Industrial & Systems Engineering
Seminar Announcement

Information Theoretic Decision Making in Markov Decision Processes with Parametric Uncertainty

Peeyush Kumar, PhD Candidate
Industrial & Systems Engineering
University of Washington

Abstract: In this talk, I will discuss our research on Markov decision processes (MDP) where the state transition probabilities are parameterized and the decision-maker does not know the true value of this parameter. The decision-maker begins with a prior belief distribution on this parameter, and pursues a Bayesian approach to update this belief as the system evolves. To maximize expected reward over the planning horizon, the decision-maker must balance the exploration versus exploitation tradeoff: select a broad variety of actions to learn the transition probabilities sufficiently well, and utilize this information to quickly zero-in on actions with high rewards. We propose Information Directed Policy Sampling (IDPS), which is an information theoretic framework that explicitly manages this tradeoff. We obtain a worst-case regret bound for IDPS by applying a generalization of Hoeffding's inequality to Markov chains. The regret per stage converges to zero as the planning horizon increases to infinity; IDPS is thus asymptotically optimal. These theoretical guarantees are supplemented with numerical results on a sequential auction-design problem, and a response-guided dosing problem. Time permitting, I will also present generalizations to MDPs with hierarchical parametric uncertainty and to partially observable MDPs.

Bio: *Peeyush Kumar* is the co-founder and Chief Technology Officer of Cohort Intelligence, an Artificial Intelligence healthcare company enabling better decision making for physicians and nurses. He is also currently a Ph.D. candidate at the Department of Industrial and Systems Engineering, University of Washington Seattle working with Prof. Archis Ghate. His broader research area is in data driven decision making under uncertainty, particularly, in the application to healthcare decision making. His research interest lie in the field of operations research, machine learning, healthcare analytics, dynamic systems under uncertainty, stochastic processes, and clinical language processing.

Tuesday, March 27, 2018
1:30 – 2:20 p.m.
MEB 235