Bo Dai

Title: The Journey of Making Linear MDP Practical

Abstract: It is common to address the curse of dimensionality in Markov decision processes (MDPs) by exploiting low-rank representations. This motivates much of the recent theoretical study on linear MDPs or low-rank MDPs, which enjoy good theoretical properties. However, most approaches require a given representation under unrealistic assumptions about the normalization of the decomposition or introduce unresolved computational challenges in practice.

We first demonstrate the linear MDPs structure exists in a generic stochastic nonlinear dynamic, but has not exploited yet. Then, we consider several efficient representation learning for reinforcement learning. The framework also admits confidence-adjusted index algorithms, enabling an efficient and principled approach to incorporating optimism or pessimism in the face of uncertainty. Finally, we further extend the structure for partially observable MDPs with practical algorithm. This journey leads to a series of algorithms with rigorously theoretical guarantees, meanwhile, achieving strong empirical performances.