

Scalable Learning in Weakly Coupled Markov Decision Processes

Monday, June 17, 2024 4:30 PM (30 minutes)

We explore a general reinforcement learning framework within a Markov decision process (MDP) consisting of a large number N of independent sub-MDPs, linked by global constraints. In the non-learning scenario, when the model meets a specific non-degenerate condition, efficient algorithms (i.e., polynomial in N) exist, achieving a performance gap smaller than \sqrt{N} relative to the linear program upper bound. Analyzing the learning scenario in relation to this upper bound forms the central topic of this work.

Presenter: YAN, Chen

Session Classification: Parallel session: Online learning in stochastic networks