

Online learning in multi-agent systems : Regret, equilibrium, and the road ahead

Panayotis Mertikopoulos

Université Grenoble Alpes, France
panayotis.mertikopoulos@imag.fr

Tutoriel de l'axe DMEI (Décision, Modélisation, Evaluation, Incertitude) et de l'action transverse DAAO (Données, Apprentissage Automatique et Optimisation) du GDR RO.

Abstract From auctions to traffic planning and the training of adversarial neural nets, the goal of multi-agent learning is always the same : to make more informed decisions that lead to better rewards over time. But who benefits in the end ? Each agent unilaterally ? All agents concurrently ? Both / None of the above ? This tutorial is intended to explore these questions by providing an introduction to online optimization and learning from a multi-agent, game-theoretic viewpoint. We will cover the basics of learning in finite games (focusing on the multiplicative/exponential weights algorithm and its properties), and the mainstay algorithms for online optimization in continuous games (online gradient descent, follow the regularizer leader, etc.). We will also discuss the impact of the information available to the players, as well as a range of concrete applications to machine learning and operations research.