EMNLP 2021

NLP for Conversational AI

Proceedings of the 3rd Workshop

November 10, 2021

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 978-1-954085-86-2

Introduction

Welcome to the 3rd Workshop on NLP for Conversational AI, at EMNLP 2021.

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 2021 is competitive. We received 92 submissions, and after a rigorous review process, we only accepted 32. There are a total of 26 accepted regular workshop papers, 1 extended abstract, and 5 cross-submissions. The workshop overall acceptance rate is about just under 35%. We hope you will enjoy NLP4ConvAI 2021 at EMNLP and contribute to the future success of our community!

NLP4ConvAI 2021 Organizers
Alexandros Papangelis, Amazon Alexa AI
Pawel Budzianowski, PolyAI
Bing Liu, Facebook
Elnaz Nouri, Microsoft Research
Abhinav Rastogi, Google Research
Yun-Nung (Vivian) Chen, National Taiwan University

Organizing Committee:

Alexandros Papangelis, Amazon Alexa AI
Pawel Budzianowski, PolyAI
Bing Liu, Facebook
Elnaz Nouri, Microsoft Research
Abhinav Rastogi, Google Research
Yun-Nung (Vivian) Chen, National Taiwan University

Program Committee:

Abhinav Arora, Facebook
Yuan Cao, Google
Iñigo Casanueva, PolyAI
Daniel Cer, Google
Ta-Chung Chi, Carnegie Mellon University
Aleksandr Chuklin, Google
Samuel Coope, PolyAI
Paul Crook, Facebook
Aditya Gupta, Google
Sonal Gupta, Facebook
Raghav Gupta, Google
Dilek Hakkani-Tur, Amazon
Chao-Wei Huang, National Taiwan University
Ni Lao, Apple
Stefan Larson, SkySync
Wenqiang Lei, National University of Singapore
Ben Levin, PolyAI
Xiujun Li, University of Washington
Zhaojiang Lin, HKUST
Yang Liu, Amazon
Xiaohu Liu, Amazon
Andrea Madotto, HKUST
Shikib Mehri, Carnegie Mellon University
Gaurav Menghani, Google
Pranab Mohanty, Facebook
Seungwhan Moon, Facebook
Wei Peng, Huawei
Julien Perez, Naver Labs
Elahe Rahimtoroghi, Google
Anirudha Rayasam, Google
Marek Rei, Imperial College London
Lina M. Rojas Barahona, Orange Labs
Adithya Sagar, Facebook
Chinnadhurai Sankar, Facebook
Lei Shu, Amazon
Shang-Yu Su, National Taiwan University
Kai Sun, Cornell University
Ryuichi Takanobu, Alibaba
Gokhan Tur, Amazon

Stefan Ultes, Daimler AG
David Vandyke, Apple
Ivan Vulić, University of Cambridge
Peng Wang, AlphaSense
Zhiguang Wang, Facebook
Chien-Sheng Wu, Salesforce
Hu Xu, Facebook
Zi Yang, Google
Yi-Ting Yeh, Carnegie Mellon University
Emine Yilmaz, University College London
Lili Yu, Facebook
Hamed Zamani, University of Massachusetts Amherst
Xiaoxue Zang, Google
Hongyuan Zhan, Facebook
Bin Zhang, Google
Su Zhu, AISpeech

Invited Speakers:

Pascale Fung, Hong Kong University of Science & Technology
Koichiro Yoshino, Institute of Physical and Chemical Research (RIKEN)
Mona Diab, George Washington University
Idan Szpektor, Google
Tsung-Hsien Wen, PolyAI

Table of Contents

Conference Program

Taking Things Personally: Third Person to First Person Rephrasing
Marcel Granero Moya and Panagiotis Agis Oikonomou Filandras

Few-Shot Intent Classification by Gauging Entailment Relationship Between Utterance and Semantic Label
Jin Qu, Kazuma Hashimoto, Wenhao Liu, Caiming Xiong and Yingbo Zhou

Personalized Extractive Summarization Using an Ising Machine Towards Real-time Generation of Efficient and Coherent Dialogue Scenarios
Hiroaki Takatsu, Takahiro Kashikawa, Koichi Kimura, Ryota Ando and Yoichi Matsuyama

Multilingual Paraphrase Generation For Bootstrapping New Features in Task-Oriented Dialog Systems
Subhadarshi Panda, Caglar Tirkaz, Tobias Falke and Patrick Lehnen

Overcoming Conflicting Data when Updating a Neural Semantic Parser
David Gaddy, Alex Kouzemtchenko, Pavankumar Reddy Muddireddy, Prateek Kolhar and Rushin Shah

