

Learning and Control in Countable State Spaces

Monday, June 17, 2024 11:00 AM (1 hour)

Abstract: We will consider policy optimization methods in reinforcement learning where the state space is countably infinite. The motivation arises from control problems in communication networks and matching markets. We consider an algorithm called Natural Policy Gradient (NPG), which is a popular algorithm for finite state spaces, and show three results in the context of countable state spaces: (i) in the case where perfect policy evaluation is possible, we show that standard NPG converges with a small modification; (ii) if the error in policy evaluation is within a factor of the true value function, we show that one can obtain bounds on the performance of the NPG algorithms; and (iii) we will discuss the ability of neural network-based function approximations to satisfy the condition in (ii) above.

Session Classification: Keynote: R. Srikant (University of Illinois Urbana-Champaign)