Research Project Proposal: Abstractions in Extensive-Form Games

Jacopo P. Gargano, jacopopio.gargano@mail.polimi.it

1. Introduction to the problem

The area of focus of our research is Algorithmic Game Theory, a field of study that aims to analyze strategic conditions and to design algorithms able to find strategies for the involved agents allowing them to reach an *equilibrium*[1]. A strategic environment is mathematically modeled through formal representations so as to describe its problems and solutions, if any. Solving a strategic problem requires the use of the theory of computation and algorithm design: the former to analyze the problem complexity and evaluate its *difficulty*[2], the latter to solve the problem - usually corresponding to finding *equilibria*[1]. Therefore, Algorithmic Game Theory is a combination of Mathematics, specifically Game Theory, and Computer Science.

1.1. Research topic and its importance

The main problem faced by Game Theory is that of game representation and resolution. Solving a game typically means finding its *Nash equilibria*[3]. Finding exact Nash equilibria is not always feasible. This is why approximated solutions of the game, corresponding to quasi-optimal strategies, are considered good. These go under the name of ϵ-Nash equilibria.

Algorithmic Game Theory is focused on designing algorithms to solve games, comprising finding Nash equilibria. In the early days of Algorithmic Game Theory, relatively small games were analyzed and their reduced size allowed them to be solved through the use of linear programming [1]. Recently however, with the introduction of the concept of regret minimization in imperfect information games [12], the most used ways to solve games for Nash equilibria are based on Counterfactual Regret Minimization (*CFR*) [12], which can solve larger games. Not surprisingly there exist many variations of it with improved performances [2, 6, 8, 9]. In general, there are several equilibrium-finding algorithms to solve a game, however a central challenge in solving games is that the game might be too *large*[4]. For instance, two-player no-limit Texas hold'em poker has more than 10^{165} nodes [7].

In order to solve the issue of complexity in decision making with large games, the concept of *abstracting* games was developed [1]. Abstractions are a method to lower game complexity while retaining all relevant information. Abstracting generally consists in building a smaller version of the game tree with a reduced number of states and actions.

In practice, sequential games are very large and their complexity prevents them to be fully represented, explored and analyzed to find equilibria. Being able to study large and infinite games through abstractions is crucial to extend the applicability of game theoretical principles to real-world problems. These include every possible strategic situation that is representable through a sequential game, including but not limited to recreational games, sports, governance and conflicts. This is why our research topic is of great significance.

[1] See Section 1.1.

[2] In terms of computational complexity (e.g. NP-hardness).

[3] *Nash equilibrium*: a strategy profile such that each player does not benefit from deviating from their strategy, keeping the strategies of all the other players fixed. According to Nash's Existence Theorem, every game with a finite number of players in which each player can choose from finitely many pure strategies has at least one Nash equilibrium.

[4] *Large game*: a game whose representation through a tree is infeasible.

1.2. The problem and its importance

Despite the remarkable contributions[5] in the field of abstractions, especially those by Brown and Sandholm [3, 4], there are still several open problems to tackle. Most real-world strategic games are too large given that the available actions belong to a continuous space. These kinds of games are known as infinite games and there is no explicit way to fully represent them. The only available option to obtain exact information about large or infinite games is to collect game samples in the form of *traces*[6] and corresponding payoffs for the players according to their preference.

Moreover, information about the game in the form of samples, obtained via actual play or simulations, is nowadays largely and easily accessible. However, little research has been carried out on *simulation-based games*, where a complete and accurate description of the game is not available, but game samples and corresponding noisy payoffs are. Specifically, to the best of our knowledge, with respect to our research topic, the only significant results were achieved by Viqueira et al. in [11]. They present a method able to learn all approximated equilibria of a simulation-based game. However, according to the authors, their algorithm can only find *pure strategy*[7] ϵ-Nash equilibria. Considering that all games admit *mixed strategy* Nash equilibria and few also admit pure ones, this open problem is of great importance in the field.

