

Learning through Interaction in Cooperative Multi-Agent Systems

Host: Marynel Vázquez



Wednesday, April 7, 2021
4:00 p.m.
Speaker: Kalesha Bullard

Zoom Presentation

Abstract: Effective communication is an important skill for enabling information exchange and cooperation in multi-agent systems, in which agents coexist in shared environments with humans and/or other artificial agents. Indeed, human domain experts can be a highly informative source of instructive guidance and feedback (supervision). My prior work explores this type of interaction in depth, as a mechanism for enabling learning for artificial agents. However, dependence upon human partners for acquiring or adapting skills has important limitations. Human time and cognitive load is typically constrained (particularly in realistic settings) and data collection from humans, though potentially qualitatively rich, can be slow and costly to acquire. Yet, the ability to learn through interaction with other agents represents another powerful mechanism for enabling interactive learning. Though other artificial agents may *also* be novices, agents can co-learn through providing each other evaluative feedback (reinforcement), given the learning task has been sufficiently structured and allows for generalization to novel settings.

This talk presents research that investigates methods for enabling agents to learn *general* communication skills through interactions with *other agents*. In particular, the talk will focus on my ongoing work within Multi-Agent Reinforcement Learning, investigating *emergent* communication protocols, inspired by communication in real-world problem settings. We present a novel problem setting and a general approach that allows for zero-shot coordination (ZSC), i.e., discovering protocols that can generalize to independently trained agents. We also explore and analyze specific difficulties associated with finding globally optimal ZSC protocols, as complexity of the communication task increases or the modality for communication changes (e.g. to implicit communication through physical movement, by a simulated robotic agent). Overall, this work opens up exciting avenues for learning *general* communication protocols in complex domains.

Bio: Kalesha Bullard is a postdoctoral researcher at Facebook AI Research. She completed her PhD in Computer Science at Georgia Institute of Technology in 2019, where her research focused within the space of interactive robot learning. During her postdoc, Kalesha has expanded her research to explore the space of multi-agent reinforcement learning, currently investigating how to enable embodied multi-agent populations to learn *general* communication protocols. More broadly, Kalesha's research interests span autonomous reasoning and decision making for artificial agents in multi-agent settings. To date, her research has focused on principled methods for enabling agents to learn through interaction with other agents (human or artificial) to achieve shared goals. Beyond research, Kalesha has participated in a number of service roles throughout her research career, recently serving on organizing and program committees for workshops associated with several top Artificial Intelligence conference venues (NeurIPS, AAAI, AAMAS). This past year, she was selected as one of the 2020 Electrical Engineering and Computer Science (EECS) Rising Stars.