

NLP4ConvAI 2024

The 6th Workshop on NLP for Conversational AI

Proceedings of the Workshop

August 16, 2024

©2024 Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
317 Sidney Baker St. S
Suite 400 - 134
Kerrville, TX 78028
USA
Tel: +1-855-225-1962
acl@aclweb.org

ISBN 979-8-89176-160-5

Introduction

We are excited to welcome you to NLP4ConvAI 2024, the 6th Annual Workshop on NLP for Conversational AI, co-located with ACL 2024 at Bangkok, Thailand.

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 five previous editions of NLP for Conversational AI workshops at ACL & EMNLP, NLP4ConvAI 2024 is a one-day workshop including keynotes, oral presentations and posters.

We received 37 submissions this year, consisting of 25 long papers and 12 short papers. We had a total of 68 program committee (PC) members. At least three PC members reviewed each of the papers. We accepted 14 papers: 14 long papers and 0 short papers. These numbers give an overall acceptance rate of 38%, with the long and short papers acceptance rate being 56% and 0% respectively. Out of the 14 accepted papers, 6 are being presented as oral presentations and the remaining in a poster session. We have also identified one best paper (Revealing User Familiarity Bias in Task-Oriented Dialogue via Interactive Evaluation) and one outstanding paper (Engineering Conversational Search Systems: A Review of Applications, Architectures, and Functional Components).

In addition, the workshop program consists of invited talks given by leading practitioners in industry and academia. 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 2024 at ACL and contribute to the future success of our community!

NLP4ConvAI 2024 Organizers

Elnaz Nouri, General and Sponsorship Chair

Abhinav Rastogi, General Chair

Georgios Spithourakis, Program Chair

Bing Liu, Publication chair

Yun-Nung (Vivian) Chen, Publication chair

Yu Li, Diversity & Publicity Chair

Alon Albalak, Shared Task Chair

Hiromi Wakaki, Shared Task Chair

Alexandros Papangelis, Advisory Board

Organizing Committee

General Chair and Sponsorship Chair

Elnaz Nouri, Microsoft Research

General Chair

Abhinav Rastogi, Google Research

Program Chair

Georgios Spithourakis, ex-PolyAI, Entrepreneur First

Publication Chair

Bing Liu, Meta

Publication Chair

Yun-Nung Chen, National Taiwan University

Diversity and Publicity Chair

Yu Li, Columbia University

Shared Task Chair

Alon Albalak, University of California, Santa Barbara

Shared Task Chair

Hiromi Wakaki, Sony Group Corporation

Advisory Board

Alexandros Papangelis, Amazon Alexa AI

Program Committee

Reviewers

Divyanshu Aggarwal
Renat Aksitov, Google
Alon Albalak
Raghuram Mandyam Annasamy, Google
Abhinav Arora, Facebook
Mukul Bhutani, Google
Simone Caldarella
Jie Cao
Simral Chaudhary
Maximillian Chen
Parag Pravin Dakle
Nina Dethlefs, University of Hull
Daniel Elkind
Ashwinkumar Ganesan, Amazon
Ananya Ganesh
Shubham Garg, Amazon
Alborz Geramifard, Meta
Parker Glenn, Fidelity Investments
Xiaojie Guo, International Business Machines
Ruohao Guo, Georgia Institute of Technology
Meiqi Guo, Google
Vivek Gupta, University of Pennsylvania, United States
Raghav Gupta, Research, Google
Shachi H. Kumar, Intel Labs
Richang Hong, Hefei University of Technology
Songbo Hu, Language Technology Lab, University of Cambridge
Chao-Wei Huang, National Taiwan University
Mihir Kale, Google
HyeongSik Kim
Abi Komma
Nitish Kulkarni, Google
Harsh Lara, Research, Google
Stefan Larson, Vanderbilt University
Benjamin Levin, PolyAI Limited
Hsien-chin Lin
Yen-Ting Lin, National Taiwan University
Wolfgang Maier, Mercedes Benz Research & Development
Dana Moukheiber
Seyed Mahed Mousavi
Lin Ning, Google
Naoki Otani, Megagon Labs
Alexandros Papangelis, Amazon
Udita Patel, Amazon
Wei Peng, Huawei Technologies Ltd.
Shiva Kumar Pentyala, Salesforce.com
Samrat Phatale, Google DeepMind

