

NILLI 2022

Novel Ideas in Learning-to-Learn through Interaction

Proceedings of the Workshop

December 8, 2022

©2022 Association for Computational Linguistics

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 978-1-959429-17-3

Introduction

Interactive environments have played a pivotal role in the development of reasoning mechanisms in intelligent species. Recent advances in language generation, multimodal learning, interactive and embodied learning with using language as a mode of instruction for learning have increased focus to addressing challenges in this growing topic of research. In the horizon, there is scope for realizing scenarios where agents with primitive task knowledge and an interact-and-learn procedure to systematically acquire knowledge through its interactions with the environment. This Novel Ideas in Learning-to-Learn through Interaction (NILLI) workshop focuses on collecting discussions to improve clarity towards the challenges in this topic of research which require modeling sophisticated continual interactive agents across diverse tasks and mediums of interactions. This interdisciplinary research topic 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_2022

Organizing Committee

Organizers

Prasanna Parthasarathi, Huawei
Koustuv Sinha, McGill University
Kalesha Bullard, DeepMind
Chinnadhurai Sankar, Meta AI
Marc-Alexandre Côté, Microsoft Research

Advisory Committee

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

Program Committee

Reviewers

Luciana Benotti

Ta-Chung Chi

Ifrac Idrees

Danni Ma, Sneha Mehta, Nikita Moghe

Thao Nguyen

Yuanze Pang, Roma Patel

Janarthanan Rajendran, Eric Rosen, Josh Roy

Sashank Santhanam, Bishal Santra, Yow-Ting Shiue

Zian Wang

Eric Yuan

Keynote Talk:

Jakob Foerster

Oxford University

Bio: Jakob Foerster is an Associate Professor at the department of engineering science at the University of Oxford. During his PhD at Oxford he helped bring deep multi-agent reinforcement learning to the forefront of AI research and interned at Google Brain, OpenAI, and DeepMind. After his PhD he worked as a research scientist at Facebook AI Research in California, where he continued doing foundational work. He was the lead organizer of the first Emergent Communication workshop at NeurIPS in 2017, which he has helped organize ever since and was awarded a prestigious CIFAR AI chair in 2019.

Keynote Talk:

Jacob Andreas

MIT

Bio: Jacob Andreas is an assistant professor at MIT. His research aims to build intelligent systems that can communicate effectively using language and learn from human guidance. Jacob earned his Ph.D. from UC Berkeley, his M.Phil. from Cambridge (where he studied as a Churchill scholar) and his B.S. from Columbia. As a researcher at Microsoft Semantic Machines, he founded the language generation team and helped develop core pieces of the technology that powers conversational interaction in Microsoft Outlook. He has been the recipient of Samsung's AI Researcher of the Year award, MIT's Kolokotronis teaching award, and paper awards at NAACL and ICML.

Keynote Talk:

Maja Rita Rudolph
Bosch Center for AI

Bio: Maja Rudolph is a Senior Research Scientist at the Bosch Center for Artificial Intelligence, where she works in the field of deep probabilistic modeling. She completed her Ph.D. in Computer Science at Columbia University, advised by David Blei, in 2018. Maja holds an MS in Electrical Engineering from Columbia University and a BS in Mathematics from MIT.

Keynote Talk:

Karthik Narasimhan
Princeton University

Bio: Karthik Narasimhan is an assistant professor in the Computer Science department at Princeton University. His research spans the areas of natural language processing and reinforcement learning, with the goal of building intelligent agents that learn to operate in the world through both their own experience (doing things) and leveraging existing human knowledge (reading about things). Karthik received his PhD from MIT in 2017, and spent a year as a visiting research scientist at OpenAI, contributing to the GPT language model, prior to joining Princeton in 2018. He is the recipient of a Google Research Scholar Award (2022), an Amazon research award (2019) and best paper awards/nominations at EMNLP (2015, 2016).

Keynote Talk:

Joyce Chai

University of Michigan

Bio: Joyce Chai is a Professor in the Department of Electrical Engineering and Computer Science at the University of Michigan. Prior to joining UM in 2019, she was a Professor in Computer Science and Engineering at Michigan State University. She has over twenty years of research experience in the areas of natural language processing, situated and multimodal dialogue systems, human-robot communication, and artificial intelligence. Her recent work, with strong connections to psycholinguistics and cognitive science, explores intersections between language, vision, and robotics to enable human-agent communication. Her research has been supported by NSF, DARPA, and ONR. She holds a Ph.D. in Computer Science from Duke University.

