicit and explicit relations (*Joint Learning*) against learning implicit relations only (*Implicit only*). *Args* refers to n-grams from individual arguments. For all experiments, we report both the maximum and average (in parenthesis) of 10 runs for fair comparison with all prior work. It is not surprising to see that gradually adding word pair features improves performance on implicit relations. When using joint learning and $WP-[1-4]$ we obtain an improvement in Macro F1-Score and Accuracy for implicit relations over previous state of the art works (Dai and Huang, 2018) and (Bai and Zhao, 2018). We also observe improvement for the expansion class in the joint learning setting, likely due to its prevalence in both implicit and explicit relations. In these cases, we observe that joint learning improves over training on just implicit relations (with a 3 point improvement in overall accuracy primarily due to a 5 point improvement in classification of the expansion class). On the other hand, we find that in just *Args* setting, we obtain state-of-the-art performance compared to prior work (Dai and Huang, 2018) for explicit relation F1-Score and accuracy, achieving 20% reduction in error rate. We don't get any benefit by combining it with word pairs (for $WP-1$), and the extra complexity for $k > 1$ makes it more difficult for the model to distinguish the effective features. This may occur because the connective itself is a very strong baseline.

5.2 Four-way Ensemble results

As the experiments described in 4.2 were conducted with 10 random initializations, we also

present the results of an ensemble created out of these 10 runs via majority voting in Table 5. Compared to (Dai and Huang, 2018), while our ensemble achieves marginal improvement of 0.59 F1 on implicit relations, it improves by around 1 point on explicit relations for both metrics, around a 20% error reduction.

5.3 Binary classification

In Table 4, we report our results on four binary classification tasks. From the results, we see that our model does better on all classes in comparison to other CNN-based architectures. Our averaged results are directly comparable to those of (Dai and Huang, 2018) and we observe improvement for the expansion class. Our model may not generalize as well on other three classes because they account for 14, 26.4, and 6% of the test set, respectively, leading to high variance across multiple runs.

5.4 Fifteen-way classification

CoNLL organized a multilingual shallow discourse parsing shared task in 2016. In this shared task (Xue et al., 2016), they consider second level types and release test and blind test sets for fifteen-way classification of explicit and implicit relations, including EntRel and AltLex relations as implicit relations. We compare our architecture against the systems that participated in that task, with results presented in Table 6. Our architecture produces very similar results in line with the results reported by various neural network based systems that participated in the task. However, we also observe that using word pair features does not lead to further improvement over using just n-gram features. One possible explanation for this

Model	Implicit		Explicit	
	Macro F1	Acc	Macro F1	Acc
Dai et al. (2018)	51.84	59.85	94.17	94.82
WP-[1-4], Args, IO	51.63	58.03	-	-
Args, JL	49.54	58.70	94.81	95.64
WP-1, Args, JL	51.90	59.94	95.16	95.95
WP-[1-4], Args, JL	52.53	61.28	94.38	95.25

Table 5: Ensemble results of four-way classification experiments. JL : “Joint Learning” and IO : “Implicit Only”

Model	F1 score			
	Implicit		Explicit	
	PDTB	Blind	PDTB	Blind
Xue et al., (2016)	40.91	37.67	90.22	78.56
Lan et al., (2017)	39.40	40.12	-	-
Args, JL	39.68	38.74	89.91	76.98
WP-[1-4], Args, JL	39.39	39.36	89.48	77.00

Table 6: Results of fifteen-way task on CoNLL 2016 test and blind test sets

trend is that word pair features capture enough semantic information to discriminate the top-level classes however it fails to separate the second level of types. Comparing against (Lan et al., 2017), we see that our model is competitive with their LSTM-based architecture in spite of the fact that they used external data to achieve these results. This also possibly indicates that it is hard to get further improvements on this task without data augmentation due to lack of enough training data for second level types in PDTB.

6 Discussion

6.1 Comparison of Model Complexity

In Table 7 we present the number of parameters of our model in the first two columns. We have con-

Model	Parameters
Ours	
Conv 2,3,4,5 50 per size	242.2k
Conv 2,4,6,8 100 per size	692k
Gate ₁	240k
Gate ₂	660k
Total	1834.2k
LSTM Model	
Bi-LSTM Layer 1	1550.4k
Bi-LSTM Layer 2	2160k
Total	3710.4k

Table 7: Comparison of model complexity. Gate₁ & Gate₂ have output of size 300. LSTMs have hidden state of size 300.

volution layers to learn n-gram features and WP-k features. Apart from these layers, we have two gate layers: Gate₁ and Gate₂ in the table. Our model has approximately 1.8 million parameters. The input embeddings to our model have dimensionality of 346 (300 (word) + 46 (POS)). Assuming the same input to the two-layered Bi-LSTM model with a hidden state of size 300, this model will have approximately 3.7 million parameters. For this comparison, we have assumed the number of parameters of the LSTM given input vectors of size m and giving output vectors of size n is $4(nm + n^2)$. Both models have dense layers for implicit and explicit relation prediction so they are ignored for these calculations.