Evgeniia Razumovskaia, University of Cambridge
Giuseppe Riccardi, University of Trento
Lina Maria Rojas-Barahona
Igor Shalyminov, Amazon
Akshat Shrivastava
Lei Shu, Google
Shubham Shukla
Georgios P. Spithourakis
Anirudh Sundar, Georgia Institute of Technology
Anh Duong Trinh, National College of Ireland and Technological University Dublin
Yi-Lin Tuan, UC Santa Barbara
Gokhan Tur, University of Illinois at Urbana-Champaign
Stefan Ultes, Otto-Friedrich Universität Bamberg
David Vandyke
Ivan Vulić, University of Cambridge and PolyAI Limited
Peidong Wang
Peratham Wiriathamabhum
Jiaxing Wu, Google
Longfei Yang, Tokyo Institute of Technology, Tokyo Institute of Technology
Qiusi Zhan, University of Illinois Urbana-Champaign
Jianguo Zhang, Salesforce AI Research

Table of Contents

<i>On the Benchmarking of LLMs for Open-Domain Dialogue Evaluation</i> John Mendonça, Alon Lavie and Isabel Trancoso	1
<i>Exploring Description-Augmented Dataless Intent Classification</i> Ruoyu Hu, Foad Khosmood and Abbas Edalat	13
<i>Revealing User Familiarity Bias in Task-Oriented Dialogue via Interactive Evaluation</i> Takyoun Kim, Jamin Shin, Young-Ho Kim, Sanghwan Bae and Sungdong Kim	37
<i>Evaluating Robustness of Open Dialogue Summarization Models in the Presence of Naturally Occurring Variations</i> Ankita Gupta, Chulaka Gunasekara, Hui Wan, Jatin Ganhotra, Sachindra Joshi and Marina Danilevsky	56
<i>Engineering Conversational Search Systems: A Review of Applications, Architectures, and Functional Components</i> Phillip Schneider, Wessel Poelman, Michael Rovatsos and Florian Matthes	73
<i>Efficient Dynamic Hard Negative Sampling for Dialogue Selection</i> Janghoon Han, Dongkyu Lee, Joongbo Shin, Hyunkyung Bae, Jeesoo Bang, Seonghwan Kim, Stanley Jungkyu Choi and Honglak Lee	89
<i>Chamain: Harmonizing Character Persona Integrity with Domain-Adaptive Knowledge in Dialogue Generation</i> Seung-Moo Yang, Jeehyun Lee and Won Ik Cho	101
<i>Faithful Persona-based Conversational Dataset Generation with Large Language Models</i> Pegah Jandaghi, Xianghai Sheng, Xinyi Bai, Jay Pujara and Hakim Sidahmed	114

Program

Friday, August 16, 2024

- 09:00 - 09:10 *Opening Remarks*
- 09:10 - 09:40 *INVITED SPEAKER (Scott Yih)*
- 09:40 - 10:10 *INVITED SPEAKER (Prof. Mausam)*
- 10:10 - 10:30 *Revealing User Familiarity Bias in Task-Oriented Dialogue via Interactive Evaluation (Best Paper)*
- 10:30 - 10:50 *Coffee Break*
- 10:50 - 12:00 *Poster Session*
- 12:00 - 13:30 *Lunch Break*
- 13:30 - 14:00 *INVITED SPEAKER (Xiang Gao)*
- 14:00 - 14:20 *Engineering Conversational Search Systems: A Review of Applications, Architectures, and Functional Components (Outstanding Paper)*
- 14:20 - 14:40 *Visualizing Dialogues: Enhancing Image Selection through Dialogue Understanding with Large Language Models*
- 14:40 - 15:00 *SOLID: Self-seeding and Multi-intent Self-instructing LLMs for Generating Intent-aware Information-Seeking Dialogs*
- 15:00 - 15:30 *INVITED SPEAKER (Ben Athiwaratkun)*
- 15:30 - 15:50 *Coffee Break*
- 15:50 - 16:20 *SHARED TASK*
- 16:20 - 16:40 *Aligning to Thousands of Preferences via System Message Generalization*
- 16:40 - 17:00 *Evaluating Robustness of Open Dialogue Summarization Models in the Presence of Naturally Occurring Variations*

Friday, August 16, 2024 (continued)

17:00 - 17:10 *Closing Remarks*