Not So Fast, Classifier – Accuracy and Entropy Reduction in Incremental Intent Classification
Lianna Hrycyk, Alessandra Zarcone and Luzian Hahn

On the Robustness of Intent Classification and Slot Labeling in Goal-oriented Dialog Systems to Real-world Noise
Sailik Sengupta, Jason Krone and Saab Mansour

Amendable Generation for Dialogue State Tracking
Xin Tian, Liankai Huang, Yingzhan Lin, Siqi Bao, Huang He, Yunyi Yang, Hua Wu, Fan Wang and Shuqi Sun

What Went Wrong? Explaining Overall Dialogue Quality through Utterance-Level Impacts
James D. Finch, Sarah E. Finch and Jinho D. Choi

XPersona: Evaluating Multilingual Personalized Chatbot
Zhaojiang Lin, Zihan Liu, Genta Indra Winata, Samuel Cahyawijaya, Andrea Madotto, Yejin Bang, Etsuko Ishii and Pascale Fung

Collaborative Data Relabeling for Robust and Diverse Voice Apps Recommendation in Intelligent Personal Assistants
Qian Hu, Thahir Mohamed, Zheng Gao, Xibin Gao, Radhika Arava, Xiyao Ma and Mohamed AbdelHady

Semi-supervised Intent Discovery with Contrastive Learning
Xiang Shen, Yinge Sun, Yao Zhang and Mani Najmabadi

No Day Set (continued)

CS-BERT: a pretrained model for customer service dialogues
Peiyao Wang, Joyce Fang and Julia Reinspach

PLATO-KAG: Unsupervised Knowledge-Grounded Conversation via Joint Modeling
Xinxian Huang, Huang He, Siqi Bao, Fan Wang, Hua Wu and Haifeng Wang

Improving Dialogue State Tracking by Joint Slot Modeling
Ting-Rui Chiang and Yi-Ting Yeh

Learning to Learn End-to-End Goal-Oriented Dialog From Related Dialog Tasks
Janarthanan Rajendran, Jonathan K. Kummerfeld and Satinder Baveja

Personalized Search-based Query Rewrite System for Conversational AI
Eunah Cho, Ziyan Jiang, Jie Hao, Zheng Chen, Saurabh Gupta, Xing Fan and Chenlei Guo

Dialogue Response Generation via Contrastive Latent Representation Learning
Shuyang Dai, Guoyin Wang, Sunghyun Park and Sungjin Lee

AuGPT: Auxiliary Tasks and Data Augmentation for End-To-End Dialogue with Pre-Trained Language Models
Jonáš Kulhánek, Vojtěch Hudeček, Tomáš Nekvinda and Ondřej Dušek

Investigating Pretrained Language Models for Graph-to-Text Generation
Leonardo F. R. Ribeiro, Martin Schmitt, Hinrich Schütze and Iryna Gurevych

Style Control for Schema-Guided Natural Language Generation
Alicia Tsai, Shereen Oraby, Vittorio Perera, Jiun-Yu Kao, Yuheng Du, Anjali Narayan-Chen, Tagyoung Chung and Dilek Hakkani-Tur

Using Pause Information for More Accurate Entity Recognition
Sahas Dendukuri, Pooja Chitkara, Joel Ruben Antony Moniz, Xiao Yang, Manos Tsagkias and Stephen Pulman

Think Before You Speak: Learning to Generate Implicit Knowledge for Response Generation by Self-Talk
Pei Zhou, Behnam Hedayatnia, Karthik Gopalakrishnan, Seokhwan Kim, Jay Pujara, Xiang Ren, Yang Liu and Dilek Hakkani-Tur

Teach Me What to Say and I Will Learn What to Pick:
Unsupervised Knowledge Selection Through Response Generation with Pretrained Generative Models
Ehsan Lotfi, Maxime De Bruyn, Jeska Buhmann and Walter Daelemans

No Day Set (continued)

Influence of user personality on dialogue task performance: A case study using a rule-based dialogue system
Ao Guo, Atsumoto Ohashi, Ryu Hirai, Yuya Chiba, Yuiko Tsunomori and Ryuichiro Higashinaka

Towards Code-Mixed Hinglish Dialogue Generation
Vibhav Agarwal, Pooja Rao and Dinesh Babu Jayagopi

Towards Zero and Few-shot Knowledge-seeking Turn Detection in Task-orientated Dialogue Systems
Di Jin, Shuyang Gao, Seokhwan Kim, Yang Liu and Dilek Hakkani-Tur