Furthermore, previous works have mainly contributed with domain-specific implementations of the presented abstraction methods. Most of them were focused on heads-up no-limit Texas hold'em Poker and simpler variants of it. A general model-free[8] approach has not been presented yet. This issue is at least as meaningful as the aforementioned more theoretical problem, if not more.

Therefore, with our research we aim to solve the following problems: real-world games are too large to be represented, no clear domain-independent abstraction approach was presented to solve these games, poker is the only main reference application of these techniques.

After designing our solution to solving large or infinite sequential games, we will experiment it by finding optimal strategies in cybersecurity scenarios. We believe that our research will have a great impact on Security, both physical and cyber. Being able to solve large games would allow the application of game-theoretic principles, that is finding optimal strategies, to any real-world meaningful strategic situation. Critical infrastructures can be protected, illegal drug, money and weapons trafficking could be drastically limited, and urban crime could be suppressed. Cybersecurity problems are more complex than physical ones, as the space of actions can be much larger, leading to infinite games, and this is why they are a perfect fit to our research.

Our research precisely aims to find a way to abstract large games and to solve them for optimal strategies. Our contributions to society would be useful and notable, resulting in a great impact.

2. Main related works

The main works related to our research topic are the ones regarding abstraction approaches in extensive-form games. These can be distinguished between those in which games are explicitly representable and are abstracted through *information* or *action* abstractions and those studying implicit games in a simulation-based fashion.

The most advanced techniques comprising abstractions were developed by Brown and Sandholm first in [3] and then in [4], contributing respectively with *Libratus* and *Pluribus*. *Libratus* features three main modules. The first computes an abstraction of the game and solves it through self-play via an improved version of Monte Carlo Counterfactual Regret Minimization *MCCFR* [6, 8], obtaining a *blueprint strategy*. The second module comes into play later in the game as a *refinement* by constructing a finer-grained abstraction for a particular part of the game that is reached during play and solves it in real time. They exploit the *nested subgame solving* technique, which consists in solving a subgame with the actions outside the precomputed abstraction included. This technique

[5] See Section 2

[6] *Game trace*: an ordered list of traversed states and undertaken actions.

[7] Strategies can be *pure* or *mixed*. Actions of a mixed strategy are taken according to a probability distribution; in a pure strategy only one action is taken and all others never are.

[8] *Model-free*: no information on the game is available besides game samples and corresponding approximated payoffs.

comes with a provable safety guarantee [5]. Finally, *Libratus* that is able to self-improve by enhancing the *blueprint strategy* filling in missing branches in the blueprint abstraction and solving those for a strategy. *Pluribus* is an enhanced version of its predecessor *Libratus* able to play six-player heads-up no-limit poker. Despite proving itself to be an undisputed winner against top players, results are not solidly supported by theory.

A theoretical contribution to simulation-based games and relative abstractions was given by Tuyls et al. with [10]. This year Viqueira et al. presented [11]. In this work the authors study simulation-based games, finding all pure strategy ϵ-Nash equilibria.

3. Research plan

3.1. The goal

The goal of our research is to develop a *bottom-up model-free abstraction approach*, supported by *theoretical guarantees*, able to find *mixed strategy Nash equilibria* in *any* extensive-form game in a *simulation-based fashion*, that is, only through game samples – in the form of traces and corresponding approximated payoffs – as the only inputs.

Therefore, our research is a blend of algorithm design and theoretical analysis, with the ultimate aim of providing a general method for any simulation-based game. As a consequence, the nature of our research is mainly theoretical and will be supported by various implementations to show that our approach is domain-independent.

3.2. The process

The following is a GANTT chart representing the main activities of our research and their estimated duration.

