COMPUTER SCIENCE & ENGINEERING
UNIVERSITY OF MICHIGAN

DISSERTATION DEFENSE

Wilka Carvalho

Discovering Structured Representations for Reinforcement Learning

Friday, July 7, 2023
3:00pm – 5:00pm
3941 Beyster
Hybrid – Zoom
Passcode: 453457

ABSTRACT: Deep reinforcement learning (Deep RL) has recently emerged as a powerful method for developing AI that can learn to select actions in the world. One key question in RL is how an agent should learn knowledge that can be transferred to new situations. In this dissertation, I hypothesize that one key to transferring knowledge is the ability to discover *structured representations* that permit relational reasoning over basic units describing the agent's experience. Recent research in computer vision and natural language processing has shown that structured neural networks with sparse and dynamic information flow enable the discovery of such structured representations, leading to faster learning and improved generalization. **The thesis of this dissertation** is that we can equip reinforcement learning agents with the ability to discover and exploit structured representations by incorporating structured neural networks with dynamic information into the core components of an RL learner. By equipping RL agents with the ability to discover structured representations, we can reduce the amount of experience the agent needs for learning and improve its ability to transfer behaviors across situations. To support this argument, I present the following evidence. **First**, I incorporate structured neural networks into an RL agent's state function. I assume access to an object detector (but no other ground-truth information) and show that we can discover useful object representations by learning a relational object-centric transition model over an (initially random) embedding space. This approach enables the discovery of object representations that capture an object's category, properties, and attributes, while achieving performance comparable to an agent with access to ground-truth object information. **Afterward**, I remove the assumption of having access to an object detector. I modify an RL agent's state function to consist of a set of neural network modules that dynamically share and update information from the observation. I demonstrate that this structured state function can discover object primitives that facilitate generalization to task variations across three diverse object-centric environments defined by language instructions, 3D objects, and object motions. **Next**, I incorporate structured neural networks into an agent's value function. Combining this with learning a predictive state representation known as successor features, I show that it enables the discovery of features that allow for generalization to combinations of tasks. **Finally**, I incorporate structured neural networks into an agent's policy. By doing so, I develop the first successor feature-based method to transfer to combinations of tasks in a 3D environment when all necessary representations are discovered. I demonstrate that our method transfers to new tasks with hundreds of millions fewer samples compared to the transfer learning baselines we compare against. **Taken together**, the contributions of this thesis demonstrate that incorporating structured neural networks into the core components of an RL learner can enable structured representation learning that both reduces the amount of experience an agent requires for learning and improves its ability to transfer behaviors across situations.

CHAIR: Prof. Satinder Singh Baveja and Prof. Honglak Lee