

ACL 2020

NLP for Conversational AI

Proceedings of the 2nd Workshop

July 9, 2020

©2020 The Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

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

978-1-952148-08-8

Introduction

Welcome to the ACL 2020 Workshop on **NLP for Conversational AI**.

Ever since the invention of the intelligent machine, hundreds and thousands of mathematicians, linguists, and computer scientists have dedicated their career to empowering human-machine communication in natural language. Although the idea is finally around the corner with a proliferation of virtual personal assistants such as Siri, Alexa, Google Assistant, and Cortana, the development of these conversational agents remains difficult and there still remain plenty of unanswered questions and challenges.

Conversational AI is hard because it is an interdisciplinary subject. Initiatives were started in different research communities, from Dialogue State Tracking Challenges to NIPS Conversational Intelligence Challenge live competition and the Amazon Alexa prize. However, various fields within the NLP community, such as semantic parsing, coreference resolution, sentiment analysis, question answering, and machine reading comprehension etc. have been seldom evaluated or applied in the context of conversational AI.

The goal of this workshop is to bring together NLP researchers and practitioners in different fields, alongside experts in speech and machine learning, to discuss the current state-of-the-art and new approaches, to share insights and challenges, to bridge the gap between academic research and real-world product deployment, and to shed the light on future directions. “NLP for Conversational AI” will be a one-day workshop including keynotes, spotlight talks, posters, and panel sessions. In keynote talks, senior technical leaders from industry and academia will share insights on the latest developments of the field. An open call for papers will be announced to encourage researchers and students to share their prospects and latest discoveries. The panel discussion will focus on the challenges, future directions of conversational AI research, bridging the gap in research and industrial practice, as well as audience-suggested topics.

With the increasing trend of conversational AI, NLP4ConvAI 2020 is competitive. We received 27 submissions, and after a rigorous review process, we only accept 15. There are total 13 accepted regular workshop papers and 2 cross-submissions or extended abstracts. The workshop overall acceptance rate is about 55.5%. We hope you will enjoy NLP4ConvAI 2020 at ACL and contribute to the future success of our community!

NLPConvAI 2020 Organizers

Tsung-Hsien Wen, PolyAI

Asli Celikyilmaz, Microsoft

Zhou Yu, UC Davis

Alexandros Papangelis, Uber AI

Mihail Eric, Amazon Alexa AI

Anuj Kumar, Facebook

Iñigo Casanueva, PolyAI

Rushin Shah, Google

Organizers:

Tsung-Hsien Wen, PolyAI (UK)
Asli Celikyilmaz, Microsoft (USA)
Zhou Yu, UC Davis (USA)
Alexandros Papangelis, Uber AI (USA)
Mihail Eric, Amazon Alexa AI (USA)
Anuj Kumar, Facebook (USA)
Iñigo Casanueva, PolyAI (UK)
Rushin Shah, Google (USA)

Program Committee:

Pawel Budzianowski, University of Cambridge and PolyAI
Sam Coope, PolyAI
Ondřej Dušek, Heriot Watt University
Nina Dethlefs, University of Hull
Daniela Gerz, PolyAI
Pei-Hao Su, PolyAI
Matthew Henderson, PolyAI
Simon Keizer, AI Lab, Vrije Universiteit Brussel
Ryan Lowe, McGill University
Julien Perez, Naver Labs
Marek Rei, Imperial
Gokhan Tür, Uber AI
Ivan Vulic, University of Cambridge and PolyAI
David Vandyke, Apple
Zhuoran Wang, Trio.io
Tiancheng Zhao, CMU
Alborz Geramifard, Facebook
Shane Moon, Facebook
Jinfeng Rao, Facebook
Bing Liu, Facebook
Pararth Shah, Facebook
Ahmed Mohamed, Facebook
Vikram Ramanarayanan, ETS
Rahul Goel, Google
Shachi Paul, Google
Seokhwan Kim, Amazon
Andrea Madotto, HKUST
Dian Yu, UC Davis
Piero Molino, Uber
Chandra Khatri, Uber
Yi-Chia Wang, Uber
Huaixiu Zheng, Uber
Fei Tao, Uber
Abhinav Rastogi, Google
Stefan Ultes, Daimler

José Lopes, Heriot Watt University
Behnam Hedayatnia, Amazon
Dilek Hakkani-Tür, Amazon
Ta-Chung Chi, CMU
Shang-Yu Su, NTU
Pierre Lison, NR
Ethan Selfridge, Interactions
Teruhisa Misu, HRI
Svetlana Stoyanchev, Toshiba
Ryuichiro Higashinaka, NTT
Hendrik Bushmeier, University of Bielefeld
Kai Yu, Baidu
Koichiro Yoshino, NAIST
Abigail See, Stanford University
Ming Sun, Amazon Alexa AI

Invited Speakers:

Yun-Nung Chen, National Taiwan University
Dilek Hakkani-Tür, Amazon Alexa AI
Jesse Thomason, University of Washington
Antoine Bordes, Facebook AI Research
Jacob Andreas, Massachusetts Institute of Technology
Jason Williams, Apple

