Adaptive Control and Intersections with Reinforcement Learning

Anuradha Annaswamy

Active-adaptive Control Laboratory
Department of Mechanical Engineering
Massachusetts Institute of Technology

Adaptive control and reinforcement learning based control are two different methods that are both commonly employed for the control of uncertain systems. Historically, adaptive control has excelled at real-time control of systems with specific model structures through adaptive rules that learn the underlying parameters while providing strict guarantees on stability, asymptotic performance, and learning. Reinforcement learning methods are applicable to a broad class of systems and are able to produce near-optimal policies for highly complex control tasks. This is often enabled by significant offline training via simulation or the collection of large input-state datasets. In both methods, the main approach used for updating the parameters is based on a parametrized policy and gradient descent-like algorithms. Related tools of analysis, convergence, and robustness in both fields have a tremendous amount of similarity as well.

This talk will examine the similarities and interconnections between adaptive control and reinforcement learning-based control. Concepts in stability, performance, and learning, common to both methods will be discussed. Building on the similarities in update laws and common concepts, new intersections and opportunities for improved algorithm analysis will be explored. Two specific examples of dynamic systems are used to illustrate the details of the two methods, their advantages, and their deficiencies. We will explore how these methods can be leveraged and integrated to lead to provably correct methods for learning in real-time with guaranteed fast convergence. Examples will be drawn from a range of engineering applications.