NILLI 2023

Novel Ideas in Learning-to-Learn through Interaction

Proceedings of the Workshop

December 7, 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 979-8-89176-057-8

Introduction

Collaborative dialogues [1, 2, 3] with automated systems through language interactions have become ubiquitous, wherein it is becoming common from setting an alarm to planning one's day through language interactions. With recent advances in dialogue research [1, 2, 3], embodied learning [4, 5, 6] and using language as a mode of instruction for learning agents [4, 7, 8] there is, now, a scope for realizing domains that can assume agents with primitive task knowledge and a continual interact-and-learn procedure to systematically acquire knowledge through verbal/non-verbal interactions [9, 10, 11, 12]. The research direction of building interactive learning agents [4, 7, 8, 13] facilitates the possibility of agents to have advanced interactions like taking instructions by being a pragmatic listener, asking for more samples, generating rationales for predictions, interactions to interpret learning dynamics, or even identifying or modifying a new task that can be used towards building effective learning-to-learn mechanisms. In a way, with verbal/non-verbal interactive medium this interdisciplinary field unifies research paradigms of lifelong learning, natural language processing, embodied learning, reinforcement learning, robot learning and multi-modal learning towards building interactive and interpretable AI.

https://www.cs.mcgill.ca/~pparth2/nilli_workshop_2023

References

[1] Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. Deep reinforcement learning for dialogue generation. Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, 2016.

[2] Mike Lewis, Denis Yarats, Yann Dauphin, Devi Parikh, and Dhruv Batra. Deal or no deal? end-to-end learning of negotiation dialogues. Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, 2017.

[3] Prasanna Parthasarathi and Joelle Pineau. Extending neural generative conversational model using external knowledge sources. Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, 2018.

[4] S. Chernova and Andrea Lockerd Thomaz. Robot learning from human teachers. In Synthesis Lectures on Artificial Intelligence and Machine Learning, 2014.

[5] Roberto Calandra, Andrew Owens, Dinesh Jayaraman, Justin Lin, Wenzhen Yuan, Jitendra Malik, Edward H. Adelson, and Sergey Levine. More than a feeling: Learning to grasp and regrasp using vision and touch. IEEE Robotics and Automation Letters, 3(4):3300-3307, Oct 2018.

[6] Cătălina Cangea, Eugene Belilovsky, Pietro Liò, and Aaron Courville. Videonavqa: Bridging the gap between visual and embodied question answering, 2019.

[7] Andrea Thomaz, Guy Hoffman, and Maya Cakmak. Computational human-robot interaction. Found. Trends Robot, 4(2-3):105-223, dec 2016.

[8] Stefanie Tellex, Nakul Gopalan, Hadas Kress-Gazit, and Cynthia Matuszek. Robots that use language. Annual Review of Control, Robotics, and Autonomous Systems, 3(1):25-55, 2020.

[9] Dzmitry Bahdanau, Felix Hill, Jan Leike, Edward Hughes, Arian Hosseini, Pushmeet Kohli, and Edward Grefenstette. Learning to understand goal specifications by modelling reward, 2018.

[10] Changan Chen, Unnat Jain, Carl Schissler, Sebastia Vicenc Amengual Gari, Ziad Al-Halah, Vamsi Krishna Ithapu, Philip Robinson, and Kristen Grauman. Soundspaces: Audio-visual navigation in 3d environments, 2019.

[11] Prasoon Goyal, Scott Niekum, and Raymond J. Mooney. Using natural language for reward shaping in reinforcement learning. Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, Aug 2019.

[12] Chuang Gan, Yiwei Zhang, Jiajun Wu, Boqing Gong, and Joshua B. Tenenbaum. Look, listen, and act: Towards audio-visual embodied navigation. 2020 IEEE International Conference on Robotics and Automation (ICRA), May 2020.

[13] Kalesha Bullard, Andrea Lockerd Thomaz, and S. Chernova. Towards intelligent arbitration of diverse active learning queries. 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 6049-6056, 2018.

