




北京大学前沿计算研究中心
Center on Frontiers of Computing Studies, Peking University

静园5号院
前沿讲座




Rethinking the theoretical foundation of reinforcement learning



Prof. Nan Jiang

School of Computer Science
University of Illinois at Urbana-Champaign

🎤 Host: 王若松 助理教授
🕒 2026年7月21日 星期二 3:00pm
📍 静园五院204室



Abstract

Given two candidate functions, can we identify which one is the true value function of a large Markov decision process (MDP), given a "benign" dataset? Trivial as it might seem, a version of the question was open for 20+ years in reinforcement learning (RL), and the core difficulties are intimately related to the training instability of modern deep RL. In this talk, I will argue that by rethinking fundamental questions like this, RL theory can provide unique perspectives and solutions to practically relevant problems that are critical to the deployment of RL in real-world scenarios.

Biography

Nan Jiang is an Associate Professor of Computer Science at University of Illinois at Urbana-Champaign. Prior to joining UIUC, he was a postdoc researcher at Microsoft Research NYC. He received his PhD in Computer Science and Engineering at University of Michigan. His research focuses on the theory of reinforcement learning. He coauthors a monograph on RL theory and holds editorial positions in the research community, including Action Editor for JMLR, Editor for FnT in ML, and Senior Area Chairs for ICML and ICLR. His contributions are recognized by Best Paper Award in AAMAS 2015, Outstanding Paper Runner-up in ICML 2022, Adobe Data Science Award in 2021, NSF CAREER Award in 2022, Google Research Scholarship in 2024, and Sloan Research Fellowship in 2024.