


SIGDIAL 2019



**20th Annual Meeting of the
Special Interest Group on Discourse and
Dialogue**



Proceedings of the Conference

11-13 September 2019
Stockholm, Sweden

In cooperation with:

Association for Computational Linguistics (ACL)

International Speech Communication Association (ISCA)

Association for the Advancement of Artificial Intelligence (AAAI)

We thank our sponsors:

Honda Research Institute

Amazon Alexa

Rasa Technologies

Toshiba Research Europe

Interactions

Spotify

Educational Testing Service (ETS)

KTH Royal Institute of Technology

Microsoft Research

Apple

Monash University

Platinum



Gold



Silver



Measuring the Power of Learning.™



MONASH
INFORMATION
TECHNOLOGY

Bronze

TOSHIBA

In cooperation with



©2019 The Association for Computational Linguistics

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 978-1-950737-61-1

Introduction

We are excited to welcome you to SIGDIAL 2019, the 20th Annual Meeting of the Special Interest Group on Discourse and Dialogue. This year the conference is being held in Stockholm, Sweden, on September 11-13, with the Satellite Event YRRSDS 2019 (Young Researchers' Roundtable on Spoken Dialog Systems), and in close temporal proximity with Interspeech 2019, held in Gratz, Austria, and SemDial 2019, held in London, UK.

The SIGDIAL conference is a premier publication venue for research in discourse and dialogue. This year, the program includes three keynote talks, six oral presentation sessions, three poster sessions including six demonstrations, a panel entitled “The Future of Dialogue Research” organized by Phil Cohen, and a special session entitled “Implications of Deep Learning for Dialogue Modeling” organized by Nigel Ward, Yun-Nung (Vivian) Chen, Tatsuya Kawahara and Gabriel Skantze.

We received a record 146 submissions this year, about one third more than the submissions received in 2018. The 146 submissions comprised 93 long papers, 43 short papers and 10 demo descriptions. All submissions received at least three reviews. When making our selections for the program, we carefully considered the reviews and the comments made during the discussions among reviewers. The members of the Program Committee did an excellent job in reviewing the submitted papers, and we thank them for their essential role in selecting the accepted papers and helping produce a high quality program for the conference. In line with the SIGDIAL tradition, our aim has been to create a balanced program that accommodates as many favourably rated papers as possible. We accepted 51 papers: 33 long papers—three of which were converted to short papers, 13 short papers, and five demo descriptions. These numbers give an overall acceptance rate of 35%, with the following rates for the different types of papers: 35% for long papers, 30% for short papers and 50% for demo descriptions. It is worth noting that the acceptance rate for long papers was significantly lower than that of previous years – a result of the unusually large number of submissions.

Each of the three conference days features one keynote and one poster session, with the remaining time given to oral presentations, demos, the panel and the special session. The oral presentations comprise 16 of the long papers and three long papers selected for the special session. The three poster sessions feature the remaining long papers, all the short papers and two work-in-progress special session papers. In terms of content, about a quarter of the accepted papers discuss datasets and evaluation issues, and approximately half employ deep learning to address problems in discourse and dialogue—a trend also exhibited in recent Language Technology conferences. Finally, this SIGDIAL features an invited demo that showcases research conducted in the department of Robotics, Perception and Learning at KTH, the host institution.

A conference of this scale requires advice, help and enthusiastic participation of many parties, and we have a big ‘thank you’ to say to all of them. Regarding the program, we thank our three keynote speakers, Dan Bohus (Microsoft Research, Redmond, Washington, US), Mirella Lapata (University of Edinburgh, UK) and Helen Meng (Chinese University of Hong Kong, China) for their inspiring talks on situated interaction, learning neural natural language interfaces, and dialogue research application to healthcare, e-commerce and education. We also thank the organizer of the panel on the Future of Dialogue Research, and the organizers of the special session on Implications of Deep Learning for Dialogue Modeling. We are grateful for their smooth and efficient coordination with the main conference. In addition, we thank Alex Papangelis, Mentoring Chair for SIGDIAL 2019, for his dedicated work on the mentoring process. The goal of mentoring is to assist authors of papers that contain important ideas but require significant stylistic modifications. In total, seven of the accepted papers received mentoring, and we thank our mentoring team for their excellent support of the authors.

We extend special thanks to our Local Chair, Gabriel Skantze, and his team, including the student

volunteers who provide on-site assistance. SIGDIAL 2019 would not have been possible without their effort in arranging the conference venue, handling registration, making banquet arrangements, numerous preparations for the conference, and last but not least, Gabriel's personal contributions, which exceeded those of a local organizer.

Mikio Nakano, our Sponsorship Chair, has conducted the massive task of recruiting and liaising with our conference sponsors, many of whom continue to contribute year after year. Sponsorships support valuable aspects of the program, such as lunches and the conference banquet. We thank Mikio for his dedicated work and his assistance with conference planning. We gratefully acknowledge the support of our sponsors: (Platinum level) Honda Research Institute, Interactions and Microsoft Research; (Gold level) Amazon Alexa, Apple, Rasa Technologies and Spotify; (Silver level) Educational Testing Service (ETS) and Monash University; and (Bronze level) Toshiba Research Europe. We also thank the KTH Royal Institute of Technology for its generous sponsorship as host.