6.2 Comparison of Training Time

We also compare the running time of our model to the model of (Dai and Huang, 2018). We compare the wall clock training time per epoch of both systems, using their released code as well as our own. For a fair comparison, we re-implemented our architecture in Pytorch to match their usage. Furthermore, the models were run on the same GPU (Tesla K80) on the same machine. We ran each model three times for five epochs. The training time of our model was 109.6 seconds on average compared to 206.17 seconds for their model.

6.3 Qualitative Analysis

We conduct a qualitative analysis in an attempt to understand the most important WP-k and n-grams learned by our architecture. We modified our architecture to get rid of all non-linear layers after the convolutional layers, which allows us to examine the effect of the word pairs and n-grams

Implicit Relations and Top Features
<p>Arg1: Alliant said it plans to use the microprocessor in future products Arg2: It declined to discuss its plans for upgrading its current product line Class: Comparison WP-k: (said : declined), (Alliant : product line), (declined : Alliant said), (upgrading : microprocessor future products), (plans : declined discuss its plans), (discuss : use the microprocessor future) Arg1 n-grams: (Alliant said), (microprocessor in future products), (plans to use the microprocessor) Arg2 n-grams: (product line), (It declined to discuss), (for upgrading its current product)</p>
<p>Arg1: I ca n't see why there would be a conflict of interest Arg2: Estimates are based on the previous price of similar works sold at auction and current market conditions, and are not affected by any knowledge of who the potential buyer could be Class: Contingency WP-k: (n't : affected), (not : conflict), (why : not affected), (not : n't see), (see : works sold auction and), (affected : why there would conflict) Arg1 n-grams: (a conflict of interest), (ca n't see why there) Arg2 n-grams: (Estimates are based on), (works sold at auction), (are not affected by any)</p>
<p>Arg1: And it allows Mr. Van de Kamp to get around campaign spending limits Arg2: He can spend the legal maximum for his campaign Class: Expansion WP-k: (And : can), (limits : spend), (allows : spend), (And : He can), (maximum : spending limits) Arg1 n-grams: (And it allows), (spending limits) Arg2 n-grams: (He can spend), (legal maximum), (his campaign)</p>
<p>Arg1: As the market dropped Friday , Robertson Stephens slashed the value of the offering by 7% Arg2: Yesterday , when similar securities rebounded , it bumped the valuation up again Class: Temporal WP-k: (As : when), (bumped : slashed), (As : Yesterday when), (Yesterday : As the market), (when : As the market dropped) Arg1 n-grams: (market dropped), (Robertson Stephens slashed the value) Arg2 n-grams: (similar securities), (Yesterday , when), (bumped the valuation up again)</p>
<p>Arg1: the fact that seven patents were infringed suggests that infringement was willful Arg2: It 's difficult to be that consistently wrong Class: Contingency WP-k: (willful : consistently), (willful : wrong), (suggests : difficult that consistently wrong), (consistently : infringed suggests that infringement) Arg1 n-grams: (the fact), (suggests that infringement was willful) Arg2 n-grams: (consistently wrong), ('s difficult to be that)</p>
<p>Arg1: and special consultants are springing up to exploit the new tool Arg2: Blair Entertainment has just formed a subsidiary – 900 Blair – to apply the technology to television Class: Expansion WP-k: (springing : formed), (exploit : formed), (springing : subsidiary Blair), (formed : springing exploit), (springing : has just formed subsidiary), (Blair : and special consultants) Arg1 n-grams: (special consultants are springing up) Arg2 n-grams: (Blair Entertainment has), (subsidiary – 900 Blair –)</p>

Table 8: Implicit examples along with top features selected from across three runs. Note that we drop very short words in the cartesian product only and not in the individual arguments.

directly on the output. Dropping the gate layers caused the F1 score averaged across the first three runs to drop from 50.9 to 50.1, indicating both that the gate layers help incorporate interactions between the model components and that our approach here is a reasonable approximation to what the model is learning. Instead of the gates, we concatenate the output of all the convolutional layers and use a final classification layer (different for implicit and explicit relations as in our full model) to train this simplified architecture. In the absence of non-linearity, we are able to map the features selected by max pooling back to the $WP-k$ and n -grams associated with their embeddings (i.e. argmax pooling rather than max pooling and mapping the selected indices back to the input). As each filter is associated with multiple feature maps ($k = 100$ for word pairs and $k = 50$ for n -grams as described in Section 4.2), we count the number of times each $WP-k$ and n -gram was selected during pooling and select the most prominent features according to their frequency.

