Learning & Optimization in Multiagent Decision-Making Systems

Instructor: Rasoul Etesami (etesami1@illinois.edu)

Seminar course to be offered at Tsinghua University in Summer 2025.

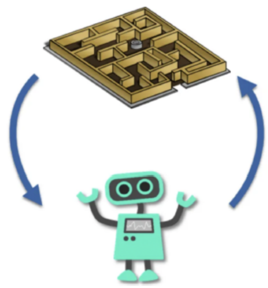

Course Description: In recent years, a broad array of studies on learning in multiagent systems has emerged across diverse fields, including operations research, computer science, engineering, and applied mathematics. On one hand, due to the lack of centralized coordination or information in many applications, distributed multiagent decision systems have become the focal point of numerous new areas, such as the growing popularity of online social networks, the analysis of large-scale network data, and challenges arising from interactions among agents in wireless networks, smart grids, formation control, robotic rendezvous, and socioeconomic systems. On the other hand, thanks to recent advances in Artificial Intelligence, it is now possible to obtain satisfactory solutions to many complex problems where classical approaches either fail or yield weak results.

This seminar course aims to explore recent advances in the analysis of multiagent decision-making systems—both with and without conflicting objectives—using optimization and learning techniques. The course will investigate such systems through the lens of optimization methods, game theory, and machine learning, with an emphasis on theoretical and algorithmic developments.

The course will begin with foundational background on single-agent nonlinear static optimization and its extension to dynamic stochastic settings, including Markov Decision Processes, the Minimum Principle, Bellman Optimality Conditions, and Policy Gradient Methods. Building upon the static single-agent setting, the course will cover topics in cooperative multiagent optimization—such as consensus over networks—as well as noncooperative multiagent decision-making, such as static games. From there, the course will move on to dynamic single-agent optimization and its multiagent counterparts, including dynamic stochastic games.

Finally, the course will address recent learning-based approaches for solving static and dynamic decision-making problems in both single-agent and multiagent systems when system parameters are unknown. These approaches include Regret Minimization, Reinforcement Learning, and learning Nash equilibria in static and dynamic games. To the extent possible, each topic will be motivated by practical applications drawn from engineering and cyber-physical systems.