

**Ph.D. in Information Technology
Thesis Defense**

**April 22nd, 2024
at 16:30**

Room BIO1 – building 21

Federico CACCIAMANI – XXXVI Cycle

Learning Optimal Equilibria and Mechanisms Under Information Asymmetry

Supervisor: Prof. Nicola Gatti

Abstract:

Multi-agent environments represent a ubiquitous case of application for game-theoretic techniques. The default game model involves the presence of a large number of agents with asymmetric information on the environment that sequentially interact between them. Algorithmic Game Theory starts from particular notions of equilibria, i.e., sets of strategies for the players such that no one has incentive to deviate from her strategy, and studies the development of algorithms for computing or approximating those. The milestones achieved by researchers in the field over the last decades made it clear how, in order to obtain successful deployments of such game-theoretic techniques in complex real-world settings, it is of uttermost importance to leverage the formulation of learning algorithms for finding approximately optimal solutions of the games. Most of the research has been focused on the design of learning algorithms for simple scenarios, e.g., two-players games, but algorithms for more general cases are still far from reaching adequate performances. The goal of this thesis is to advance the research in this sense. In particular, we investigate different multi-agent scenarios, which we differentiate based on the role of the players holding information on the environment, focusing on the definition of suitable learning algorithms for finding optimal players' strategies. In the first part of the thesis, we study cases in which the agents holding information are active, i.e., they can leverage their information to take informed actions in the game. In this context, we tackle two distinct cases: team games, which model cases in which two teams of agents compete one against the other, as well as the broader class of general-sum games, in which we do not make any particular assumption on the players. For team games, we define a simple transformation that uses a correlation protocol based on public information for obtaining a compact formulation of the teams' strategy sets. The transformation yields an equivalent two-players zero-sum game, which can be naturally used to obtain the first no-regret learning-based algorithm for computing equilibria in team games. Then, inspired by previous literature, we lay the ground for the formulation in the context of team games of popular techniques that proved crucial for achieving strong performances in two-players games,

i.e., stochastic regret minimization and subgame solving. For general-sum games, instead, we observe that the mainstream approach that is being used, which consists in the use of decentralized and coupled learning dynamics for approximating different types of correlated equilibria, suffers from major drawbacks as it does not offer any guarantee on the type of equilibrium reached. To mitigate this issue, we take the perspective of a mediator issuing action recommendations to the players and design a centralized learning dynamic that guarantees convergence to the set of optimal correlated equilibria in sequential games. The second part of the thesis is devoted to the study of cases in which the agents holding information are passive, i.e., they cannot directly take actions in the game, but can only report their information (possibly untruthfully) in order to influence the behavior of another uninformed party. This setting corresponds to the case of information acquisition, in which we take the perspective of the uninformed agent (which we call the principal) that is interested in gathering information from the agents, incentivizing their behavior by means of mechanisms composed by an action policy and/or payment functions. In this context, we separately study the cases in which the principal's mechanisms are composed exclusively by action policies and payment schemes and, for both cases, we provide algorithms for learning optimal mechanisms via interactions with the agents.

Martino BERNASCONI DE LUCA – XXXVI Cycle

Online Learning with Uncertain Constraints: Cumulative and Replenishable Violations

Supervisor: Prof. Francesco Trovò

Abstract:

Learning how to act optimally in complex environments constitutes a crucial milestone in the field of artificial intelligence. In recent years, considerable attention has been directed towards investigating online learning in microeconomic settings, such as Bayesian Persuasion, Mechanism Design, and many other economic scenarios. This line of research is primarily motivated by the fact that it allows overcoming the often unrealistic assumptions of perfect knowledge regarding the exact configuration of the environment. However, most learning paradigms tend to overlook the presence of constraints that are ubiquitous in many real-world applications. These constraints are usually uncertain, and one cannot simply restrict the decision space a priori but needs to adapt online to the observed constraint violations. In this thesis, we take a step towards examining the impact of uncertain constraints on the learning process. Depending on the application, constraints may encapsulate various real-world requirements. For instance, in autonomous driving, safety prerequisites are typically demanded and can be defined by a set of constraints. Similarly, in online advertising, constraints may be imposed on the total budget to spend.

Nonetheless, these constraints manifest very distinct characteristics. In safety constraints, the objective is to avoid violating the constraints entirely, while in budget constraints, some overspending in certain phases and underspending in others are permissible, as long as the overall budget constraint is satisfied. The former are referred to as per-round constraints, while the latter are termed long-term constraints. Therefore, when considering these two complementary settings, we necessitate two different measures for quantifying the effectiveness of the designed algorithms in satisfying the constraints. Namely, this thesis investigates two types of constraint violations, namely cumulative violations and replenishable violations, which are appropriate for per-round and long-term constraints, respectively. In the former, we solely consider the constraint violations incurred by the algorithm, whereas in the latter, we permit the violations to be offset by rounds in which the constraints are satisfied. In this thesis, we analyze the diverse effects of constraints on learning agents in various microeconomic settings. Particularly, in the presence of per-round constraints, we explore various sequential decision models and demonstrate that achieving zero constraint violations is impossible. However, we show that attaining asymptotically small cumulative violations is possible. Conversely, in the presence of long-term constraints, in settings resembling budget constraints, we can establish results that would otherwise be unattainable if no violations were allowed.

PhD Committee

Prof. Nicolò Cesa-Bianchi, Politecnico di Milano

Prof.ssa Tatjana Chavdarova, University of California

Prof. Gabriele Farina, MIT