

ACL 2019

NLP for Conversational AI

Proceedings of the 1st Workshop

August 1, 2019
Florence, Italy

©2019 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

ISBN 978-1-950737-37-6

Introduction

Welcome to the ACL 2019 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 2019 is competitive. We received 68 submissions, and after a rigorous review process, we only accept 25. There are total 16 accepted regular workshop papers and 7 cross-submissions or extended abstracts. The workshop overall acceptance rate is about 36.8%.

We hope you will enjoy NLP4ConvAI 2019 at ACL and contribute to the future success of our community!

NLPConvAI 2019 Organizers

Tania Bedrax-Weiss, Google AI

Yun-Nung (Vivian) Chen, National Taiwan University

Dilek Hakkani-Tur, Amazon Alexa

Anuj Kumar, Facebook

Mike Lewis, Facebook AI

Thang-Minh Luong, Google Brain

Pei-Hao (Eddy) Su, PolyAI

Tsung-Hsien (Shawn) Wen, PolyAI

Organizers:

Tania Bedrax-Weiss, Google AI
Yun-Nung Chen, National Taiwan University
Dilek Hakkani-Tur, Amazon Alexa
Anuj Kumar, Facebook
Mike Lewis, Facebook AI
Thang-Minh Luong, Google Brain
Pei-Hao Su, PolyAI
Tsung-Hsien Wen, PolyAI

Program Committee:

Abhinav Arora, Anusha Balakrishnan, Pawel Budzianowski, Inigo Casanueva, Wenhua Chen, Ta-Chung Chi, Sam Coope, Nina Dethlefs, Mihail Eric, Ashish Garg, Daniela Gerz, Raghav Gupta, Adithya Gurram, Matt Henderson, Minlie Huang, Ting-Hao Huang, Simon Keizer, Shang-Wen Li, Bing Liu, Yi Luan, Gaurav Menghani, Mrinal Mohit, Seungwhan Moon, Jekaterina Novikova, Alexandros Papangelis, Julien Perez, Elahe Rahimtoroghi, Deepak Ramachandran, Abhinav Rastogi, Marek Rei, Lina M Rojas, Chinnadhurai Sankar, Pararth Shah, Georgios Spithourakis, Shang-Yu Su, Ming Sun, Trieu H Trinh, Gokhan Tur, Stefan Ultes, David Vandyke, Ivan Vendrov, Ivan Vulic, Peng Wang, Yi-Chia Wang, Zhuoran Wang, Jason Williams, Chien-Sheng Wu, Zi Yang, Yi-Ting Yeh, Tiancheng Zhao

Meta-Reviewers:

Tania Bedrax-Weiss, Google AI
Anuj Kumar, Facebook
Pei-Hao Su, PolyAI
Ivan Vulić, PolyAI
Tsung-Hsien Wen, PolyAI

Invited Speaker:

Yejin Choi, University of Washington
Jianfeng Gao, Microsoft Research AI
Matt Henderson, PolyAI
Verena Rieser, Heriot-Watt University
Ruhi Sarikaya, Amazon Alexa
Jason Weston, Facebook

