THE HONG KONG UNIVERSITY OF SCIENCE & TECHNOLOGY

Department of Mathematics

MATH/IEDA JOINT SEMINAR

Towards a Unified Theory for Continuous-Time Reinforcement Learning

By

Dr. Yanwei JIA
Columbia University

Abstract

While the existing reinforcement learning (RL) theory and algorithms have been predominantly developed for discrete-time Markov decision processes (MDPs), real-world problems abound in which RL agents can and indeed need to interact with environments in ultra-high frequency or outright continuously in time. Algorithms devised for MDPs often become unstable in the high-frequency regime and there seems to lack a rigorous and overarching theoretical foundation to guide the design and analysis of RL algorithms in this case. In this talk, I will report a series of recent work that endeavor to fill this gap. The research is characterized by a unified martingale perspective to study RL in continuous time and spaces under the entropy-regularized exploratory formulation. Through this martingale lens, we can recover and interpret many well-known RL algorithms, but more importantly can motivate new problems along with their solutions. The first part of the talk introduces the martingale characterization of policy evaluation and temporal-difference errors as a fundamental step for model-free learning. The second part explains why policy gradient is mathematically a policy evaluation problem. The third part presents the continuous-time counterpart of the Q-learning theory. As the conventional Q-function collapses in continuous time, we consider its first-order approximation and coin the term "(little) q-function". We jointly characterize the q-function and the value function by martingale conditions and, consequently, policy evaluation algorithms can be applied. The last part reports the results of extensive experiments on applying our RL algorithms to financial asset allocation, including simulations, backtesting and real-time portfolios that demonstrate outstanding performance of the resulting data-driven investment strategies.

Biography

Dr. Yanwei Jia is an associate research scientist and adjunct assistant professor in the Department of Industrial Engineering and Operations Research at Columbia University. He obtained Ph.D. from National University of Singapore in 2020. His research interest falls broadly into financial engineering and decision making problems, focusing on FinTech and data analytics. His recent research aims to develop fundamental theory on continuous-time reinforcement learning, and to solve problems in financial engineering, such as asset allocation and algorithmic trading.

Date : **18 August 2023 (Friday)**
Time : **2:00pm**
Venue : **Room 4472 (Lifts 25/26)**

All are Welcome!