We present six implicit examples,³ in Table 8 and the corresponding top $WP-k$ and n -gram features. We selected these examples by running the simplified architecture three times and selecting implicit examples which were classified correctly during all the runs in the *Joint Learning* setting.

We find the following general properties in the examples we studied:

- We consistently observed that smaller filters learn either verb-to-verb mappings or adjective-to-adjective mappings. In examples one, four and five, (said : declined), (bumped : slashed) and (willful : wrong) are selected respectively. The first two pairs capture antonymy and last one maps adjectives.
- Larger filters tend to align important words (verbs and nouns) in either argument to phrases in the other argument. In the third example, (maximum: spending limits) is selected, along with (affected : why there would conflict) in the second example, and (plans : declined discuss its plans) from the first, among others.
- For the third example, while the true class is *Expansion*, which the *Joint Learning* model classifies correctly, the *Implicit Only* model

³Although we learn implicit and explicit relations jointly, we focus on only implicit relations due to space constraints.

labels it as *Contingency*. The *Joint Learning* model selects functional word pair interactions such as (And : can), which may be more indicative of an Expansion relation due to the presence of the connective “And” at the start of the first argument. We also observe the word pair (Kamp : spend) is not selected as a top feature in the *Joint Learning* setting, while it is selected in the *Implicit Only* scenario. As it includes a proper noun, it is unlikely to generalize as a useful feature. Finally, *Joint Learning* identifies the semantically coherent $WP-2$ (maximum : spending limits). This pair does not appear in the *Implicit Only* case.

7 Conclusion

We proposed an approach to learn implicit relations by incorporating word pair features as a novel way to capture the interaction between the arguments, a distinct approach compared to the popular attention-based approaches used with Bi-LSTM based models. We also show that joint learning of implicit and explicit relations is beneficial to implicit relations. Our results show that our model is able to surpass or match the performance of a Bi-LSTM based model using paragraph level context.

For future work, we plan to explore data augmentation techniques. As our best performance is on the expansion class, which is also the largest, if we are able to obtain more data we might improve our performance on smaller classes as well. We will thus investigate extending our joint learning method to include resources beyond the PDTB.

Another possible avenue is to replace word embeddings with contextualized embeddings to study the efficacy of the latter with our architecture. Pre-trained language models like BERT (Jacob et al., 2018) have been recently used to achieve state-of-the-art results on sentence pair classification tasks. As a future step we will experiment with our model on top of these contextual representations, which would likely enhance the performance while still maintaining the interpretability.

Acknowledgments

We would like to thank Kathleen McKeown, Dimitris Alikaniotis and anonymous reviewers for their constructive feedback

