

Lipschitz Lifelong Reinforcement Learning: transferring value functions across MDPs

Friday, June 21, 2024 9:30 AM (1 hour)

Abstract: How close are the optimal value functions of two Markov decision processes that share the same state and action spaces but have different dynamics and rewards? In this talk, we will consider the problem of knowledge transfer when an agent is facing a series of reinforcement learning (RL) tasks. We will introduce a novel metric between Markov decision processes (MDPs) and establish that close MDPs have close optimal value functions. These theoretical results lead us to a value-transfer method for Lifelong RL, which we use to build a PAC-MDP algorithm with improved convergence rate. Beyond value transfer, this talk will open up on challenges and opportunities deriving from such an analysis.

Session Classification: Keynote: Emmanuel Rachelson (ISAE-SUPAERO)