Koichiro Yoshino, our publicity chair, was tireless in the design and maintenance of the SIGDIAL 2019 website, cheerfully coping with multiple and constant changes; and Stefan Ultes, our publication chair, capped the long organizational process by putting together these high quality conference proceedings. We thank the SIGdial board, both current and emeritus officers, Gabriel Skantze, Mikio Nakano, Vikram Ramanarayanan, Ethan Selfridge, Kallirroi Georgila, Jason Williams and Amanda Stent, for their advice and support from beginning to end.

We once again thank our program committee members for committing their time to help us select an excellent technical program. Finally, we thank all the authors who submitted to the conference and all the conference participants for making SIGDIAL 2019 a success and for growing the research areas of discourse and dialogue with their fine work.

Satoshi Nakamura, General Chair

Milica Gašić and Ingrid Zukerman, Program Co-Chairs

Committee Listings

General Chair:

Satoshi Nakamura, Nara Institute of Science and Technology, Japan

Program Chairs:

Milica Gašić, Heinrich Heine University Düsseldorf, Germany

Ingrid Zukerman, Monash University, Australia

Local Chair:

Gabriel Skantze, KTH Royal Institute of Technology, Sweden

Sponsorship Chair:

Mikio Nakano, Honda Research Institute, Japan

Mentoring Chair:

Alexandros Papangelis, Uber AI, United States

Publication Chair:

Stefan Ultes, Daimler AG, Germany

Publicity Chair:

Koichiro Yoshino, Nara Institute of Science and Technology, Japan

SIGdial Officers:

President: Gabriel Skantze, KTH Royal Institute of Technology, Sweden

Vice President: Mikio Nakano, Honda Research Institute, Japan

Secretary: Vikram Ramanarayanan, Educational Testing Service (ETS) Research, United States

Treasurer: Ethan Selfridge, Interactions, United States

President Emeritus: Jason Williams, Apple, United States

Vice President Emeritus: Kallirroi Georgila, University of Southern California, United States

Program Committee:

Masahiro Araki, Kyoto Institute of Technology, Japan

Ron Artstein, USC Institute for Creative Technologies, United States

Timo Baumann, University of Hamburg, Germany

Frederic Bechet, Aix Marseille Université - LIS/CNRS, France

Steve Beet, Aculab plc, United Kingdom

Jose Miguel Benedi, Universitat Politècnica de València, Spain

Luciana Benotti, Universidad Nacional de Córdoba, Argentina

Nate Blaylock, Nuance Communications, United States

Dan Bohus, Microsoft Research, United States

Johan Boye, KTH, Sweden

Paweł Budzianowski, University of Cambridge, United Kingdom

Hendrik Buschmeier, Bielefeld University, Germany

Andrew Caines, University of Cambridge, United Kingdom
Giuseppe Carenini, University of British Columbia, Canada
Christophe Cerisara, Université de Lorraine, CNRS, LORIA, France
Joyce Chai, Michigan State University, United States
Senthil Chandramohan, Microsoft, United States
Lin Chen, Head of AI, Cambia Health Solutions, United States
Philip Cohen, Monash University, Australia
Mark Core, University of Southern California, United States
Heriberto Cuayahuitl, University of Lincoln, United Kingdom
Vera Demberg, Saarland University, Germany
Nina Dethlefs, University of Hull, United Kingdom
David DeVault, University of Southern California, United States
Barbara Di Eugenio, University of Illinois at Chicago, United States
Giuseppe Di Fabbrizio, VUI, Inc., United States
Maxine Eskenazi, Carnegie Mellon University, United States
Keelan Evanini, Educational Testing Service, United States
Mauro Falcone, Fondazione Ugo Bordoni, Italy
Raquel Fernandez, ILLC, University of Amsterdam, Netherlands
Kallirroi Georgila, University of Southern California, United States
Alborz Geramifard, Facebook, United States
Jonathan Ginzburg, Université Paris-Diderot (Paris 7), France
Ivan Habernal, Technische Universität Darmstadt, Germany
Dilek Hakkani-Tur, Amazon Alexa AI, United States
Helen Hastie, Heriot-Watt University, United Kingdom
Michael Heck, Heinrich Heine University, Germany
Ryuichiro Higashinaka, NTT Media Intelligence Labs., Japan
Takuya Hiraoka, NEC Central Research Laboratories, Japan
Keikichi Hirose, University of Tokyo, Japan
David M. Howcroft, Heriot-Watt University, United Kingdom
Michimasa Inaba, University of Electro-Communications, Japan
Koji Inoue, Kyoto University, Japan
David Janiszek, Université Paris Descartes, France
Kristiina Jokinen, AIRC, AIST, Japan
Arne Jonsson, Linkping University, Sweden
Pamela Jordan, University of Pittsburgh, United States
Filip Jurcicek, Apple Inc., United Kingdom
Tatsuya Kawahara, Kyoto University, Japan
Simon Keizer, Emotech North, United Kingdom
Casey Kennington, Boise State University, United States
Chandra Khatri, Uber AI, United States
Norihide Kitaoka, Toyohashi University of Technology, Japan
Kazunori Komatani, Osaka University, Japan
Stefan Kopp, Bielefeld University, Germany
Fabrice Lefevre, Avignon University, France
James Lester, North Carolina State University, United States
Junyi Jessy Li, University of Texas at Austin, United States
Pierre Lison, Norwegian Computing Centre, Norway
Diane Litman, University of Pittsburgh, United States
Bing Liu, Facebook, United States

