Gergely Neu

Distances for Markov chains from sample streams

Bisimulation metrics are powerful tools for measuring similarities between stochastic processes, and specifically Markov chains. Recent advances have uncovered that bisimulation metrics are, in fact, optimal-transport distances, which has enabled the development of fast algorithms for computing such metrics with provable accuracy and runtime guarantees. However, these recent methods, as well as all previously known methods, assume full knowledge of the transition dynamics. This is often an impractical assumption in most real-world scenarios, where typically only sample trajectories are available. In this work, we propose a stochastic optimization method that addresses this limitation and estimates bisimulation metrics based on sample access, without requiring explicit transition models. Our approach is derived from a new linear programming (LP) formulation of bisimulation metrics, which we solve using a stochastic primal-dual optimization method. We provide theoretical guarantees on the sample complexity of the algorithm and validate its effectiveness through a series of empirical evaluations.