


Department of Statistics Seminar Series



Jinhang Chai
**Department of Operations Research
and Financial Engineering**
Princeton University

Deep Transfer Q-Learning for Offline Non-Stationary Reinforcement Learning

Wednesday, November 19th, 2025

12:10-1:30 PM

Hill 552

Zoom Meeting: Meeting ID: 983 2864 8232

Password: 665123

<https://rutgers.zoom.us/j/98328648232?pwd=gcHifQ9OouKZnUY0EHLe8sjNXaV1oj.1>

Light refreshments will be served in Hill 452, 11:40am

Abstract:

In dynamic decision-making scenarios across business and healthcare, leveraging sample trajectories from diverse populations can significantly enhance reinforcement learning (RL) performance for specific target populations, especially when sample sizes are limited. While existing transfer learning methods primarily focus on linear regression settings, they lack direct applicability to reinforcement learning algorithms. This paper pioneers the study of transfer learning for dynamic decision scenarios modeled by non-stationary finite-horizon Markov decision processes, utilizing neural networks as powerful function approximators and backward inductive learning. We demonstrate that naive sample pooling strategies, effective in regression settings, fail in Markov decision processes. To address this challenge, we introduce a novel “re-weighted targeting procedure” to construct “transferable RL samples” and propose “transfer deep Q-learning”, enabling neural network approximation with theoretical guarantees. We assume that the reward functions are transferable and deal with both situations in which the transition densities are transferable or non-transferable. Our analytical techniques for transfer learning in neural network approximation and transition density transfers have broader implications, extending to supervised transfer learning with neural networks and domain shift scenarios. Empirical experiments on both synthetic and real datasets corroborate the advantages of our method, showcasing its potential for improving decision-making through strategically constructing transferable RL samples in non-stationary reinforcement learning contexts.

Bio:

Jinhang Chai is a fifth-year PhD student in the Department of Operations Research and Financial Engineering at Princeton University, working with Professor Jianqing Fan. Before that, he received a bachelor's degree from School of Mathematical Sciences at Peking University, with a double major in Economics. His research interest centers on statistical decision-making and high-dimensional statistics.