Eduardo Lleida Solano, University of Zaragoza, Spain
José Lopes, Heriot Watt University, United Kingdom
Ramon Lopez-Cozar, University of Granada, Spain
Annie Louis, University of Edinburgh, United Kingdom
Nurul Lubis, Heinrich Heine University, Germany
Matthew Marge, Army Research Laboratory, United States
Helen Meng, Chinese University of Hong Kong, China
Teruhisa Misu, Honda Research Institute USA, United States
Satoshi Nakamura, Nara Institute of Science and Technology, Japan
Mikio Nakano, Honda Research Institute, Japan
Shashi Narayan, Google, United Kingdom
Douglas O’Shaughnessy, INRS-EMT (Univ. of Quebec), Canada
Mari Ostendorf, University of Washington, United States
Alexandros Papangelis, Uber AI, United States
Volha Petukhova, Saarland University, Germany
Paul Piwek, The Open University, United Kingdom
Andrei Popescu-Belis, HEIG-VD / HES-SO, Switzerland
Rashmi Prasad, Interactions Corporation, United States
Matthew Purver, Queen Mary University of London, United Kingdom
Vikram Ramanarayanan, Educational Testing Service R&D, United States
Antoine Raux, Apple, United States
Ehud Reiter, University of Aberdeen, United Kingdom
Norbert Reithinger, DFKI GmbH, Germany
Giuseppe Riccardi, University of Trento, Italy
Antonio Roque, Tufts University, United States
Carolyn Rose, Carnegie Mellon University, United States
Sophie Rosset, LIMSI, CNRS, Universit Paris-Saclay, France
Sakriani Sakti, Nara Institute of Science and Technology (NAIST) / RIKEN AIP, Japan
Ruhi Sarikaya, Amazon, United States
Niko Schenk, Goethe University Frankfurt am Main, Germany
David Schlangen, Bielefeld University, Germany
Ethan Selfridge, Interactions Corp, United States
Gabriel Skantze, KTH Speech Music and Hearing, Sweden
Manfred Stede, University of Potsdam, Germany
Georg Stemmer, Intel Corp., Germany
Matthew Stone, Rutgers University, United States
Svetlana Stoyanchev, Interactions Corporation, United States
Kristina Striegnitz, Union College, United States
Hiroaki Sugiyama, NTT Communication Science Labs., Japan
Antnio Teixeira, DETI/IEETA, University of Aveiro, Portugal
Joel Tetreault, Grammarly, United States
Sofia Thunberg, Linköping University, Sweden
Takenobu Tokunaga, Tokyo Institute of Technology, Japan
David Traum, University of Southern California, United States
Bo-Hsiang Tseng, University of Cambridge, United Kingdom
Gokhan Tur, Uber, United States
Stefan Ultes, Daimler AG, Germany
David Vandyke, Apple, United Kingdom
Marilyn Walker, University of California Santa Cruz, United States

Hsin-Min Wang, Academia Sinica, Taiwan
Nigel Ward, University of Texas at El Paso, United States
Jason D Williams, Microsoft Research, United States
Yen-chen Wu, University of Cambridge, United Kingdom
Koichiro Yoshino, Nara Institute of Science and Technology, Japan
Steve Young, Cambridge University, United Kingdom
Kai Yu, Shanghai Jiao Tong University, China
Zhou Yu, University of California, Davis, United States

Secondary Reviewers: Alexandre Denis, Sahar Ghannay, Churn-Jung Liao, Yi Ma, Sonja Stange,
Hiroki Tanaka, Martin Villalba

Mentors:

Dimitrios Alikaniotis, Grammarly, United States
Timo Baumann, University of Hamburg, Germany
Hendrik Buschmeier, Bielefeld University, Germany
Ivan Habernal, Technische Universität Darmstadt, Germany
Kornel Laskowski, Carnegie Mellon University, United States
Stefan Ultes, Daimler AG, Germany
David Vandyke, Apple, United Kingdom

Invited Speakers:

Dan Bohus, Microsoft Research, United States
Mirella Lapata, University of Edinburgh, United Kingdom
Helen Meng, Chinese University of Hong Kong, China

