Learning Distributed Equilibria in Linear-Quadratic Stochastic Differential Games: An α-Potential Approach

Philipp Plank[a] *and Yufei Zhang*[a]

[a] Imperial College London

Multi-agent reinforcement learning has shown impressive practical success in solving dynamic games, with applications ranging from autonomous driving and robotics to neuroscience and economics. Despite this empirical progress, theoretical understanding of the convergence of these learning algorithms in dynamic and stochastic environments remains limited.
In this talk, we take a first step toward bridging the gap by studying a fundamental class of N-player dynamic games: linear-quadratic stochastic differential games. We focus on independent policy-gradient learning, where each player employs a distributed policy that depends only on its own state and updates its policy independently by following the gradient of its own objective. We show that the game admits an α-potential structure, where α is determined by the degree of pairwise interaction asymmetry. We further characterize an affine distributed equilibrium beyond the mean field setting for both symmetric and asymmetric interactions. For pairwise-symmetric interactions, we show that independent policy-gradient methods converge globally to an affine equilibrium at a linear rate. For asymmetric interactions, we prove that independent projected policy-gradient algorithms converge linearly to an approximate equilibrium, with the suboptimality proportional to the degree of asymmetry.