

Exploiting latent structure and bisimulation metrics for better generalization in reinforcement learning

Host: Marynel Vázquez



Wednesday February 24, 2021

4:00 p.m.

Speaker: Amy Zhang

Zoom Presentation

Abstract:

The advent of deep learning has shepherded unprecedented progress in various fields of machine learning. Despite recent advances in deep reinforcement learning (RL) algorithms, however, there is no method today that exhibits anywhere near the generalization that we have seen in computer vision and NLP. Indeed, one might ask whether deep RL algorithms are even capable of the kind of generalization that is needed for open-world environments. This challenge is fundamental and will not be solved with incremental algorithmic advances.

In this talk, we propose to incorporate different assumptions that better reflect the real world and allow the design of novel algorithms with theoretical guarantees to address this fundamental problem. We first present how state abstractions can accelerate reinforcement learning from rich observations, such as images, without relying either on domain knowledge or pixel-reconstruction. Our goal is to learn state abstractions that both provide for effective downstream control and invariance to task-irrelevant details. We use bisimulation metrics to quantify behavioral similarity between states, and learn robust latent representations which encode only the task-relevant information from observations. We provide theoretical guarantees for the learned approximate abstraction and extend this notion to families of tasks with varying dynamics.

Bio:

I am a final year PhD candidate at McGill University and the Mila Institute, co-supervised by Profs. Joelle Pineau and Doina Precup. I am also a researcher at Facebook AI Research. My work focuses on bridging theory and practice through learning approximate state abstractions and learning representations for generalization in reinforcement learning. I previously obtained an M.Eng. in EECS and dual B.Sci. degrees in Mathematics and EECS from MIT.