Table of Contents

<i>Deep Reinforcement Learning For Modeling Chit-Chat Dialog With Discrete Attributes</i> Chinnadhurai Sankar and Sujith Ravi	1
<i>Improving Interaction Quality Estimation with BiLSTMs and the Impact on Dialogue Policy Learning</i> Stefan Ultes	11
<i>Lifelong and Interactive Learning of Factual Knowledge in Dialogues</i> Sahisnu Mazumder, Bing Liu, Shuai Wang and Nianzu Ma	21
<i>Few-Shot Dialogue Generation Without Annotated Data: A Transfer Learning Approach</i> Igor Shalymov, Sungjin Lee, Arash Eshghi and Oliver Lemon	32
<i>SIM: A Slot-Independent Neural Model for Dialogue State Tracking</i> Chenguang Zhu, Michael Zeng and Xuedong Huang	40
<i>Simple, Fast, Accurate Intent Classification and Slot Labeling for Goal-Oriented Dialogue Systems</i> Arshit Gupta, John Hewitt and Katrin Kirchhoff	46
<i>Time Masking: Leveraging Temporal Information in Spoken Dialogue Systems</i> Rylan Conway and Mathias Lambert	56
<i>To Combine or Not To Combine? A Rainbow Deep Reinforcement Learning Agent for Dialog Policies</i> Dirk Väh and Ngoc Thang Vu	62
<i>Contextualized Representations for Low-resource Utterance Tagging</i> Bhargavi Paranjape and Graham Neubig	68
<i>Capturing Dialogue State Variable Dependencies with an Energy-based Neural Dialogue State Tracker</i> Anh Duong Trinh, Robert J. Ross and John D. Kelleher	75
<i>Leveraging Non-Conversational Tasks for Low Resource Slot Filling: Does it help?</i> Samuel Louvan and Bernardo Magnini	85
<i>Collaborative Multi-Agent Dialogue Model Training Via Reinforcement Learning</i> Alexandros Papangelis, Yi-Chia Wang, Piero Molino and Gokhan Tur	92
<i>Scoring Interactional Aspects of Human-Machine Dialog for Language Learning and Assessment using Text Features</i> Vikram Ramanarayanan, Matthew Mulholland and Yao Qian	103
<i>Spoken Conversational Search for General Knowledge</i> Lina M. Rojas Barahona, Pascal Bellec, Benoit Besset, Martinho Dossantos, Johannes Heinecke, munshi asadullah, Olivier Leblouch, Jeanyves. Lancien, Geraldine Damnati, Emmanuel Mory and Frederic Herledan	110
<i>Graph2Bots, Unsupervised Assistance for Designing Chatbots</i> Jean-Leon Bouraoui, Sonia Le Meitour, Romain Carbou, Lina M. Rojas Barahona and Vincent Lemaire	114
<i>On a Chatbot Conducting Dialogue-in-Dialogue</i> Boris Galitsky, Dmitry Ilvovsky and Elizaveta Goncharova	118
<i>DeepCopy: Grounded Response Generation with Hierarchical Pointer Networks</i> Semih Yavuz, Abhinav Rastogi, Guan-Lin Chao and Dilek Hakkani-Tur	122

<i>Towards End-to-End Learning for Efficient Dialogue Agent by Modeling Looking-ahead Ability</i> Zhuoxuan Jiang, Xian-Ling Mao, Ziming Huang, Jie Ma and Shaochun Li	133
<i>Unsupervised Dialogue Spectrum Generation for Log Dialogue Ranking</i> Xinnuo Xu, Yizhe Zhang, Lars Liden and Sungjin Lee	143
<i>Tree-Structured Semantic Encoder with Knowledge Sharing for Domain Adaptation in Natural Language Generation</i> Bo-Hsiang Tseng, Paweł Budzianowski, Yen-chen Wu and Milica Gasic	155
<i>Structured Fusion Networks for Dialog</i> Shikib Mehri, Tejas Srinivasan and Maxine Eskenazi	165
<i>Flexibly-Structured Model for Task-Oriented Dialogues</i> Lei Shu, Piero Molino, Mahdi Namazifar, Hu Xu, Bing Liu, Huaixiu Zheng and Gokhan Tur ..	178
<i>FriendsQA: Open-Domain Question Answering on TV Show Transcripts</i> Zhengzhe Yang and Jinho D. Choi	188
<i>Foundations of Collaborative Task-Oriented Dialogue: What’s in a Slot?</i> Philip Cohen	198
<i>Speaker-adapted neural-network-based fusion for multimodal reference resolution</i> Diana Kleingarn, Nima Nabizadeh, Martin Heckmann and Dorothea Kolossa	210
<i>Learning Question-Guided Video Representation for Multi-Turn Video Question Answering</i> Guan-Lin Chao, Abhinav Rastogi, Semih Yavuz, Dilek Hakkani-Tur, Jindong Chen and Ian Lane	215
<i>Zero-shot transfer for implicit discourse relation classification</i> Murathan Kurfalı and Robert Östling	226
<i>A Quantitative Analysis of Patients’ Narratives of Heart Failure</i> Sabita Acharya, Barbara Di Eugenio, Andrew Boyd, Richard Cameron, Karen Dunn Lopez, Pamela Martyn-Nemeth, Debaleena Chattopadhyay, Pantea Habibi, Carolyn Dickens, Haleh Vatani and Amer Ardati	232
<i>TDDiscourse: A Dataset for Discourse-Level Temporal Ordering of Events</i> Aakanksha Naik, Luke Breitfeller and Carolyn Rose	239
<i>Real Life Application of a Question Answering System Using BERT Language Model</i> Francesca Alloatti, Luigi Di Caro and Gianpiero Sportelli	250
<i>Hierarchical Multi-Task Natural Language Understanding for Cross-domain Conversational AI: HER-MIT NLU</i> Andrea Vanzo, Emanuele Bastianelli and Oliver Lemon	254
<i>Dialog State Tracking: A Neural Reading Comprehension Approach</i> Shuyang Gao, Abhishek Sethi, Sanchit Agarwal, Tagyoung Chung and Dilek Hakkani-Tur	264
<i>Cross-Corpus Data Augmentation for Acoustic Addressee Detection</i> Oleg Akhtiamov, Ingo Siegert, Alexey Karpov and Wolfgang Minker	274
<i>A Scalable Method for Quantifying the Role of Pitch in Conversational Turn-Taking</i> Kornel Laskowski, Marcin Włodarczyk and Mattias Heldner	284

