NLP4ConvAI 2023

The 5th Workshop on NLP for Conversational AI

Proceedings of the Workshop

July 14, 2023

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 978-1-959429-97-5

Introduction

We are excited to welcome you to NLP4ConvAI 2023, the 5th Annual Workshop on NLP for Conversational AI, co-located with ACL 2023 at Toronto, Canada.

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 in conversational AI, and to shed light on future directions. Following the success of the four previous editions of NLP for Conversational AI workshops at ACL & EMNLP, NLP4ConvAI 2023 is a one-day workshop including keynotes, oral presentations and posters.

We received 53 submissions this year, consisting of 38 long papers and 15 short papers. We had a total of 54 program committee (PC) members. At least three PC members reviewed each of the papers. We accepted 20 papers: 15 long papers and 5 short papers. These numbers give an overall acceptance rate of 38%, with the long and short papers acceptance rate being 39% and 33% respectively. Out of the 20 accepted papers, six are being presented as oral presentations and the remaining in a poster session. We have also identified one best paper (Generating Video Game Scripts with Style) and two outstanding papers (On the Underspecification of Situations in Open-domain Conversational Datasets, and Conversational Recommendation as Retrieval: A Simple, Strong Baseline).

In addition, the workshop program consists of five invited talks given by leading practitioners in industry and academia. We thank our five keynote speakers, Diyi Yang (Stanford University), Larry Heck (Georgia Institute of Technology), Vipul Raheja (Grammarly), Nurul Lubis (Heinrich Heine University Düsseldorf) and Jason Weston (Meta AI) for their inspiring, informative and thought provoking talks. We would also like to thank all the authors for submitting their work at the workshop, the program committee members for diligently reviewing the submissions and giving valuable feedback to the authors, and the ACL organizing committee for supporting us throughout the process.

We hope you will enjoy NLP4ConvAI 2023 at ACL and contribute to the future success of our community!

NLP4ConvAI 2023 Organizers
Abhinav Rastogi, General Chair
Georgios Spithourakis, Program Chair
Yun-Nung (Vivian) Chen and Bing Liu, Publication chairs
Yu Li, Diversity & Publicity Chair
Elnaz Nouri, Sponsorship Chair
Alon Albalak, Shared Task Chair
Alexandros Papangelis, Advisory Board

Organizing Committee

General Chair

Abhinav Rastogi, Google Research

Program Chair

Georgios Spithourakis, ex-PolyAI, Entrepreneur First

Publication Chairs

Yun-Nung Chen, National Taiwan University
Bing Liu, Meta

Diversity and Publicity Chair

Yu Li, Columbia University

Sponsorship Chair

Elnaz Nouri, Microsoft Research

Shared Task Chair

Alon Albalak, University of California, Santa Barbara

Advisory Board

Alexandros Papangelis, Amazon Alexa AI

Program Committee

Yun-Nung Chen, National Taiwan University
Bing Liu, Meta
Elnaz Nouri
Alexandros Papangelis, Amazon
Abhinav Rastogi, Google
Georgios P. Spithourakis, ex-PolyAI, Entrepreneur First
Yu Li, Columbia University
Alon Albalak, University of California, Santa Barbara
Abhinav Arora, Meta
Kartikeya Badola, Google Research
Mukul Bhutani, Carnegie Mellon University, Apple, Google
Alexandra Birch, University of Edinburgh
Jie Cao, University of Colorado
Yuan Cao, Google Brain
Guan-Lin Chao, Microsoft
Maximillian Chen, Columbia University
Nina Dethlefs, University of Hull
Ashwinkumar Ganesan, Amazon Alexa AI
Shubham Garg, Amazon.com
Christian Geishauser, Heinrich Heine University Duesseldorf
Alborz Geramifard, Facebook AI
Liane Guillou, The University of Edinburgh
Raghav Gupta, Google Inc.
Shachi H Kumar, Intel Labs
Dilek Hakkani-Tur, Amazon Alexa AI
Michael Heck, Heinrich Heine University
David M. Howcroft, Edinburgh Napier University
Songbo Hu, University of Cambridge
Chao-Wei Huang, National Taiwan University
Rishabh Joshi, Google
Mihir Kale, Google
Saarthak Khanna, Amazon
Seokhwan Kim, Amazon Alexa AI
Stefan Larson, Vanderbilt University
Hsien-chin Lin, Heinrich Heine University
Liangchen Luo, Google
Wolfgang Maier, Mercedes-Benz AG
Sneha Mehta, Bloomberg L.P.
Udita Patel, Amazon.com
Siddhesh Pawar, Google
Baolin Peng, Tencent AI Lab
Wei Peng, Huawei Technologies
Shiva Kumar Pentyala, Salesforce AI
Samrat Phatale, Google Research
Evgeniia Razumovskaia, University of Cambridge
Lina M. Rojas Barahona, Orange Innovation Research
Igor Shalyminov, Amazon AWS

