Michael Muehlebach

On the Sample-Complexity of Online Reinforcement Learning: Packing, Priors, and Pontryagin

The talk studies the sample complexity of online reinforcement learning in the general setting of nonlinear dynamical systems with continuous state and action spaces. The analysis accommodates a large class of dynamical systems ranging from a finite set of nonlinear candidate models to models with bounded and Lipschitz continuous dynamics, to systems that are parametrized by a compact and real-valued set of parameters. We derive policy regret guarantees for each scenario and recover earlier results that were exclusively derived for linear time-invariant dynamical systems. The last part of the talk adopts a broader point of view and discusses the relation between feedforward and feedback controllers. We conclude by highlighting applications to real-world robotic systems including autonomous racing, magnetic navigation systems, and a table tennis playing robot that is actuated by pneumatic artificial muscles.