

Davide Grossi

Learning to Cooperate under Uncertain Incentive Alignment

We apply a multi-agent reinforcement learning approach to study the emergence of cooperation among reinforcement learning agents that interact in a game. Crucially, agents are however uncertain about the extent to which their incentives are aligned, or misaligned, in the game. In this framework, we explore through computational experiments under what conditions further features of the model (like communication, reputation mechanisms, varying risk attitudes) may support learning of more cooperative policies.

This is joint work with: Nicole Orzan (University of Groningen), Erman Acar (University of Amsterdam) and Roxana Radulescu (Utrecht University).