



THE HONG KONG UNIVERSITY OF SCIENCE & TECHNOLOGY

Department of Mathematics

SEMINAR ON STATISTICS AND DATA SCIENCE

Towards Optimal Sample Complexities in Offline Reinforcement Learning and Markov Games

By

Prof. Yuxin CHEN

University of Pennsylvania

Abstract

Emerging reinforcement learning (RL) applications necessitate the design of sample-efficient solutions in order to accommodate the explosive growth of problem dimensionality. Despite the empirical success, however, our understanding about the statistical limits of RL remains highly incomplete. In this talk, I will present some recent progress towards settling the sample complexity in two RL scenarios. The first one is concerned with offline or batch RL, which performs learning using only pre-collected data without further exploration. We prove that model-based offline RL --- a plug-in approach that leverages the pessimism principle with Bernstein-style penalty --- achieves minimal-optimal sample complexity without any burn-in cost. The second scenario is concerned with multi-agent RL in zero-sum Markov games, assuming access to a generative model (a.k.a. simulator). We develop a new algorithm --- built upon the integration of adaptive sampling, online learning, and the optimism principle --- that overcomes the curse of multi-agents and the barrier of long horizon simultaneously. Our results emphasize the prolific interplay between high-dimensional statistics, online learning, and game theory. (See <https://arxiv.org/abs/2204.05275> and <https://arxiv.org/abs/2208.10458> for more details).

This is based on joint work with Gen Li, Laixi Shi, Yuling Yan, Yuejie Chi, Jianqing Fan, and Yuting Wei

Biography

Yuxin Chen is currently an associate professor in the Department of Statistics and Data Science at the University of Pennsylvania. Before joining UPenn, he was an assistant professor of electrical and computer engineering at Princeton University. He completed his Ph.D. in Electrical Engineering at Stanford University, and was also a postdoc scholar at Stanford Statistics. His current research interests include high-dimensional statistics, nonconvex optimization, and reinforcement learning. He has received the Alfred P. Sloan Research Fellowship, the ICCM best paper award (gold medal), the AFOSR and ARO Young Investigator Awards, the Google Research Scholar Award, and was selected as a finalist for the Best Paper Prize for Young Researchers in Continuous Optimization. He has also received the Princeton Graduate Mentoring Award.

Date : 7 December 2022 (Wednesday)

Time : 10:30am – 11:30am

Zoom Meeting : <https://hkust.zoom.us/j/95008648547> (Passcode: hkust)

All are Welcome!