THE HONG KONG POLYTECHNIC UNIVERSITY
香港理工大學

DEPARTMENT OF APPLIED MATHEMATICS
應 用 數 學 系

The Hong Kong Polytechnic University
Department of Applied Mathematics

Seminar

Demystifying the efficiency of reinforcement learning: A statistical perspective

By

Dr. Gen Li
The Chinese University of Hong Kong

Abstract

Reinforcement learning (RL) is frequently modeled as learning and decision making in a Markov decision process (MDP). A core objective of RL is to search for a policy — based on a collection of noisy data samples — that approximately maximizes expected cumulative rewards in an MDP, without direct access to a precise description of the underlying model. In contemporary applications, it is increasingly more common to encounter environments with prohibitively large state and action space, thus exacerbating the challenge of collecting enough samples to learn the model. In this talk, we present three recent works to show how to break the sample size barrier in reinforcement learning, including offline reinforcement learning, reward-agnostic exploration, and multi-agent Markov game. These results might shed light on the efficacy of these algorithms in more complicated scenarios.

Date: 29 August 2023 (Tuesday)
Time: 16:00-17:00 (Hong Kong Standard Time GMT +8)
Venue: FJ303
Speaker: Dr. Gen Li, The Chinese University of Hong Kong
Host: Dr. Ruijian Han, The Hong Kong Polytechnic University

*** * * ALL ARE WELCOME * * ***