Approximate Planning and Learning for Partially Observed Systems

Aditya Mahajan
McGill University
Canada

Reinforcement learning (RL) provides a conceptual framework for designing agents which learn to act optimally in an unknown environment. RL has been successfully used in various applications ranging from robotics, industrial automation, finance, healthcare, and natural language processing. The success of RL is based on a solid foundation of combining the theory of exact and approximate Markov decision processes (MDPs) with iterative algorithms that are guaranteed to learn an exact or approximate action-value function and/or an approximately optimal policy. However, for the most part, the research on RL theory is focused primarily on systems with full state observations.

In various applications including robotics, finance, and healthcare, the agent only gets a partial observation of the state of the environment. In this talk, we describe a new framework for approximate planning and learning for partially observed systems. This framework is based on the notion of approximate information state. The talk will highlight the strong theoretical foundations of this framework, illustrate how many of the existing approximation results can be viewed as a special case of approximate information state, and provide empirical evidence which suggests that this approach works well in practice.