Table of Contents

<i>Using Alternate Representations of Text for Natural Language Understanding</i> Venkat Varada, Charith Peris, Yangsook Park and Christopher Dipersio	1
<i>On Incorporating Structural Information to improve Dialogue Response Generation</i> Nikita Moghe, Priyesh Vijayan, Balaraman Ravindran and Mitesh M. Khapra	11
<i>CopyBERT: A Unified Approach to Question Generation with Self-Attention</i> Stalin Varanasi, Saadullah Amin and Guenter Neumann	25
<i>How to Tame Your Data: Data Augmentation for Dialog State Tracking</i> Adam Summerville, Jordan Hashemi, James Ryan and william ferguson	32
<i>Efficient Intent Detection with Dual Sentence Encoders</i> Iñigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson and Ivan Vulić	38
<i>Accelerating Natural Language Understanding in Task-Oriented Dialog</i> Ojas Ahuja and Shrey Desai	46
<i>DLGNet: A Transformer-based Model for Dialogue Response Generation</i> Olabiyi Oluwatobi and Erik Mueller	54
<i>Data Augmentation for Training Dialog Models Robust to Speech Recognition Errors</i> Longshaokan Wang, Maryam Fazel-Zarandi, Aditya Tiwari, Spyros Matsoukas and Lazaros Polymenakos	63
<i>Automating Template Creation for Ranking-Based Dialogue Models</i> Jingxiang Chen, Heba Elfardy, Simi Wang, Andrea Kahn and Jared Kramer	71
<i>From Machine Reading Comprehension to Dialogue State Tracking: Bridging the Gap</i> Shuyang Gao, Sanchit Agarwal, Di Jin, Tagyoung Chung and Dilek Hakkani-Tur	79
<i>Improving Slot Filling by Utilizing Contextual Information</i> Amir Pouran Ben Veyseh, Franck Dernoncourt and Thien Huu Nguyen	90
<i>Learning to Classify Intents and Slot Labels Given a Handful of Examples</i> Jason Krone, Yi Zhang and Mona Diab	96
<i>MultiWOZ 2.2 : A Dialogue Dataset with Additional Annotation Corrections and State Tracking Baselines</i> Xiaoxue Zang, Abhinav Rastogi and Jindong Chen	109
<i>Sketch-Fill-A-R: A Persona-Grounded Chit-Chat Generation Framework</i> Michael Shum, Stephan Zheng, Wojciech Kryscinski, Caiming Xiong and Richard Socher	118
<i>Probing Neural Dialog Models for Conversational Understanding</i> Abdelrhman Saleh, Tovly Deutsch, Stephen Casper, Yonatan Belinkov and Stuart Shieber	132

Workshop Program

July 9, 2020

- 06:00** *Opening Remarks*
- 06:10 *Invited Talk*
Yun-Nung Chen
- 06:40 *Invited Talk*
Antonie Bordes
- 07:10 *Invited Talk*
Jacob Andreas
- 07:40 *Using Alternate Representations of Text for Natural Language Understanding*
Venkat Varada, Charith Peris, Yangsook Park and Christopher Dipersio
- 07:50 *On Incorporating Structural Information to improve Dialogue Response Generation*
Nikita Moghe, Priyesh Vijayan, Balaraman Ravindran and Mitesh M. Khapra
- 08:00 *CopyBERT: A Unified Approach to Question Generation with Self-Attention*
Stalin Varanasi, Saadullah Amin and Guenter Neumann
- 08:10 *How to Tame Your Data: Data Augmentation for Dialog State Tracking*
Adam Summerville, Jordan Hashemi, James Ryan and william ferguson
- 08:20 *Efficient Intent Detection with Dual Sentence Encoders*
Iñigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson and Ivan Vulić
- 08:30 *Accelerating Natural Language Understanding in Task-Oriented Dialog*
Ojas Ahuja and Shrey Desai
- 08:40 *DLGNet: A Transformer-based Model for Dialogue Response Generation*
Olabiyi Oluwatobi and Erik Mueller
- 08:50 *Data Augmentation for Training Dialog Models Robust to Speech Recognition Errors*
Longshaokan Wang, Maryam Fazel-Zarandi, Aditya Tiwari, Spyros Matsoukas and Lazaros Polymenakos

July 9, 2020 (continued)

- 09:00** **Panel**
- 10:00 *Invited Talk*
Jesse Thomason
- 10:30 *Invited Talk*
Dilek Hakkani-Tür
- 11:00 *Invited Talk*
Jason Williams
- 11:30 *Automating Template Creation for Ranking-Based Dialogue Models*
Jingxiang Chen, Heba Elfardy, Simi Wang, Andrea Kahn and Jared Kramer
- 11:40 *From Machine Reading Comprehension to Dialogue State Tracking: Bridging the Gap*
Shuyang Gao, Sanchit Agarwal, Di Jin, Tagyoung Chung and Dilek Hakkani-Tur
- 11:50 *Improving Slot Filling by Utilizing Contextual Information*
Amir Pouran Ben Veyseh, Franck Dernoncourt and Thien Huu Nguyen
- 12:00 *Learning to Classify Intents and Slot Labels Given a Handful of Examples*
Jason Krone, Yi Zhang and Mona Diab
- 12:10 *MultiWOZ 2.2 : A Dialogue Dataset with Additional Annotation Corrections and State Tracking Baselines*
Xiaoxue Zang, Abhinav Rastogi and Jindong Chen
- 12:20 *Sketch-Fill-A-R: A Persona-Grounded Chit-Chat Generation Framework*
Michael Shum, Stephan Zheng, Wojciech Kryscinski, Caiming Xiong and Richard Socher
- 12:30 *Probing Neural Dialog Models for Conversational Understanding*
Abdelrhman Saleh, Tovly Deutsch, Stephen Casper, Yonatan Belinkov and Stuart Shieber
- 12:40** ***Closing Remarks***