NJIT | INSTITUTE FOR DATA SCIENCE

Data Science Seminar Series

In Collaboration with
The Department of Data Science

Resilient and (Near)-Optimal Algorithms for Reinforcement Learning and Nonconvex Optimization

Hosted by Mengjia Xu

Yi Zhou

University of Utah

Date: Monday, March 25, 2024
Time: 2:30 PM - 3:30 PM (Coffee served at 2:15 PM)
Location: GITC Building Room 3600 (3rd floor Seminar Room)
Web Link: Zoom Meeting Link

In modern machine learning, reinforcement learning and nonconvex optimization stand out as pivotal themes, relying on big data and stochastic algorithms to train models with optimized performance. Recent trends underscore the growing need for resilient models to, e.g., improve decision-making in both cooperative and competitive environments and enhance model robustness to data distribution shift. This talk will present resilient stochastic algorithms for addressing multi-agent reinforcement learning (MARL) and nonconvex optimization problems. Specifically, In the context of MARL with cooperative agents, a fully decentralized algorithm will be presented to learn the optimal policy with near-optimal computation and communication complexities. Transitioning to MARL with competitive agents, a decentralized primal-dual algorithm framework will be introduced to learn correlated equilibrium with near-optimal computational efficiency. Lastly, a new class of generalized-smooth nonconvex problems will be introduced, which covers broad machine learning applications with resiliency requirements. Subsequently, two optimal gradient-based algorithms will be presented to solve these problems efficiently. The talk aims to envision a foundational and comprehensive complexity theory to resilient machine learning algorithms.

Yi Zhou is an assistant professor affiliated with the Department of ECE at the University of Utah. Before that, he worked as a post-doctorate research associate in the Department of ECE at Duke University. He obtained a Ph.D. in ECE from The Ohio State University in 2018. His research interests include deep learning, reinforcement learning, statistical machine learning, nonconvex and distributed optimization, and statistical signal processing. His work on convergence analysis of cubic regularization under metric regularity is selected as spotlight paper at NeurIPS 2018. He is a recipient of the 2023 NSF Career Award.