Akshat Shrivastava, Facebook
Shubham Shukla, Independent Researcher
Kai Sun, Meta
Anh Duong Trinh, National College of Ireland
Miroslav Tushev, Amazon
Stefan Ultes, University of Bamberg
David Vandyke, Apple
Nikolas Vitsakis, University of Edinburgh, Heriott-Watt University
Peidong Wang, Microsoft
Peratham Wiriyathammabhum, Self
Chien-Sheng Wu, Salesforce
Hongyuan Zhan, Meta AI
Jianguo Zhang, Salesforce Research
Lukas Zilka, Google

Keynote Talk: Inclusive Conversational AI for Positive Impact

Diyi Yang
Stanford University
2023-07-14 09:10:00 – Room: **Harbour B**

Abstract: Conversational AI has revolutionized the way we interact with technology, holding the potential to create positive impact on a variety of domains. In this talk, we present two studies that develops inclusive conversational AI techniques to empower users in different contexts for social impact. The first one looks at linguistic prejudice with a participatory design approach to develop dialect-inclusive language tools for low-resourced dialects in conversational question answering, together with efficient adaptation of models trained on Standard American English (SAE) to different dialects. The second work introduces CARE, an interactive conversational agent that supports peer counselors by generating personalized suggestions. CARE diagnoses suitable counseling strategies and provides tailored example responses during training, empowering counselors to respond effectively. These works showcase the potential of how inclusive language technologies can address language and communication barriers and foster positive impact.

Bio: Diyi Yang is an assistant professor in the Computer Science Department at Stanford University. Her research goal is to understand the social aspects of language and build socially responsible NLP systems for social impact. Her work has received multiple best paper nominations or awards at top NLP and HCI conferences (e.g., ACL, EMNLP, SIGCHI, and CSCW). She is a recipient of IEEE AI 10 to Watch (2020), the Intel Rising Star Faculty Award (2021), the Samsung AI Researcher of the Year (2021), the Microsoft Research Faculty Fellowship (2021), and the NSF CAREER Award (2022).

Keynote Talk: Build it for One @ Right Place Right Time: Leveraging Context in Conversational Systems

Larry Heck
Georgia Institute of Technology
2023-07-14 09:40:00 – Room: **Harbour B**

Abstract: Recent years have seen significant advances in conversational systems, particularly with the advent of attention-based language models pre-trained on large datasets of unlabeled natural language text. While the breadth of the models has led to fluid and coherent dialogues over a broad range of topics, they can make mistakes when high precision is required. High precision is not only required when specialized skills are involved (legal/medical/tax advice, computations, etc.), but also to avoid seemingly trivial mistakes such as commonsense and other relevant 'in-the-moment' context. Much of this context centers on and should be derived from the user's perspective. This talk will explore prior and current work on leveraging this user-centric context (build it for one) and the user's specific situation (right place right time) to improve the accuracy and utility of conversational systems.

Bio: Larry Heck is a Professor in ECE and Interactive Computing, co-Executive Director of the AI Hub, Farmer Chair of Advanced Computing Concepts, and a GRA Eminent Scholar at Georgia Tech. He is a Fellow of the IEEE, inducted into the Academy of Distinguished Engineers at Georgia Tech, and named a Distinguished Engineer at Texas Tech. After receiving the PhD EE from Georgia Tech, he joined SRI, followed by VP of Research at Nuance, VP of Search and Advertising at Yahoo!, Chief Speech Scientist and Distinguished Engineer at Microsoft, Principal Scientist with Google Research, and CEO of Viv Labs and SVP at Samsung.

Keynote Talk: Building Better Writing Assistants In the Era of Conversational LLMs

Vipul Raheja
Grammarly
2023-07-14 13:30:00 – Room: **Harbour B**

Abstract: Text revision is a complex, iterative process. It is no surprise that human writers are unable to simultaneously comprehend multiple demands and constraints of the task of text revision when producing well-written texts, as they are required to cover the content, follow linguistic norms, set the right tone, follow discourse conventions, etc. This presents a massive challenge and opportunity for intelligent writing assistants, which have undergone an enormous shift in their abilities in the past few years and months via large language models. In addition to the quality of editing suggestions, writing assistance has undergone a monumental shift in terms of being a one-sided, push-based paradigm, to now being a natural language-based, conversational exchange of input and feedback. However, writing assistants still lack in terms of their quality, personalization, and overall usability, limiting the value they provide to users. In this talk, I will present my research, challenges, and insights on building intelligent and interactive writing assistants for effective communication, navigating challenges pertaining to quality, personalization, and usability.

Bio: Vipul Raheja is an Applied Research Scientist at Grammarly. He works on developing robust and scalable approaches centered around improving the quality of written communication, leveraging Natural Language Processing and Machine Learning. His research interests lie at the intersection of large language models and controllable text generation for writing assistance. He also co-organizes the Workshop on Intelligent and Interactive Writing Assistants (In2Writing). He received his Masters in Computer Science from Columbia University and in the past, worked at IBM Research, and x.ai on building conversational scheduling assistants.