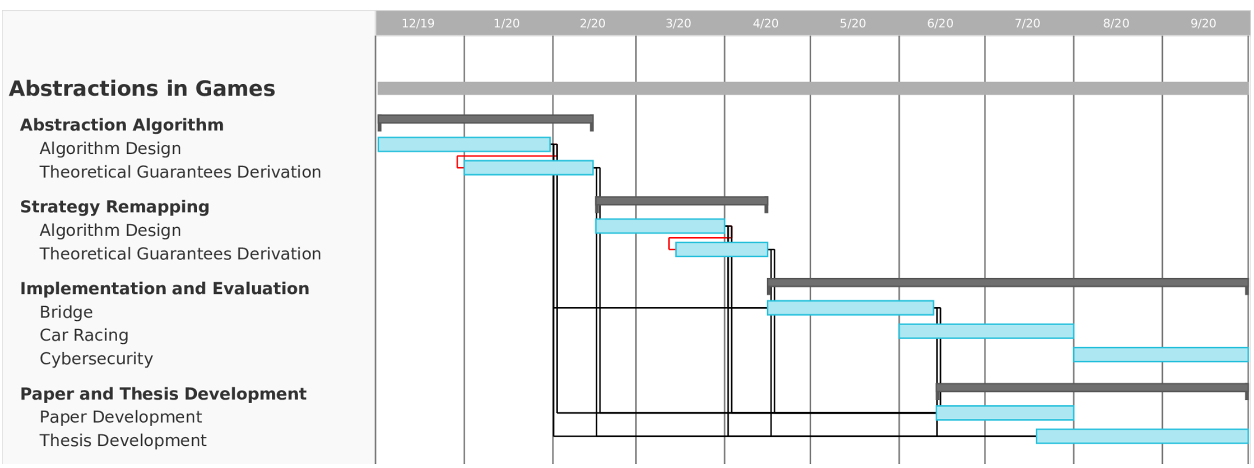

Figure 1: GANTT Chart of Research Activities

3.3. Research evaluation

The metrics that will be used to evaluate the output of our research will be based on the following:

- Convergence rate of algorithms
- Quality of the approximations
- Performance of experimental implementations, calculated through application specific results (e.g. percentage of wins when playing bridge)

References

[1] Billings, D., Burch, N., Davidson, A., Holte, R., Schaeffer, J., Schauenberg, T., and Szafron, D. Approximating game-theoretic optimal strategies for full-scale poker. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence* (San Francisco, CA, USA, 2003), IJCAI'03, Morgan Kaufmann Publishers Inc., pp. 661–668.

[2] Brown, N., Lerer, A., Gross, S., and Sandholm, T. Deep counterfactual regret minimization. *arXiv preprint arXiv:1811.00164* (2018).

[3] Brown, N., and Sandholm, T. Superhuman ai for heads-up no-limit poker: Libratus beats top professionals. *Science 359*, 6374 (2018), 418–424.

[4] Brown, N., and Sandholm, T. Superhuman ai for multiplayer poker. *Science 365*, 6456 (2019), 885–890.

[5] Burch, N., Johanson, M., and Bowling, M. Solving imperfect information games using decomposition. In *Twenty-Eighth AAAI Conference on Artificial Intelligence* (2014).

[6] Gibson, R., Lanctot, M., Burch, N., Szafron, D., and Bowling, M. Generalized sampling and variance in counterfactual regret minimization. In *Twenty-Sixth AAAI Conference on Artificial Intelligence* (2012).

[7] Johanson, M. Measuring the size of large no-limit poker games. *arXiv preprint arXiv:1302.7008* (2013).

[8] Lanctot, M., Waugh, K., Zinkevich, M., and Bowling, M. Monte carlo sampling for regret minimization in extensive games. In *Advances in neural information processing systems* (2009), pp. 1078–1086.

[9] Tammelin, O. Solving large imperfect information games using cfr+. *arXiv preprint arXiv:1407.5042* (2014).

[10] Tuyls, K., Perolat, J., Lanctot, M., Leibo, J. Z., and Graepel, T. A generalised method for empirical game theoretic analysis, 2018.

[11] Viqueira, E. A., Cousins, C., Upfal, E., and Greenwald, A. Learning equilibria of simulation-based games, 2019.

[12] Zinkevich, M., Johanson, M., Bowling, M., and Piccione, C. Regret minimization in games with incomplete information. In *Advances in neural information processing systems* (2008), pp. 1729–1736.