<i>A Large-Scale User Study of an Alexa Prize Chatbot: Effect of TTS Dynamism on Perceived Quality of Social Dialog</i>	
Michelle Cohn, Chun-Yen Chen and Zhou Yu	293
<i>Influence of Time and Risk on Response Acceptability in a Simple Spoken Dialogue System</i>	
Andisheh Partovi and Ingrid Zukerman	307
<i>Characterizing the Response Space of Questions: a Corpus Study for English and Polish</i>	
Jonathan Ginzburg, Zulipiye Yusupujiang, Chuyuan Li, Kexin Ren and Paweł Łupkowski	320
<i>From Explainability to Explanation: Using a Dialogue Setting to Elicit Annotations with Justifications</i>	
Nazia Attari, Martin Heckmann and David Schlangen	331
<i>Prediction of User Emotion and Dialogue Success Using Audio Spectrograms and Convolutional Neural Networks</i>	
Athanasios Lykartsis and Margarita Kotti	336
<i>Modelling Adaptive Presentations in Human-Robot Interaction using Behaviour Trees</i>	
Nils Axelsson and Gabriel Skantze	345
<i>Coached Conversational Preference Elicitation: A Case Study in Understanding Movie Preferences</i>	
Filip Radlinski, Krisztian Balog, Bill Byrne and Karthik Krishnamoorthi	353
<i>A Crowd-based Evaluation of Abuse Response Strategies in Conversational Agents</i>	
Amanda Cercas Curry and Verena Rieser	361
<i>A Dynamic Strategy Coach for Effective Negotiation</i>	
Yiheng Zhou, He He, Alan W Black and Yulia Tsvetkov	367
<i>Investigating Evaluation of Open-Domain Dialogue Systems With Human Generated Multiple References</i>	
Prakhar Gupta, Shikib Mehri, Tiancheng Zhao, Amy Pavel, Maxine Eskenazi and Jeffrey Bigham	379
<i>User Evaluation of a Multi-dimensional Statistical Dialogue System</i>	
Simon Keizer, Ondřej Dušek, Xingkun Liu and Verena Rieser	392
<i>Dialogue Act Classification in Team Communication for Robot Assisted Disaster Response</i>	
Tatiana Anikina and Ivana Kruijff-Korbayova	399
<i>Multi-Task Learning of System Dialogue Act Selection for Supervised Pretraining of Goal-Oriented Dialogue Policies</i>	
Sarah McLeod, Ivana Kruijff-Korbayova and Bernd Kiefer	411
<i>B. Rex: a dialogue agent for book recommendations</i>	
Mitchell Abrams, Luke Gessler and Matthew Marge	418
<i>SpaceRefNet: a neural approach to spatial reference resolution in a real city environment</i>	
Dmytro Kalpakchi and Johan Boye	422
<i>Which aspects of discourse relations are hard to learn? Primitive decomposition for discourse relation classification</i>	
Charlotte Roze, Chloé Braud and Philippe Muller	432
<i>Discourse Relation Prediction: Revisiting Word Pairs with Convolutional Networks</i>	
Siddharth Varia, Christopher Hidey and Tuhin Chakrabarty	442

Conference Program

11 September 2019

09:00–09:20 *Welcome*

09:20–10:20 *Keynote 1 - Learning Natural Language Interfaces with Neural Models*
Mirella Lapata

10:20–10:50 *Coffee Break*

10:50–12:05 **Session 1 - Policy and Knowledge**

Deep Reinforcement Learning For Modeling Chit-Chat Dialog With Discrete Attributes

Chinnadhurai Sankar and Sujith Ravi

Improving Interaction Quality Estimation with BiLSTMs and the Impact on Dialogue Policy Learning

Stefan Ultes

Lifelong and Interactive Learning of Factual Knowledge in Dialogues

Sahisnu Mazumder, Bing Liu, Shuai Wang and Nianzu Ma

12:05–13:20 *Lunch*

11 September 2019 (continued)

13:20–15:10 Poster and Demos 1

Few-Shot Dialogue Generation Without Annotated Data: A Transfer Learning Approach

Igor Shalyminov, Sungjin Lee, Arash Eshghi and Oliver Lemon

SIM: A Slot-Independent Neural Model for Dialogue State Tracking

Chenguang Zhu, Michael Zeng and Xuedong Huang

Simple, Fast, Accurate Intent Classification and Slot Labeling for Goal-Oriented Dialogue Systems

Arshit Gupta, John Hewitt and Katrin Kirchhoff

Time Masking: Leveraging Temporal Information in Spoken Dialogue Systems

Rylan Conway and Mathias Lambert

To Combine or Not To Combine? A Rainbow Deep Reinforcement Learning Agent for Dialog Policies

Dirk V ath and Ngoc Thang Vu

Contextualized Representations for Low-resource Utterance Tagging

Bhargavi Paranjape and Graham Neubig

Capturing Dialogue State Variable Dependencies with an Energy-based Neural Dialogue State Tracker

Anh Duong Trinh, Robert J. Ross and John D. Kelleher

Leveraging Non-Conversational Tasks for Low Resource Slot Filling: Does it help?