Keynote Talk: Dialogue Evaluation via Offline Reinforcement Learning and Emotion Prediction

Nurul Lubis
Heinrich Heine University Düsseldorf
2023-07-14 15:00:00 – Room: **Harbour B**

Abstract: Task-oriented dialogue systems aim to fulfill user goals, such as booking hotels or searching for restaurants, through natural language interactions. They are ideally evaluated through interaction with human users. However, this is unattainable to do at every iteration of the development phase due to time and financial constraints. Therefore, researchers resort to static evaluation on dialogue corpora. Although they are more practical and easily reproducible, they do not fully reflect real performance of dialogue systems. Can we devise an evaluation that keeps the best of both worlds? In this talk I explore the usage of offline reinforcement learning and emotion prediction for dialogue evaluation that is practical, reliable, and strongly correlated with human judgements.

Bio: Nurul Lubis received the B.Eng. degree (cum laude) in 2014 from Bandung Institute of Technology, Bandung, Indonesia and the M.Eng. and Dr.Eng. degrees in 2017 and 2019, respectively, from Nara Institute of Science and Technology (NAIST), Nara, Japan. She received the NAIST Best Student Award in 2019. She was a recipient of the Japanese Ministry of Education, Culture, Sports, Science and Technology (MEXT) scholarship from 2014 to 2019. She was a research intern at Honda Research Institute Japan, Co. Ltd., Saitama, Japan and is currently a postdoctoral researcher at the Heinrich Heine University Düsseldorf, Düsseldorf, Germany. Her research interests include emotion in spoken language, affective dialogue systems, and dialogue policy optimization with reinforcement learning and variational methods.

Keynote Talk: Improving Open Language Models by Learning from Organic Interactions

Jason Weston
Meta AI
2023-07-14 15:50:00 – Room: **Harbour B**

Abstract: We discuss techniques that can be used to learn how to improve AIs (dialogue models) by interacting with organic users "in the wild". Training models with organic data is challenging because interactions with people in the wild include both high quality conversations and feedback, as well as adversarial and toxic behavior. We thus study techniques that enable learning from helpful teachers while avoiding learning from people who are trying to trick the model into unhelpful or toxic responses. We present BlenderBot 3x, an update on the conversational model BlenderBot 3, trained on 6M such interactions from participating users of the system. BlenderBot 3x is both preferred in conversation to BlenderBot 3, and is shown to produce safer responses in challenging situations. We then discuss how we believe continued use of these techniques – and improved variants – can lead to further gains.

Bio: Jason Weston is a research scientist at Meta AI, USA and a Visiting Research Professor at NYU. He earned his PhD in machine learning at Royal Holloway, University of London and AT&T Research in Red Bank, NJ (advisors: Alex Gammerman, Volodya Vovk and Vladimir Vapnik) in 2000. From 2002-2003 he was a research scientist at the Max Planck Institute for Biological Cybernetics, Tuebingen, Germany. From 2003-2009 he was a research staff member at NEC Labs America, Princeton. From 2009-2014 he was a research scientist at Google, NY. Jason's papers include best paper awards at ICML and ECML, and a Test of Time Award for his work A Unified Architecture for Natural Language Processing: Deep Neural Networks with Multitask Learning, ICML 2008 (with Ronan Collobert).

Table of Contents

Program

Friday, July 14, 2023

09:00 - 09:10 *Opening Remarks*

09:10 - 09:40 *Inclusive Conversational AI for Positive Impact (Diyi Yang)*

09:40 - 10:10 *Build it for One @ Right Place Right Time: Leveraging Context in Conversational Systems (Larry Heck)*

10:10 - 10:30 *Generating Video Game Scripts with Style (Best Paper)*

10:30 - 10:50 *Coffee Break*

10:50 - 12:00 *Poster Session*

12:00 - 13:30 *Lunch Break*

13:30 - 14:00 *Building Better Writing Assistants In the Era of Conversational LLMs (Vipul Raheja)*

14:00 - 14:20 *Response Generation in Longitudinal Dialogues: Which Knowledge Representation Helps?*

14:20 - 14:40 *On the Underspecification of Situations in Open-domain Conversational Datasets (Outstanding Paper)*

14:40 - 15:00 *Correcting Semantic Parses with Natural Language through Dynamic Schema Encoding*

15:00 - 15:30 *Dialogue Evaluation via Offline Reinforcement Learning and Emotion Prediction (Nurul Lubis)*

15:30 - 15:50 *Coffee Break*

15:50 - 16:20 *Improving Open Language Models by Learning from Organic Interactions (Jason Weston)*

16:20 - 16:40 *Conversational Recommendation as Retrieval: A Simple, Strong Baseline (Outstanding Paper)*

16:40 - 17:00 *A Survey of Challenges and Methods in the Computational Modeling of Multi-Party Dialog*

Friday, July 14, 2023 (continued)

17:00 - 17:10 *Closing Remarks*