Essays in learning, optimization and game theory

DOI:
10.26481/dis.20210222bd

Summary and General Discussion

This thesis is divided into three Chapters. Chapter 2 deals with learning in games. Chapter 3 is devoted to optimization. Finally Chapter 4 focuses on game theory.

In Chapter 2 we examine the long-term behavior of regret-minimizing agents in time-varying games with continuous action spaces. In its most basic form, (external) regret minimization guarantees that an agent's cumulative payoff is no worse in the long run than that of the agent's best fixed action in hindsight. Going beyond this worst-case guarantee, we consider a dynamic regret variant that compares the agent's accrued rewards to those of *any* sequence of play. By properly adapting a restart procedure pioneered by Besbes et al. [7], we show that players are able to avoid dynamic regret against any test sequence whose total variation grows sublinearly with the horizon of play. In particular, specializing to a wide-class of no-regret strategies based on mirror descent, we derive explicit rates of dynamic regret minimization, both in expectation and in high probability. We then leverage these results to show that players are able to stay close to Nash equilibrium in time-varying monotone games – and even converge to equilibrium if the sequence of stage games admits a limit.

While the information structure is relevant in some contexts, as the reviewers of the article mentioned, this structure is not well-adapted when the information has a different form than the gradient of a function. More precisely, in many applications, agents take decisions based on information (amount of money, quantity of products, ...). Further work has to be done in order to adapt this kind of information structure, and study properties of derivative-free algorithms. Another drawback is that all the agents have to follow the same algorithm. In real life, such an algorithm would probably be a mediator. Thus, it would be more natural to target a correlated equilibrium more than a Nash equilibrium. Such a study would be of great interest.

In Chapter 3 we use a class of strongly convergent primal-dual schemes for solving variational inequalities defined by a Lipschitz continuous and pseudo-monotone map in infinite-dimensional Hilbert spaces, which have been studied by Dennis Meier in [22]. This novel numerical scheme is based on Tseng's forward-backward-forward scheme, which is known to display weak convergence, unless very strong global monotonicity assumptions are made on the involved operators. We test the performance of the algorithm in the computationally challenging task to find dynamic user equilibria in traffic networks and verify that our scheme is at least competitive to state-of-the-art solvers, and in some cases even improves upon them.

In general, the commuting operator of a network is not monotone, but the theoretical convergence of the algorithm relies on this assumption. Further work should done, both in order to relax the strong assumption of monotonicity of the operator, and on finding classes of networks that have a monotone operator.

In Chapter 4 we introduce a discrete-time search game, in which two players compete to find an object first. The object moves according to a time-varying Markov chain on finitely many states. The players know the Markov chain and the initial probability distribution of the object, but do not observe the current state of the object. The players are active in turns. The active player chooses a state, and this choice is observed by the other player. If the object is in the chosen state, this player wins and the game ends. Otherwise, the object moves according to the Markov chain and the game continues at the next period.

We show that this game admits a value, and for any error-term $\varepsilon > 0$, each player has a pure (subgame-perfect) ε-optimal strategy. Interestingly, a 0-optimal strategy does not always exist. The ε-optimal strategies are robust in the sense that they are 2ε-optimal on all finite but sufficiently long horizons, and also 2ε-optimal in the discounted version of the game provided that the discount factor is close to 1. We derive results on the analytic and structural properties of the value and the ε-optimal strategies. Moreover, we examine the performance of the finite truncation strategies, which are easy to calculate and to implement. We devote special attention to the important time-homogeneous case, where additional results hold.

Although many variations of this game are interesting (finitely many players, unknown probability distributions, players move over a graph, the object does not want to be found), two very relevant computational questions arise : how can we compute efficiently optimal strategies for the two players? Is it possible to achieve linear convergence rate of optimal strategies in the finite horizon game to the optimal ones in the infinite horizon game for general sequences of permutation matrices?