Samuel Louvan and Bernardo Magnini

Collaborative Multi-Agent Dialogue Model Training Via Reinforcement Learning

Alexandros Papangelis, Yi-Chia Wang, Piero Molino and Gokhan Tur

Scoring Interactional Aspects of Human-Machine Dialog for Language Learning and Assessment using Text Features

Vikram Ramanarayanan, Matthew Mulholland and Yao Qian

Spoken Conversational Search for General Knowledge

Lina M. Rojas Barahona, Pascal Bellec, Benoit Besset, Martinho Dossantos, Johannes Heinecke, munshi asadullah, Olivier Leblouch, Jeanyves Lancien, Geraldine Damnati, Emmanuel Mory and Frederic Herledan

11 September 2019 (continued)

Graph2Bots, Unsupervised Assistance for Designing Chatbots

Jean-Leon Bouraoui, Sonia Le Meitour, Romain Carbou, Lina M. Rojas Barahona and Vincent Lemaire

On a Chatbot Conducting Dialogue-in-Dialogue

Boris Galitsky, Dmitry Ilvovsky and Elizaveta Goncharova

15:10–15:40 *Coffee Break*

15:40–16:55 **Session 2 (Special Session) - Implications of Deep Learning for Dialogue Modeling**

DeepCopy: Grounded Response Generation with Hierarchical Pointer Networks

Semih Yavuz, Abhinav Rastogi, Guan-Lin Chao and Dilek Hakkani-Tur

Towards End-to-End Learning for Efficient Dialogue Agent by Modeling Looking-ahead Ability

Zhuoxuan Jiang, Xian-Ling Mao, Ziming Huang, Jie Ma and Shaochun Li

Unsupervised Dialogue Spectrum Generation for Log Dialogue Ranking

Xinnuo Xu, Yizhe Zhang, Lars Liden and Sungjin Lee

16:55–17:55 *Panel: The Future of Dialogue Research*

Organizer: Phil Cohen

Vikram Ramanarayanan, Educational Testing Service (ETS) Research, USA

Sujith Savi, Google, USA

Gabriel Skantze, KTH Royal Institute of Technology, Sweden

18:15–19:45 *Reception*

12 September 2019

09:00–10:00 *Keynote 2 - Situated Interaction*
Dan Bohus

10:00–10:30 *Coffee Break*

10:30–11:45 **Session 3 - Generation and End-to-end Dialogue Systems**

Tree-Structured Semantic Encoder with Knowledge Sharing for Domain Adaptation in Natural Language Generation

Bo-Hsiang Tseng, Paweł Budzianowski, Yen-chen Wu and Milica Gasic

Structured Fusion Networks for Dialog

Shikib Mehri, Tejas Srinivasan and Maxine Eskenazi

Flexibly-Structured Model for Task-Oriented Dialogues

Lei Shu, Piero Molino, Mahdi Namazifar, Hu Xu, Bing Liu, Huaixiu Zheng and Gokhan Tur

11:45–13:00 *Lunch*

13:00–14:15 **Poster and Demos 2**

FriendsQA: Open-Domain Question Answering on TV Show Transcripts

Zhengzhe Yang and Jinho D. Choi

Foundations of Collaborative Task-Oriented Dialogue: What's in a Slot?

Philip Cohen

Speaker-adapted neural-network-based fusion for multimodal reference resolution

Diana Kleingarn, Nima Nabizadeh, Martin Heckmann and Dorothea Kolossa

Learning Question-Guided Video Representation for Multi-Turn Video Question Answering

Guan-Lin Chao, Abhinav Rastogi, Semih Yavuz, Dilek Hakkani-Tur, Jindong Chen and Ian Lane

12 September 2019 (continued)

Zero-shot transfer for implicit discourse relation classification

Murathan Kurfalı and Robert Östling

A Quantitative Analysis of Patients' Narratives of Heart Failure

Sabita Acharya, Barbara Di Eugenio, Andrew Boyd, Richard Cameron, Karen Dunn Lopez, Pamela Martyn-Nemeth, Debaleena Chattopadhyay, Pantea Habibi, Carolyn Dickens, Haleh Vatani and Amer Ardati

TDDiscourse: A Dataset for Discourse-Level Temporal Ordering of Events

Aakanksha Naik, Luke Breitfeller and Carolyn Rose

Real-time Generation of Unambiguous Spatial Referring Expressions

Fethiye Irmak Dogan, Sinan Kalkan and Iolanda Leite

Real Life Application of a Question Answering System Using BERT Language Model

Francesca Alloatti, Luigi Di Caro and Gianpiero Sportelli

14:15–15:05 Session 4 - Understanding and Dialogue State Tracking

Hierarchical Multi-Task Natural Language Understanding for Cross-domain Conversational AI: HERMIT NLU

Andrea Vanzo, Emanuele Bastianelli and Oliver Lemon

Dialog State Tracking: A Neural Reading Comprehension Approach

Shuyang Gao, Abhishek Sethi, Sanchit Agarwal, Tagyoung Chung and Dilek Hakkani-Tur

15:05–15:35 Coffee Break

12 September 2019 (continued)

15:35–16:25 Session 5 - Acoustics

Cross-Corpus Data Augmentation for Acoustic Addressee Detection

Oleg Akhramov, Ingo Siegert, Alexey Karpov and Wolfgang Minker

A Scalable Method for Quantifying the Role of Pitch in Conversational Turn-Taking

Kornel Laskowski, Marcin Włodarczyk and Mattias Heldner

16:25–17:10 Sponsor Session

18:30–21:00 *Banquet at Vasa Museum*

13 September 2019

09:00–10:00 *The Many Facets of Dialog*
Helen Meng

10:00–10:30 *Coffee Break*

10:30–11:45 **Session 6 - Evaluation and Data**

A Large-Scale User Study of an Alexa Prize Chatbot: Effect of TTS Dynamism on Perceived Quality of Social Dialog