Table of Contents

<i>A Repository of Conversational Datasets</i>	
Matthew Henderson, Paweł Budzianowski, Iñigo Casanueva, Sam Coope, Daniela Gerz, Girish Kumar, Nikola Mrkšić, Georgios Spithourakis, Pei-Hao Su, Ivan Vulić and Tsung-Hsien Wen	1
<i>A Simple but Effective Method to Incorporate Multi-turn Context with BERT for Conversational Machine Comprehension</i>	
Yasuhito Ohsugi, Itsumi Saito, Kyosuke Nishida, Hisako Asano and Junji Tomita	11
<i>Augmenting Neural Response Generation with Context-Aware Topical Attention</i>	
Nouha Dziri, Ehsan Kamaloo, Kory Mathewson and Osmar Zaiane	18
<i>Building a Production Model for Retrieval-Based Chatbots</i>	
Kyle Swanson, Lili Yu, Christopher Fox, Jeremy Wohlwend and Tao Lei	32
<i>Co-Operation as an Asymmetric Form of Human-Computer Creativity. Case: Peace Machine</i>	
Mika Härmäläinen and Timo Honkela	42
<i>Conversational Response Re-ranking Based on Event Causality and Role Factored Tensor Event Embedding</i>	
Shohei Tanaka, Koichiro Yoshino, Katsuhito Sudoh and Satoshi Nakamura	51
<i>DSTC7 Task 1: Noetic End-to-End Response Selection</i>	
Chulaka Gunasekara, Jonathan K. Kummerfeld, Lazaros Polymenakos and Walter Lasecki	60
<i>End-to-End Neural Context Reconstruction in Chinese Dialogue</i>	
Wei Yang, Rui Qiao, Haocheng Qin, Amy Sun, Luchen Tan, Kun Xiong and Ming Li	68
<i>Energy-Based Modelling for Dialogue State Tracking</i>	
Anh Duong Trinh, Robert Ross and John Kelleher	77
<i>Evaluation and Improvement of Chatbot Text Classification Data Quality Using Plausible Negative Examples</i>	
Kit Kuksenok and Andriy Martyniv	87
<i>Improving Long Distance Slot Carryover in Spoken Dialogue Systems</i>	
Tongfei Chen, Chetan Naik, Hua He, Pushpendre Rastogi and Lambert Mathias	96
<i>Insights from Building an Open-Ended Conversational Agent</i>	
Khyatti Gupta, Meghana Joshi, Ankush Chatterjee, Sonam Damani, Kedhar Nath Narahari and Puneet Agrawal	106
<i>Learning to Explain: Answering Why-Questions via Rephrasing</i>	
Allen Nie, Erin Bennett and Noah Goodman	113
<i>Multi-turn Dialogue Response Generation in an Adversarial Learning Framework</i>	
Oluwatobi Olabiyi, Alan O Salimov, Anish Khazane and Erik Mueller	121
<i>Relevant and Informative Response Generation using Pointwise Mutual Information</i>	
Junya Takayama and Yuki Arase	133
<i>Responsive and Self-Expressive Dialogue Generation</i>	
Kozo Chikai, Junya Takayama and Yuki Arase	139

Conference Program

Thursday, August 1, 2019

09:00–10:30 Morning Session I

09:00–09:15 *Welcome and Opening Remarks*

09:15–10:00 *Invited Talk 1: Should Conversational AI Use Neural Response Generation?*
Verena Rieser (Heriot-Watt University)

10:00–10:30 *Lightning Talks*

10:30–11:00 Coffee Break and Poster Session I

11:00–12:30 Morning Session II

11:00–11:45 *Invited Talk 2*
Matt Henderson (PolyAI)

11:45–12:30 *Invited Talk 3: The Design and Implementation of XiaoIce, An Empathetic Social Chatbot*
Jianfeng Gao (Microsoft Research AI)

12:30–14:00 Lunchtime and Poster Session II

Thursday, August 1, 2019 (continued)

14:00–15:30 Afternoon Session I

14:00–14:45 *Invited Talk 4*
Yejin Choi (University of Washington)

14:45–15:30 *Invited Talk 5: Putting Together the Threads of Conversational AI?*
Jason Weston (Facebook)

15:30–16:00 Coffee Break and Poster Session III

16:00–18:00 Afternoon Session II

16:00–16:45 *Invited Talk 6: Natural Self-Learning Contextual Conversational Systems*
Ruhi Sarikaya (Amazon Alexa)

16:45–17:45 Panel Discussion

17:45–18:00 Closing Remarks

Regular Workshop Papers

A Repository of Conversational Datasets

Matthew Henderson, Paweł Budzianowski, Iñigo Casanueva, Sam Coope, Daniela Gerz, Girish Kumar, Nikola Mrkšić, Georgios Spithourakis, Pei-Hao Su, Ivan Vulić and Tsung-Hsien Wen

A Simple but Effective Method to Incorporate Multi-turn Context with BERT for Conversational Machine Comprehension

