SHORT TERM SCIENTIFIC MISSION (STSM) – SCIENTIFIC REPORT

The STSM applicant submits this report for approval to the STSM coordinator

Action number: CA16228
STSM title: Time-Varying games
STSM start and end date: 14/01/2018 to 21/01/2018
Grantee name: Mathias Staudigl

PURPOSE OF THE STSM/

The STSM initiated a new research line on Time-varying games. In time-varying optimization problems, it is common to use diagonal minimization procedures where the agent takes a step along a standard minimization method (e.g. projected (sub)gradient, mirror descent, forward-backward splitting, etc.) and confronts the next function with the same procedure. These methods are called diagonal/staircase methods and, in contrast to standard online learning procedures, they are not designed to minimize a mean-like measure of the sequence of objective functions (such as the regret); instead, their target is to track down a temporal solution. For such an approach to be successful, a fundamental assumption is that the functions evolve "sufficiently slowly": this enables the implicit use of past information to control the gap between the optimized payoff and the actually obtained payoff over time. This problem is relatively well-explored when the payoff functions converge to a well-defined limit, but very little is known when the functions evolve slowly without converging. To that end, we intend to exploit and extend the recent results of who quantify the notion of ``slow evolution'' via the notion of a ``variation budget''. Using this notion, we will derive a learning procedure that is capable of tracking the best dynamic policy in hindsight within a vanishing fraction of this variation budget. The Aim of this STSM is to initiate the important research on Time-Varying Games. In particular, we focus on the analysis of low regret algorithms in stochastic Nash games and more general time-varying games, not requiring any assumptions on the data generating process. The techniques used in this STSM are Stochastic Approximation, Mirror-Descent Algorithms, Probability Theory and Online Optimization. The applicant for the STSM will present the first results on Time-varying games to the groups at the Host institution. This will lead to the writing of a working paper which in turn will be submitted in 2018 for publication. Based on these discussions, we will discuss how to enrich the model, but also focus on other models of learning in games, based on replicator dynamics.

DESCRIPTION OF WORK CARRIED OUT DURING THE STSMS

Day 1: Presentation of the preliminary results of the applicant; Setting up the general model framework; Cross-checking with the existing literature.

Day 2: Talk at the Game Theory Seminar of the Applicant at Tel-Aviv University. Meeting with experts at Tel-aviv university to discuss the idea to get feedback.

Day 3: Writing up of the model framework and prove of the low static regret result.

Day 4: Talk of the applicant at Bar-Ilan University. Discussion with local experts; Presentation of results.

Day 5: Working out a work plan to continue the research. Setting up further projects on hybrid replicator dynamics

Day 6 and 7: Work on generalizations of the model, and get insights on the scaling of the dynamic regret, building on the results of Day 3, where the bounds on the dynamic regret have been proven.

DESCRIPTION OF THE MAIN RESULTS OBTAINED

We have introduced the model of time-varying games. The main innovation of this model is to bound the dynamic regret for general convex games. We have established the important theorem that whenever players agree to play a strategy guaranteeing low static regret, it automatically generates a low dynamic regret when the algorithm is enriched with a suitable restart procedure. Remarkably, this result does not depend on the specific strategies the players are choosing and also not on the specific utility function players are facing. Nature can be as adversarial, provided it chooses the sequence of utility functions from a set of convex smooth games with bounded variation. It suffices that players agree to use a low regret algorithm at each stage. Explicit bounds in terms of the rate of convergence are obtainable when one specific algorithm, such as the Mirror-prox is assumed. This result generalizes recent studies on the scaling of non-stationary stochastic optimization, as introduced in recent work by Besbes et. al. (2015) to problems written in terms of non-stationary variational inequalities. It remains to investigate concrete implementations to stochastic Nash games, like Cournot-Oligopoly, or routing games.

FUTURE COLLABORATIONS (if applicable)

(max.500 words)

Potential future collaboration on Replicator dynamics with type-structured populations is envisioned.