



DISSERTATION DEFENSE

Jongwook Choi

Learning Representations for Efficient Exploration and Goal- Conditioned Reinforcement Learning

Thursday, May 30, 2024

4:00pm – 6:00pm

Virtual – [Zoom](#)

Passcode: UM-CSE



ABSTRACT: Deep reinforcement learning (RL) is a general-purpose computational framework for learning sequential decision-making agents, with the promise that agents can learn useful behaviors and solve the task through trial-and-error by maximizing rewards. One key fundamental problem in deep RL is representation learning — discovering and extracting useful or task-relevant information from raw data (e.g., observations and agent’s actions) that can make solving downstream tasks more efficient and tractable. In this dissertation, I propose and discuss several methods and principles pertaining to representation learning for RL, with a focus on state and temporal abstraction, enabling more efficient exploration, skill discovery, and learning of goal-conditioned policies for hierarchical agents.

First, I begin with a self-supervised approach to learn a state representation using the idea of contingency-awareness: the agent’s knowledge about which aspect of the environment is controllable. I present a novel attentive dynamics model that identifies controllable elements of the environment which can efficiently abstract the search space for exploration, and show that it enables strong exploration performance in difficult, hard-exploration Atari game environments featuring sparse rewards. Next, I discuss a novel perspective and foundation that unifies goal-conditioned RL and variational empowerment methods for unsupervised skill discovery based on the principle of mutual information maximization into a single family of methods. The proposed framework, variational goal-conditioned RL, allows us to interpret variational empowerment methods as a principled approach for learning latent goal representations and goal-reaching reward functions, while also enabling practical techniques and improvements brought from each other. Finally, I present another instance of skill learning for temporal abstraction: entity-centric skill learning in continuous control environments with multiple entities. By utilizing a structured goal representation and a novel intrinsic reward based on counterfactual reasoning and dynamics models, I demonstrate that one can learn pairwise object interaction behaviors without relying on any external rewards. Overall, this dissertation contributes to the advancement of deep RL by addressing state representation learning and skill learning problems, which can help build more autonomous systems for real-world problems with less human supervision.

CHAIR: Prof. Honglak Lee