References

- Fatemeh Torabi Asr and Vera Demberg. 2015. Uniform information density at the level of discourse relations: Negation markers and discourse connective omission. In *IWCS 2015*, page 118.
- Hongxiao Bai and Hai Zhao. 2018. Deep enhanced representation for implicit discourse relation recognition. *CoRR*, abs/1807.05154.
- Or Biran and Kathleen McKeown. 2013. Aggregated word pair features for implicit discourse relation disambiguation. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 69–73, Sofia, Bulgaria. Association for Computational Linguistics.
- Sasha Blair-Goldensohn, Kathleen McKeown, and Owen Rambow. 2007. Building and refining rhetorical-semantic relation models. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*, pages 428–435, Rochester, New York. Association for Computational Linguistics.
- Jifan Chen, Qi Zhang, Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2016. Implicit discourse relation detection via a deep architecture with gated relevance network. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1726–1735, Berlin, Germany. Association for Computational Linguistics.
- Zeyu Dai and Ruihong Huang. 2018. Improving implicit discourse relation classification by modeling inter-dependencies of discourse units in a paragraph. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 141–151. Association for Computational Linguistics.
- Devlin Jacob, Chang Ming-Wei, Lee Kenton, and Toutanova Kristina. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Arxiv*. Google Research.
- Yangfeng Ji, Gongbo Zhang, and Jacob Eisenstein. 2015. Closing the gap: Domain adaptation from explicit to implicit discourse relations. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2219–2224, Lisbon, Portugal. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Man Lan, Jianxiang Wang, Yuanbin Wu, Zheng-Yu Niu, and Haifeng Wang. 2017. Multi-task attention-based neural networks for implicit discourse relationship representation and identification. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1299–1308. Association for Computational Linguistics.
- Wenqiang Lei, Xuancong Wang, Meichun Liu, Ilija Ilievski, Xiangnan He, and Min-Yen Kan. 2017. Swim: A simple word interaction model for implicit discourse relation recognition. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 4026–4032.
- Yang Liu and Sujian Li. 2016. Recognizing implicit discourse relations via repeated reading: Neural networks with multi-level attention. *CoRR*, abs/1609.06380.
- Yang Liu, Sujian Li, Xiaodong Zhang, and Zhifang Sui. 2016. Implicit discourse relation classification via multi-task neural networks. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI’16*, pages 2750–2756. AAAI Press.
- Daniel Marcu and Abdessamad Echihabi. 2002. An unsupervised approach to recognizing discourse relations. In *Proceedings of 40th Annual Meeting of the Association for Computational Linguistics*, pages 368–375, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Emily Pitler, Annie Louis, and Ani Nenkova. 2009. Automatic sense prediction for implicit discourse relations in text. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 683–691, Suntec, Singapore. Association for Computational Linguistics.
- Emily Pitler, Mridhula Raghupathy, Hena Mehta, Ani Nenkova, Alan Lee, and Aravind Joshi. 2008. Easily identifiable discourse relations. In *Coling 2008: Companion volume: Posters*, pages 87–90, Manchester, UK. Coling 2008 Organizing Committee.
- Rashmi Prasad, Nikhil Dinesh, Alan Lee, Eleni Miltasakaki, Livio Robaldo, Aravind Joshi, and Bonnie Webber. 2008. The penn discourse treebank 2.0. In *Proceedings of LREC*.
- Lianhui Qin, Zhisong Zhang, and Hai Zhao. 2016. A stacking gated neural architecture for implicit discourse relation classification. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2263–2270. Association for Computational Linguistics.
- Lianhui Qin, Zhisong Zhang, Hai Zhao, Zhiting Hu, and Eric Xing. 2017. Adversarial connective-exploiting networks for implicit discourse relation classification. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1006–1017, Vancouver, Canada. Association for Computational Linguistics.

- Hannah Rohde and William S. Horton. 2010. Why or what next? eye movements reveal expectations about discourse direction.
- Samuel Rönnqvist, Niko Schenk, and Christian Chiarcos. 2017. [A recurrent neural model with attention for the recognition of Chinese implicit discourse relations](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 256–262, Vancouver, Canada. Association for Computational Linguistics.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: A simple way to prevent neural networks from overfitting](#). *J. Mach. Learn. Res.*, 15(1):1929–1958.
- Nianwen Xue, Hwee Tou Ng, Sameer Pradhan, Atapol Rutherford, Bonnie Webber, Chuan Wang, and Hongmin Wang. 2016. [CoNLL 2016 shared task on multilingual shallow discourse parsing](#). In *Proceedings of the CoNLL-16 shared task*, pages 1–19, Berlin, Germany. Association for Computational Linguistics.

