Analysis Seminar

Reinforcement Learning for Mean-field Games with Average Reward Criterion

By

Kaan Alkır
(Bilkent)

Abstract: We study the inverse reinforcement learning (IRL) problem for discrete-time, infinite-horizon mean-field games (MFG) with an average-reward criterion. Unlike the forward setting, where the reward function is known, IRL assumes access only to expert demonstrations that are optimal under some unknown reward function. The objective is to recover the reward structure that explains these expert behaviors. Our approach is based on the maximum causal entropy principle, which selects the least biased policy among those consistent with the observed demonstrations. We show that the resulting non-convex formulation is equivalently reformulated as a convex optimization problem over occupation measures. Furthermore, we establish that the dual objective is smooth and strongly convex over compact sets and derive a variational representation using a log-partition formulation. Finally, we propose a first-order algorithm for solving the dual problem and recovering an entropy-maximizing equilibrium policy.

Date: Tuesday, April 22, 2025
Time: 18:00 – 19:00
Place: ZOOM

To request the event link, please send a message to goncha@fen.bilkent.edu.tr