Keynote Talk:

Mark Riedl

Georgia Institute of Technology

Bio: Dr. Mark Riedl is a professor in the Georgia Tech School of Interactive Computing and associate director of the Georgia Tech Machine Learning Center. Dr. Riedl's research focuses on human-centered artificial intelligence—the development of artificial intelligence and machine learning technologies that understand and interact with human users in more natural ways. Dr. Riedl's recent work has focused on story understanding and generation, computational creativity, explainable AI, and teaching virtual agents to behave safely.

Table of Contents

Program

Thursday, December 8, 2022

09:00 - 09:05 *Opening Remarks*

09:05 - 09:50 *Invited Talk 1*

09:50 - 10:30 *Lightning Talks (4)*

AERBIF: Actionable Entities Recognition Benchmark for Interactive Fiction

Ivan P. Yamshchikov and Alexey Tikhonov

Multimodal Contextualized Plan Prediction for Embodied Task Completion

Mert Inan, Aishwarya Padmakumar, Spandana Gella, Patrick Lange and Dilek Hakkani-Tur

Thompson sampling for interactive Bayesian optimization of dynamic masking-based language model pre-training

Iñigo Urteaga, Moulay-Zaidane Draïdia, Tomer Lancewicki and Shahram Khadivi

ReAct: Synergizing Reasoning and Acting in Language Models

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan and Yuan Cao

10:30 - 10:45 *Coffee Break*

10:45 - 11:15 *Lightning Talks (3)*

Joint Audio/Text Training for Transformer Rescorer of Streaming Speech Recognition

Suyoun Kim, Ke Li, Lucas Kabela, Rongqing Huang, Jiedan Zhu, Ozlem Kalinli and Duc Le

Revisiting the Roles of Text” in Text Games

Mo Yu, Josh Tenenbaum, Chuang Gan, Shunyu Yao and Yi Gu

ComFact: A Benchmark for Linking Contextual Commonsense Knowledge

Antoine Bosselut, Yuki Mitsufuji, Hiromi Wakaki, Saya Kanno, Jena D. Hwang and Silin Gao

11:15 - 12:00 *Invited Talk 2*

Thursday, December 8, 2022 (continued)

12:00 - 13:00 *Lunch Break*

13:00 - 13:45 *Invited Talk 3*

13:45 - 14:30 *Invited Talk 4*

14:30 - 15:40 *Lightning Talks (7)*

LEMON: Language-Based Environment Manipulation via Execution-Guided Pre-training

Jian-Guang LOU, Ting Liu, Yu Zhang, Bei Chen, Qian Liu and Qi Shi

Learn What Is Possible, Then Choose What Is Best: Disentangling One-To-Many Relations in Language Through Text-based Games

Ke Zhou and Benjamin Towle

Reason first, then respond: Modular Generation for Knowledge-infused Dialogue

Jason Weston, Arthur Szlam, Jack Urbanek, Kurt Shuster and Leonard Adolphs

Lexi: Self-Supervised Learning of the UI Language

Oriana Riva, Chitta Baral, Kushal Arora, Shweti Mahajan and Pratyay Banerjee

Context-aware Information-theoretic Causal De-biasing for Interactive Sequence Labeling

Ani Nenkova, Ricardo Henao, Shuai Li, Handong Zhao, Ruiyi Zhang, Tong Yu, Rui Wang and Junda Wu

Predicting Long-Term Citations from Short-Term Linguistic Influence

Jacob Eisenstein, David Bamman and Sandeep Soni

StuBot: Learning by Teaching a Conversational Agent Through Machine Reading Comprehension

Hana Lee and Nayoung Jin

15:40 - 16:00 *Coffee Break*

16:00 - 16:45 *Invited Talk 5*

Thursday, December 8, 2022 (continued)

16:45 - 17:30 *Invited Talk 6*

17:30 - 17:40 *Closing Remarks*

Author Index