Michelle Cohn, Chun-Yen Chen and Zhou Yu

Influence of Time and Risk on Response Acceptability in a Simple Spoken Dialogue System

Andisheh Partovi and Ingrid Zukerman

Characterizing the Response Space of Questions: a Corpus Study for English and Polish

Jonathan Ginzburg, Zulipiye Yusupujiang, Chuyuan Li, Kexin Ren and Paweł Łupkowski

11:45–13:00 *Lunch*

13 September 2019 (continued)

13:00–14:50 Poster and Demos 3

From Explainability to Explanation: Using a Dialogue Setting to Elicit Annotations with Justifications

Nazia Attari, Martin Heckmann and David Schlangen

Prediction of User Emotion and Dialogue Success Using Audio Spectrograms and Convolutional Neural Networks

Athanasios Lykartsis and Margarita Kotti

Modelling Adaptive Presentations in Human-Robot Interaction using Behaviour Trees

Nils Axelsson and Gabriel Skantze

Coached Conversational Preference Elicitation: A Case Study in Understanding Movie Preferences

Filip Radlinski, Krisztian Balog, Bill Byrne and Karthik Krishnamoorthi

A Crowd-based Evaluation of Abuse Response Strategies in Conversational Agents

Amanda Cercas Curry and Verena Rieser

A Dynamic Strategy Coach for Effective Negotiation

Yiheng Zhou, He He, Alan W Black and Yulia Tsvetkov

Investigating Evaluation of Open-Domain Dialogue Systems With Human Generated Multiple References

Prakhar Gupta, Shikib Mehri, Tiancheng Zhao, Amy Pavel, Maxine Eskenazi and Jeffrey Bigham

User Evaluation of a Multi-dimensional Statistical Dialogue System

Simon Keizer, Ondřej Dušek, Xingkun Liu and Verena Rieser

Dialogue Act Classification in Team Communication for Robot Assisted Disaster Response

Tatiana Anikina and Ivana Kruijff-Korbayova

Multi-Task Learning of System Dialogue Act Selection for Supervised Pretraining of Goal-Oriented Dialogue Policies

Sarah McLeod, Ivana Kruijff-Korbayova and Bernd Kiefer

B. Rex: a dialogue agent for book recommendations

Mitchell Abrams, Luke Gessler and Matthew Marge

13 September 2019 (continued)

14:20–14:50 *Coffee Break (during Poster and Demos 3)*

14:50–16:05 **Session 7 - Discourse**

SpaceRefNet: a neural approach to spatial reference resolution in a real city environment

Dmytro Kalpakchi and Johan Boye

Which aspects of discourse relations are hard to learn? Primitive decomposition for discourse relation classification

Charlotte Roze, Chloé Braud and Philippe Muller

Discourse Relation Prediction: Revisiting Word Pairs with Convolutional Networks

Siddharth Varia, Christopher Hidey and Tuhin Chakrabarty

16:05–16:20 *Short Coffee Break*

16:20–17:20 *Business meeting, Awards and Closing*

Keynote Abstracts

Keynote 1 - Learning Natural Language Interfaces with Neural Models

Mirella Lapata

Professor, School of Informatics, University of Edinburgh, UK

Abstract

In Spike Jonze's futuristic film "Her", Theodore, a lonely writer, forms a strong emotional bond with Samantha, an operating system designed to meet his every need. Samantha can carry on seamless conversations with Theodore, exhibits a perfect command of language, and is able to take on complex tasks. She filters his emails for importance, allowing him to deal with information overload, she proactively arranges the publication of Theodore's letters, and is able to give advice using common sense and reasoning skills.

In this talk I will present an overview of recent progress on learning natural language interfaces which might not be as clever as Samantha but nevertheless allow users to interact with various devices and services using every day language. I will address the structured prediction problem of mapping natural language utterances onto machine-interpretable representations and outline the various challenges it poses. For example, the fact that the translation of natural language to formal language is highly non-isomorphic, data for model training is scarce, and natural language can express the same information need in many different ways. I will describe a general modeling framework based on neural networks which tackles these challenges and improves the robustness of natural language interfaces.

Biography

Mirella Lapata is professor of natural language processing in the School of Informatics at the University of Edinburgh. Her research focuses on getting computers to understand, reason with, and generate natural language. She is the first recipient (2009) of the British Computer Society and Information Retrieval Specialist Group (BCS/IRSG) Karen Sparck Jones award and a Fellow of the Royal Society of Edinburgh. She has also received best paper awards in leading NLP conferences and has served on the editorial boards of the Journal of Artificial Intelligence Research, the Transactions of the ACL, and Computational Linguistics. She was president of SIGDAT (the group that organizes EMNLP) in 2018.

Keynote 2 - Situated Interaction

Dan Bohus *Senior Principal Researcher, Perception and Interaction Group, Microsoft Research, Redmond, Washington, US*

Abstract

Physically situated dialog is a complex, multimodal affair that goes well beyond the spoken word. When interacting with each other, people incrementally coordinate their actions to simultaneously resolve several different problems: they manage engagement, coordinate on taking turns, recognize intentions, and establish and maintain common ground as a basis for contributing to the conversation. A wide array of non-verbal signals are brought to bear. Proximity and body pose, attention and gaze, head nods and hand gestures, prosody and facial expressions, all play important roles in the intricate, mixed-initiative, fluidly coordinated process we call interaction. And just like a couple of decades ago advances in speech recognition opened up the field of spoken dialog systems, today advances in vision and other perceptual technologies are again opening up new horizons – we are starting to be able to build machines that can understand these social signals and the physical world around them, and begin to participate in physically situated interactions and collaborations with people.

