2025 IE Summer Internship Project

Project title	Reinforcement Learning - Algorithm and Convergence Guarantees
Project description	Reinforcement learning (RL) has become increasingly impactful in solving sequential decision-making problems, from AlphaGo to recent large language models. However, its reliance on heuristics, the computational challenges posed by the curse of dimensionality, and the complexities arising from multi-agent interactions underscore the need for rigorous theoretical foundations, which lie at the core of my research. One of the most practical RL algorithms is the actor-critic framework, where the actor is responsible for policy improvement and the critic for policy evaluation. However, unlike typical value-based algorithms such as variance-reduced Q-learning (which has been shown to achieve minimax optimal sample complexity), policy-space algorithms such as natural actor-critic are far from theoretically optimal—particularly when implemented in a two-timescale manner rather than a two-loop manner. The goal of this project is to achieve minimax optimal sample complexity with (natural) actor-critic algorithms, possibly through improved algorithm design or advanced analysis techniques.
Host professor name	Zaiwei Chen
Professor/lab websites	https://sites.google.com/view/zaiweichen/home
Contact information to whom applicants can direct questions	For students interested in the opportunity, feel free to email chen5252@purdue.edu with any questions.
Other comments	Potential applicants should possess a solid mathematical background in analysis and probability. Prior experience with Markov decision processes and reinforcement learning is preferred but not required.

PURDUE UNIVERSITY | Edwardson School of Industrial Engineering