Author Index

- Abrams, Mitchell, 418
Acharya, Sabita, 232
Agarwal, Sanchit, 264
Akhtiamov, Oleg, 274
Alloatti, Francesca, 250
Anikina, Tatiana, 399
Ardati, Amer, 232
asadullah, munshi, 110
Attari, Nazia, 331
Axelsson, Nils, 345
- Balog, Krisztian, 353
Bastianelli, Emanuele, 254
Bellec, Pascal, 110
Besset, Benoit, 110
Bigham, Jeffrey, 379
Black, Alan W, 367
Bouraoui, Jean-Leon, 114
Boyd, Andrew, 232
Boye, Johan, 422
Braud, Chloé, 432
Breitfeller, Luke, 239
Budzianowski, Paweł, 155
Byrne, Bill, 353
- Cameron, Richard, 232
Carbou, Romain, 114
Cercas Curry, Amanda, 361
Chakrabarty, Tuhin, 442
Chao, Guan-Lin, 122, 215
Chattopadhyay, Debaleena, 232
Chen, Chun-Yen, 293
Chen, Jindong, 215
Choi, Jinho D., 188
Chung, Tagyoung, 264
Cohen, Philip, 198
Cohn, Michelle, 293
Conway, Rylan, 56
- Damnati, Geraldine, 110
Di Caro, Luigi, 250
Di Eugenio, Barbara, 232
Dickens, Carolyn, 232
Dossantos, Martinho, 110
Dunn Lopez, Karen, 232
- Dušek, Ondřej, 392
- Eshghi, Arash, 32
Eskenazi, Maxine, 165, 379
- Galitsky, Boris, 118
Gao, Shuyang, 264
Gasic, Milica, 155
Gessler, Luke, 418
Ginzburg, Jonathan, 320
Goncharova, Elizaveta, 118
Gupta, Arshit, 46
Gupta, Prakhar, 379
- Habibi, Pantea, 232
Hakkani-Tur, Dilek, 122, 215, 264
He, He, 367
Heckmann, Martin, 210, 331
Heinecke, Johannes, 110
Heldner, Mattias, 284
Herledan, Frederic, 110
Hewitt, John, 46
Hidey, Christopher, 442
Huang, Xuedong, 40
Huang, Ziming, 133
- Ilvovsky, Dmitry, 118
- Jiang, Zhuoxuan, 133
- Kalpakchi, Dmytro, 422
Karpov, Alexey, 274
Keizer, Simon, 392
Kelleher, John D., 75
Kiefer, Bernd, 411
Kirchhoff, Katrin, 46
Kleingarn, Diana, 210
Kolossa, Dorothea, 210
Kotti, Margarita, 336
Krishnamoorthi, Karthik, 353
Kruijff-Korbayova, Ivana, 399, 411
Kurfali, Murathan, 226
- Lambert, Mathias, 56
Lancien, Jeanyves., 110
Lane, Ian, 215

Laskowski, Kornel, 284
Le Meitour, Sonia, 114
Leblouch, Olivier, 110
Lee, Sungjin, 32, 143
Lemaire, Vincent, 114
Lemon, Oliver, 32, 254
Li, Chuyuan, 320
Li, Shaochun, 133
Liden, Lars, 143
Liu, Bing, 21, 178
Liu, Xingkun, 392
Louvan, Samuel, 85
Łupkowski, Paweł, 320
Lykartsis, Athanasios, 336

Ma, Jie, 133
Ma, Nianzu, 21
Magnini, Bernardo, 85
Mao, Xian-Ling, 133
Marge, Matthew, 418
Martyn-Nemeth, Pamela, 232
Mazumder, Sahisnu, 21
McLeod, Sarah, 411
Mehri, Shikib, 165, 379
Minker, Wolfgang, 274
Molino, Piero, 92, 178
Mory, Emmanuel, 110
Mulholland, Matthew, 103
Muller, Philippe, 432

Nabizadeh, Nima, 210
Naik, Aakanksha, 239
Namazifar, Mahdi, 178
Neubig, Graham, 68

Östling, Robert, 226

Papangelis, Alexandros, 92
Paranjape, Bhargavi, 68
Partovi, Andisheh, 307
Pavel, Amy, 379

Qian, Yao, 103

Radlinski, Filip, 353
Ramanarayanan, Vikram, 103
Rastogi, Abhinav, 122, 215
Ravi, Sujith, 1
Ren, Kexin, 320
Rieser, Verena, 361, 392
Rojas Barahona, Lina M., 110, 114
Rose, Carolyn, 239
Ross, Robert J., 75

Roze, Charlotte, 432

Sankar, Chinnadhurai, 1
Schlangen, David, 331
Sethi, Abhishek, 264
Shalyminov, Igor, 32
Shu, Lei, 178
Siegert, Ingo, 274
Skantze, Gabriel, 345
Sportelli, Gianpiero, 250
Srinivasan, Tejas, 165

Trinh, Anh Duong, 75
Tseng, Bo-Hsiang, 155
Tsvetkov, Yulia, 367
Tur, Gokhan, 92, 178

Ultes, Stefan, 11

Vanzo, Andrea, 254
Varia, Siddharth, 442
Vatani, Haleh, 232
Väth, Dirk, 62
Vu, Ngoc Thang, 62

Wang, Shuai, 21
Wang, Yi-Chia, 92
Wlodarczak, Marcin, 284
Wu, Yen-chen, 155

Xu, Hu, 178
Xu, Xinnuo, 143

Yang, Zhengzhe, 188
Yavuz, Semih, 122, 215
Yu, Zhou, 293
Yusupujang, Zulipiye, 320

Zeng, Michael, 40
Zhang, Yizhe, 143
Zhao, Tiancheng, 379
Zheng, Huaixiu, 178
Zhou, Yiheng, 367
Zhu, Chenguang, 40
Zukerman, Ingrid, 307