

Quantifying the likelihood of collusion by provably convergent reinforcement learning

Thursday, June 20, 2024 2:30 PM (30 minutes)

Recent advances in decentralized multiagent reinforcement learning (MARL) have led to the development of algorithms that are provably convergent in a variety of Markov game subclasses. One of these is the Decentralized Q-learning (DQ) algorithm by Arslan and Yüksel (2017) which is provably convergent in weakly acyclic games. In this talk, I will present a new characterization of weak acyclicity and use it to show that the prisoner's dilemma with a memory of one period is weakly acyclic. This new characterization naturally leads to an identification of the basins of attraction of all possible strategy equilibria of the DQ algorithm. Since only a subset of strategy equilibria leads to robust collusion, we can use this to quantify the likelihood of observing algorithmic collusion. In addition, I will discuss the effect that fluctuations in the learning process and the addition of a third intermediate action to the prisoner's dilemma have on the likelihood of collusion.

Primary author: MEYLAHN, Janusz (University of Twente)

Presenter: MEYLAHN, Janusz (University of Twente)

Session Classification: Parallel session: Algorithmic collusion: Foundations for understanding the emergence of anticompetitive behaviour