Organizing Committee

Organizers

Prasanna Parthasarathi, Huawei
Koustuv Sinha, Meta AI
Khyathi Raghavi Chandu, Carnegie Mellon University
Chinnadhurai Sankar, SliceX AI
Marc-Alexandre Côté, Microsoft Research

Advisory Committee

Adina Williams, Meta AI
Sarath Chandar, Polytechnique de Montréal
Joelle Pineau, McGill University

Keynote Talk:

Yoav Artzi
Cornell University

Bio: Yoav Artzi is an Associate Professor in the Department of Computer Science and Cornell Tech at Cornell University. His research focuses on developing learning methods for natural language understanding and generation in automated interactive systems. He received an NSF CAREER award, and his work was acknowledged by awards and honorable mentions at ACL, EMNLP, NAACL, and IROS. Yoav holds a B.Sc. from Tel Aviv University and a Ph.D. from the University of Washington.

Keynote Talk:

Alison Smith-Renner
DataMinr

Bio: Alison Smith-Renner is a research scientist & engineer who designs, builds, and evaluates intelligent systems and interactive visualizations for data exploration, analysis, and augmented decision making. Alison's research lies at the intersection of AI/ML and human-computer interaction, building explainable and interactive AI/ML systems to engender trust, improve performance, and support human-machine collaboration.

Keynote Talk:

Linxi "Jim" Fan
NVIDIA AI

Bio: Jim Fan is a research scientist at NVIDIA AI. Jim's primary focus is to develop generally capable autonomous agents. To tackle this grand challenge, research efforts span foundation models, policy learning, robotics, multimodal learning, and large-scale systems. Jim obtained their Ph.D. degree at Stanford Vision Lab, advised by Prof. Fei-Fei Li. Previously, Jim did research internships at NVIDIA, Google Cloud AI, OpenAI, Baidu Silicon Valley AI Lab, and Mila-Quebec AI Institute. Jim was the Valedictorian of Class 2016 and a recipient of the Illig Medal at Columbia University.

Keynote Talk: Interacting with Code LLMs

Daniel Fried
Carnegie Mellon University

Bio: Daniel Fried is an assistant professor at the Language Technologies Institute in the School of Computer Science at Carnegie Mellon University, working on natural language processing. Daniel's work focuses on enabling people to use language to interact with computers to carry out useful tasks in the world. One recurring theme is pragmatics: viewing language as an action that people take in context to affect their communicative partners. Daniel is excited about domains where computers can complement human abilities, and recently, has been focusing on code generation, aiming to make programming more communicative.

Program

Thursday, December 7, 2023

08:30 - 08:35 *Opening Remarks*

08:35 - 09:20 *Invited Talk 1*

09:20 - 10:05 *Invited Talk 2*

10:05 - 10:50 *Invited Talk 3*

10:50 - 11:15 *Coffee Break (25 minutes)*

11:15 - 12:00 *Invited Talk 4*

12:00 - 13:30 *Lunch Break (90 minutes)*

13:30 - 15:30 *Lightning Talks (Session 1 - 12 talks)*

Improving Visually Grounded Continual Language Learning with Selective Module Specialization
Kyra Ahrens, Lennart Bengtson, Jae Hee Lee and Stefan Wermter

A Zero-Shot Language Agent for Computer Control with Structured Reflection
Tao Li, Gang Li, Zhiwei Deng, Bryan Wang and Yang Li

Task-Attentive Transformer Architecture for Continual Learning of Vision-and-Language Tasks Using Knowledge Distillation
Yuliang Cai, Jesse Thomason and Mohammad Rostami

Explainable Claim Verification via Knowledge-Grounded Reasoning with Large Language Models
Haoran Wang and Kai Shu

Does Listener Gaze in Face-to-Face Interaction Follow the Entropy Rate Constancy Principle: An Empirical Study
Yu Wang and Hendrik Buschmeier