In this talk, using a number of research vignettes from my work, I will draw attention to some of the challenges and opportunities that lie ahead of us in this exciting space. In particular, I will discuss issues with managing engagement and turn-taking in multiparty open-world settings, and more generally highlight the importance of timing and fine-grained coordination in situated interaction. Finally, I will conclude by describing a framework that promises to simplify the development of physically situated interactive systems and enable more research and faster progress in this area.

Biography

Dan Bohus is a Senior Principal Researcher in the Perception and Interaction Group at Microsoft Research. His work centers on the study and development of computational models for physically situated spoken language interaction and collaboration. The long term question that shapes his research agenda is how can we enable interactive systems to reason more deeply about their surroundings and seamlessly participate in open-world, multiparty dialog and collaboration with people? Prior to joining Microsoft Research, Dan obtained his Ph.D. from Carnegie Mellon University.

Keynote 3 - The Many Facets of Dialog

Helen Meng *SProfessor, Department of Systems Engineering and Engineering Management, Chinese University of Hong Kong, China*

Abstract

Dialog is a most fascinating form of human communication. The back-and-forth exchanges convey the speaker's message to the listener, and the listener can derive information about the speaker's thoughts, intent, well-being, emotions and much more. This talk presents an overview of dialog research that concerns our group at The Chinese University of Hong Kong. In the domain of education and learning, we have been recording in-class student group discussions in the flipped-classroom setting of a freshman elite mathematics course. We investigate features in the weekly, within-group dialogs that may relate to class performance and learning efficacy. In the domain of e-commerce, we are developing dialog models based on approximately 20 million conversation turns, to support a virtual shopping assistant in customer inquiries and orders, logistics tracking, etc. In the domain of health and wellbeing, we are capturing and analysing dialogs between health professionals (or their virtual equivalent) and subjects in cognitive screening tests. We also conduct research in both semantic interpretation and dialog state tracking, as well as affective design of virtual conversational assistants. For the former, we have developed a Convex Polytopic Model for extracting a knowledge representation from user inputs in dialog turns by generating a compact convex polytope to enclose all the data points projected to a latent semantic space. The polytope vertices represent extracted semantic concepts. Each user input can then be "interpreted" as a sequence of polytope vertices which represent the user's goals and dialog states. For the latter, we have developed a multimodal, multi-task, deep learning framework to infer the user's emotive state and emotive state change simultaneously. This enables virtual conversational assistants to understand the emotive state in the user's input and to generate an appropriate emotive system response in the dialog turn, which will further influence the user's emotive state in the subsequent dialog turn. Such an affective design will be able to enhance user experience in conversational dialogs with intelligent virtual assistants.

Biography

Helen Meng is Patrick Huen Wing Ming Professor of Systems Engineering and Engineering Management at The Chinese University of Hong Kong (CUHK). She is the Founding Director of the CUHK Ministry of Education (MoE)-Microsoft Key Laboratory for Human-Centric Computing and Interface Technologies (since 2005), Tsinghua-CUHK Joint Research Center for Media Sciences, Technologies and Systems (since 2006), and Co-Director of the Stanley Ho Big Data Decision Analytics Research Center (since 2013). Previously, she served as CUHK Faculty of Engineering's Associate Dean (Research), Chairman of the Department of Systems Engineering and Engineering Management, Editor-in-Chief of the IEEE Transactions on Audio, Speech and Language Processing, Member of the IEEE Signal Processing Society Board of Governors, ISCA Board Member and presently Member of the ISCA International Advisory Council. She was elected APSIPA's inaugural Distinguished Lecturer 2012-2013 and ISCA Distinguished Lecturer 2015-2016. Her awards include the Ministry of Education Higher Education Outstanding Scientific Research Output Award 2009, Hong Kong Computer Society's inaugural Outstanding ICT Woman Professional Award 2015, Microsoft Research Outstanding Collaborator Award 2016 (1 in 32 worldwide), IEEE ICME 2016 Best Paper Award, IBM Faculty Award 2016, HKPWE Outstanding Women Professionals and Entrepreneurs Award 2017 (1 in 20 since 1999), Hong Kong ICT Silver Award 2018 in Smart Inclusion, and the CogInfoComm2018 Best Paper Award. Helen received all her degrees from MIT. Her research interests include big data decision analytics, and artificial intelligence especially for speech and language technologies to support multilingual and multimodal human-computer interaction. Helen has given invited / keynote presentations including INTERSPEECH 2018 Plenary Talk, World Economic Forum Global Future Council 2018, Taihe Workshop on Building Stakeholder Networks on AI Ethics and Governance 2019 and the World Peace Forum 2019. She has served in numerous Government appointments, including

Chairlady of the Research Grants Council's Assessment Panel for Competitive Research Funding Schemes for the Local Self-financing Degree Sector, Chairlady of the Working Party on Manpower Survey of the Information/Innovation Technology Sector (since 2013), as well as Steering Committee Member of Hong Kong's Electronic Health Record (eHR) Sharing. Helen is a Fellow of HKCS, HKIE, IEEE and ISCA.