University of California, Irvine
Statistics Seminar

Reinforcement Learning under Unmeasured Confounding

Zhengling Qi
Assistant Professor of Decision Sciences
George Washington University

4 p.m., Thursday, Nov. 10, 2022
6011 Donald Bren Hall

In practical reinforcement learning (RL), a representation of the full state which makes the system Markovian and therefore amenable to most existing RL algorithms is not known a priori. Decision makers are often facing so-called partial observability of the state information, which significantly hinders the task of RL. Motivated by recent advances in causal inference, we study batch RL in the face of unmeasured confounders using auxiliary variables. A number of non-parametric identification results are established, based on which several promising policy optimization algorithms are proposed with finite-sample regret guarantees. Further, if time permits, I will discuss the phenomenon named "blessing from experts" and introduce the framework of super reinforcement learning in the batch setting.