

Robust Markov Decision Processes with Data-driven, Distance-based Ambiguity Sets

Dr. Archis Ghatge

Professor & Associate Chair

Department of Industrial and Systems Engineering

University of Washington

Abstract: We will consider finite- and infinite-horizon Markov decision processes (MDPs) where the decision-maker does not know the true state-transition probabilities. The decision-maker assumes that they belong to certain ambiguity sets, and chooses actions that maximize the worst-case expected total discounted reward over all transition probabilities from these sets. We will work under a rectangular setup wherein the ambiguity set for the whole problem is a Cartesian product of ambiguity sets for individual state-action pairs across stages. Specifically, the ambiguity set for any state-action-stage triplet is a ball --- it includes all probability mass functions (pmfs) within a certain distance from an empirical transition pmf. This empirical pmf is constructed using historical, independent observations of the next state reached from each state-action pair in each stage. We will show that the optimal values of the resulting robust MDPs (RMDPs) converge to the optimal value of the true MDP, if the radii of the ambiguity balls vanish to zero as the sample-size diverges to infinity. We will also establish that the robust optimal value provides a lower bound on the value of the robust optimal policy in the true MDP, with a high probability. These two results rely on two particular sufficient conditions. We will discuss examples to illustrate that these sufficient conditions cannot be dropped from the hypotheses of our results. Several commonly used distances, including Total Variation, Burg, Hellinger, and Wasserstein, satisfy these two conditions. Moreover, the inner-optimization problems within the robust Bellman's equations can be solved efficiently for many well-known distances. This is joint work with ISE PhD student Siva Ramani.

Bio: Archis is a Professor in the Department of Industrial & Systems Engineering at the University of Washington in Seattle. He joined the University of Washington as an Assistant Professor in 2006 after receiving a PhD in Industrial and Operations Engineering from the University of Michigan in 2006, and an MS in Management Science and Engineering from Stanford in 2003. He completed his undergraduate education at the Indian Institute of Technology, Bombay, India, in 2001.