PlugMed: Improving Specificity in Patient-Centered Medical Dialogue Generation using In-Context Learning
Chengfeng Dou, Zhi Jin, Wenpin Jiao, Haiyan Zhao, Yongqiang Zhao and Zhengwei Tao

Thursday, December 7, 2023 (continued)

Dior-CVAE: Diffusion Priors in Variational Dialog Generation
Tianyu Yang, Thy Thy Tran and Iryna Gurevych

GATE: Grounded Argument and Task Extraction for Embodied Agents
Chayan Sarkar, Avik Mitra, Pradip Pramanick and Tapas Nayak

Reasoning about Commonsense Norms on Defeasible Visual Cues and Moral Dilemmas
Seungju Han, Seungbeen Lee, Junhyeok Kim, Jack Hessel, Liwei Jiang, Jiwan Chung, Yejin Son, Yejin Choi and Youngjae Yu

CLIN: A Continually Learning Language Agent for Rapid Task Adaptation and Generalization
Bodhisattwa Prasad Majumder, Bhavana Dalvi Mishra, Peter Jansen, Oyvind Tafjord, Niket Tandon, Li Zhang, Chris Callison-Burch and Peter Clark

MathDial: A Dialogue Tutoring Dataset with Rich Pedagogical Properties Grounded in Math Reasoning Problems
Jakub Macina, Nico Daheim, Sankalan Pal Chowdhury, Tanmay Sinha, Manu Kapur, Iryna Gurevych and Mrinmaya Sachan

DialGuide: Aligning Dialogue Model Behavior with Developer Guidelines
Prakhar Gupta, Yang Liu, Di Jin, Behnam Hedayatnia, Spandana Gella, Sijia Liu, Patrick Lange, Julia Hirschberg and Dilek Hakkani-Tur

15:30 - 16:00 *Coffee Break (30 minutes)*

16:00 - 17:20 *Lightning Talks (Session 2 - 8 talks)*

RSVP: Customer Intent Detection via Agent Response Contrastive and Generative Pre-Training
Yu-Chien Tang, Wei Yao Wang, An Zi Yen and Wen Chih Peng

Long-Horizon Dialogue Understanding for Role Identification in the Game of Avalon with Large Language Models
Simon Stepputtis, Joseph Campbell, Yaqi Xie, Zhengyang Qi, Wenxin Sharon Zhang, Ruiyi Wbng, Sanketh Rangreji, Charles Michael Lewis and Katia P. Sycara

Large Language Models as Source Planner for Personalized Knowledge-grounded Dialogues
Hongru WANG, Minda Hu, Yang Deng, Rui Wang, Fei Mi, Weichao Wang, Yasheng Wang, Wai-Chung Kwan, Irwin King and Kam-Fai Wong

Time-Considerable Dialogue Models via Reranking by Time Dependency
Yuiko Tsunomori, Masakazu Ishihata and Hiroaki Sugiyama

Improving Conversational Recommendation Systems via Bias Analysis and Language-Model-Enhanced Data Augmentation
Xi Wang, Hossein A. Rahmani, Jiqun Liu and Emine Yilmaz

Multi-User MultiWOZ: Task-Oriented Dialogues among Multiple Users
Yohan Jo, Xinyan Zhao, Arijit Biswas, Nikoletta Basiou, Vincent Auvray, Nikolaos Malandrakis, Angeliki Metallinou and Alexandros Potamianos

STEER: Unified Style Transfer with Expert Reinforcement
Skyler Hallinan, Faeze Brahman, Ximing Lu, Jaehun Jung, Sean Welleck and Yejin Choi

Measuring the Knowledge Acquisition-Utilization Gap in Pretrained Language Models
Amirhossein Kazemnejad, Mehdi Rezagholizadeh, Prasanna Parthasarathi and Sarath Chandar

17:20 - 17:30 *Closing Remarks*