Subsidy design for better social outcomes

Maria-Florina Balcan, Matteo Pozzi, Dravyansh Sharma

Overcoming the impact of selfish behavior of rational players in multiagent systems is a fundamental problem in game theory. Without any intervention from a central agent, strategic users take actions in order to maximize their personal utility, which can lead to extremely inefficient overall system performance, often indicated by a high Price of Anarchy. Recent work (Lin et al. 2021) investigated and formalized yet another undesirable behavior of rational agents, that of avoiding freely available information about the game for selfish reasons, leading to worse social outcomes. A central planner can significantly mitigate these issues by injecting a subsidy to reduce certain costs associated with the system and obtain net gains in the system performance. Crucially, the planner needs to determine how to allocate this subsidy effectively.

We formally show that designing subsidies that perfectly optimize the social good, in terms of minimizing the Price of Anarchy or preventing the information avoidance behavior, is computationally hard under standard complexity theoretic assumptions. On the positive side, we show that we can learn provably good values of subsidy in repeated games coming from the same domain. This data-driven subsidy design approach avoids solving computationally hard problems for unseen games by learning over polynomially many games. We also show that optimal subsidy can be learned with no-regret given an online sequence of games, under mild assumptions on the cost matrix. Our study focuses on two distinct games: a Bayesian extension of the well-studied fair cost-sharing game, and a component maintenance game with engineering applications.

Link: https://arxiv.org/abs/2409.03129