

Exploiting Structure In Reinforcement Learning

Tuesday, June 18, 2024 11:00 AM (1 hour)

Abstract: While reinforcement learning has achieved impressive success in applications such as game-playing and robotics, there is work yet to be done to make RL truly practical for optimizing policies for real-world systems. In particular, many systems exhibit clear structure that RL algorithms currently don't know how to exploit efficiently. As a result, domain-specific heuristics often dominate RL both in system performance and resource consumption. In this talk we will discuss types of structures that may arise in real-world systems, and approaches to incorporate such structure in the design of RL algorithms. Examples of structured MDPs include models that exhibit linearity with respect to a low dimensional representation, models that exhibit smoothness in the parameters or the trajectories, and models whose dynamics can be decomposed into exogenous versus endogenous components.

Session Classification: Keynote: Christina Lee Yu (Cornell University)