“No-Regret Learning Dynamics for Extensive-Form Correlated Equilibrium”

The best way to wrap your head around our work is to consider one of the many problems in which a central planner would like to steer users towards an outcome with high social welfare by guiding the consumption of shared, finite resources. This is common in the gig economy we all live in today: think for example of ride-sharing or food delivery platforms, where you have drivers that provide services to customers in different locations, and the whole market is centralized through a single app that every agent connects to.

In order to steer the system towards high social welfare, one might be tempted to impose a pre-optimized behavior to all the agents in the system. This is harder than it seems, as it requires not only to optimize behavior, but also to come up with legal enforcement paths should the agents decide to deviate from the required behavior.

A softer type of coordination is through a centralized planner that accounts for, and allows, the fact that agents can deviate from the recommended behavior if they want to. This soundly removes the need for enforcement: each agent now is free to selfishly optimize their own behavior. As part of the incentives dynamics, should an agent deviate at any point, the mediator typically reserves the right to stop issuing recommendations to that specific agent. Despite its non-enforcing nature, this type of coordination is already enough to steer the system to social welfare that would be largely impossible through other game-theoretic concepts that do not allow for a central planner, including the Nash equilibrium, which is the most famous equilibrium in game theory.

The study of the incentives in this system, and constructing an incentive-compatible system where we can prove that no agent is better off ever deviating from the recommendations is challenging. It goes back to the seminal work on correlated equilibrium by Robert Aumann in 1974, who was later awarded a Nobel prize for his work on game-theoretic cooperation.

Since then, much effort has been devoted to scaling up the computation of correlated equilibria. When designing algorithms that compute correlated equilibria, it is desirable to optimize the behavior and incentives of each agent independently from the other agents. This decentralization is fundamental to scale well as the number of agents grows. Second, it’s desirable to preserve agents’ privacy during the learning process: agents should not need to report precise utilities back to the central planner. A landmark result by Hart and Mas-Colell in 2000 showed that correlated equilibria can be found in such a decentralized way by letting all agents behave independently, according to simple learning dynamics. Their result was recently awarded the 2020 SIGecom test of time award.

Unfortunately, the work by Hart & Mas-Colell only applies to one-shot interactions, in which each agent is supposed to only interact once with the system, and agents are not allowed to adjust their behavior based on their observations because they are assumed to act simultaneously, as if it was a rock-paper-scissor kind of interaction. This renders the breakthrough by Hart&Mas-Colell hard to apply when it comes to many real-world interactions. Extending their decentralized approach to the general case where agents act more than once and can adjust their behavior to their observations has been an open question since then.

In our paper, we finally close this open problem. We hope it will ease future applications of game-theoretic concepts that guarantee high social welfare, such as correlated equilibrium, in commercial applications.

Media Contact:
Andrea Celli
email: andrea.celli@polimi.it

About the Authors:

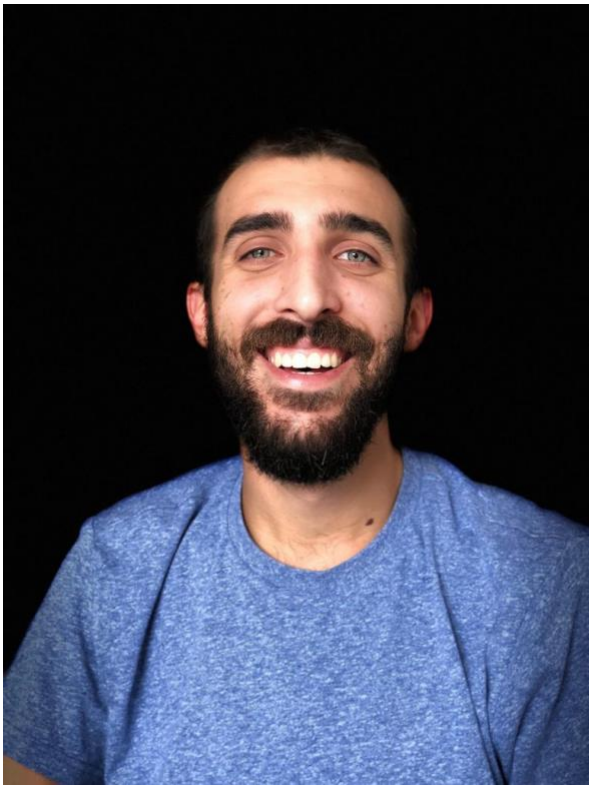

Andrea Celli (presenter)
Title: Ph.D.
Institutions: Politecnico di Milano (now at Facebook Core Data Science)
Bio: Andrea Celli is a postdoctoral researcher at Facebook Core Data Science. Prior to that, he was a postdoctoral researcher at Politecnico di Milano. He has a Ph.D. in Computer Science from the same institution, where he was advised by Prof. Nicola Gatti. He was a visiting researcher at Carnegie Mellon University, supervised by Prof. Tuomas Sandholm. His primary research interests focus on equilibrium computation problems and, more broadly, problems at the intersection of Computer Science and Economics.

Alberto Marchesi
Titles: Ph.D.
Institutions: Politecnico di Milano
Bio: Alberto Marchesi is a postdoc research assistant at the Department of Electronics, Information and Bioengineering of Politecnico di Milano. He received his Ph.D. degree in Computer Science from Politecnico di Milano, where he was advised by Prof. Nicola Gatti. His research focuses on AI and Machine Learning, with a particular emphasis on how to combine AI techniques with economic paradigms, in order to build artificial systems capable of strategic decision making. He is also interested in algorithm design, computational complexity, and optimization.

Gabriele Farina
Titles: Ph.D. Student
Institutions: Carnegie Mellon University
Bio: Gabriele Farina is a fifth-year Ph.D. student in the Computer Science Department at Carnegie Mellon University, where he is advised by Tuomas Sandholm. His primary research interests focus on optimization methods for sequential decision making and convex-concave saddle point problems, with applications to equilibrium finding in games. He is a Facebook Fellow in the area of economics and computer science.

Nicola Gatti
Titles: Associate Professor
Institutions: Politecnico di Milano
Bio: Nicola Gatti is an Associate Professor of Computer Engineering. His main research interests are the design and analysis of optimization and machine learning algorithms for strategic problems such as game theoretic situations, social choice problems, and bandits. On these topics, he published over 130 papers. He was awarded best young Italian researcher in AI in 2011. He has been Associate Editor of Journal of Artificial Intelligence Research since 2015. Currently, he is a member of the managing board of the Italian Laboratory on AI and Intelligent Systems, of the Italian Association for AI, and of the International Foundation for Autonomous Agents and Multiagent Systems. He is also co-director of the Italian Observatory in AI and PI of 6 research projects funded by private or public institutions.