Yasuhito Ohsugi, Itsumi Saito, Kyosuke Nishida, Hisako Asano and Junji Tomita

Augmenting Neural Response Generation with Context-Aware Topical Attention

Nouha Dziri, Ehsan Kamalloo, Kory Mathewson and Osmar Zaiane

Building a Production Model for Retrieval-Based Chatbots

Kyle Swanson, Lili Yu, Christopher Fox, Jeremy Wohlwend and Tao Lei

Thursday, August 1, 2019 (continued)

Co-Operation as an Asymmetric Form of Human-Computer Creativity. Case: Peace Machine

Mika Hämmäläinen and Timo Honkela

Conversational Response Re-ranking Based on Event Causality and Role Factored Tensor Event Embedding

Shohei Tanaka, Koichiro Yoshino, Katsuhito Sudoh and Satoshi Nakamura

DSTC7 Task 1: Noetic End-to-End Response Selection

Chulaka Gunasekara, Jonathan K. Kummerfeld, Lazaros Polymenakos and Walter Lasecki

End-to-End Neural Context Reconstruction in Chinese Dialogue

Wei Yang, Rui Qiao, Haocheng Qin, Amy Sun, Luchen Tan, Kun Xiong and Ming Li

Energy-Based Modelling for Dialogue State Tracking

Anh Duong Trinh, Robert Ross and John Kelleher

Evaluation and Improvement of Chatbot Text Classification Data Quality Using Plausible Negative Examples

Kit Kuksenok and Andriy Martyniv

Improving Long Distance Slot Carryover in Spoken Dialogue Systems

Tongfei Chen, Chetan Naik, Hua He, Pushpendre Rastogi and Lambert Mathias

Insights from Building an Open-Ended Conversational Agent

Khyatti Gupta, Meghana Joshi, Ankush Chatterjee, Sonam Damani, Kedhar Nath Narahari and Puneet Agrawal

Learning to Explain: Answering Why-Questions via Rephrasing

Allen Nie, Erin Bennett and Noah Goodman

Multi-turn Dialogue Response Generation in an Adversarial Learning Framework

Oluwatobi Olabiyi, Alan O Salimov, Anish Khazane and Erik Mueller

Relevant and Informative Response Generation using Pointwise Mutual Information

Junya Takayama and Yuki Arase

Responsive and Self-Expressive Dialogue Generation

Kozo Chikai, Junya Takayama and Yuki Arase

Thursday, August 1, 2019 (continued)

Cross-Submissions / Extended Abstracts

Adapting NLG methods to social robotics

Simon Meoni

Do Neural Dialog Systems Use the Conversation History Effectively? An Empirical Study

Chinnadhurai Sankar, Sandeep Subramanian, Chris J Pal, Sarath Chandar and Yoshua Bengio

FlowDelta: Modeling Flow Information Gain in Reasoning for Conversational Machine Comprehension

Yi-Ting Yeh and Yun-Nung Chen

Getting To Know You: Extracting User Attributes from Conversations for Personalized Dialogue Agents

Chien-Sheng Wu, Andrea Madotto, Zhaojiang Lin, Peng Xu and Pascale Fung

Improving Neural Conversational Models with Entropy-Based Data Filtering

Richárd Csáky, Patrik Purgai and Gábor Recski

Joint Dual Learning for Language Understanding and Generation

Shang-Yu Su, Chao-Wei Huang and Yun-Nung Chen

OpenDialKG: Explainable Conversational Reasoning with Attention-based Walks over Knowledge Graphs

Seungwhan Moon, Pararth Shah, Anuj Kumar and Rajen Subba

Personalizing Dialogue Agents via Meta-Learning

Zhaojiang Lin, Andrea Madotto, Chien-Sheng Wu and Pascale Fung

Transferable Multi-Domain State Generator for Task-Oriented Dialogue Systems

Chien-Sheng Wu, Andrea Madotto, Ehsan Hosseini-Asl, Caiming Xiong, Richard Socher and Pascale Fung