Personality-Based Representations of Imperfect-Recall Games

Extended Abstract

Andrea Celli
Politecnico di Milano
Milan, Italy
andrea.celli@polimi.it

Giulia Romano
Politecnico di Milano
Milan, Italy
giulia.romano@polimi.it

Nicola Gatti
Politecnico di Milano
Milan, Italy
nicola.gatti@polimi.it

ABSTRACT

Games with imperfect recall are a powerful model of strategic interactions that allows agents to forget less important details of the past. Nevertheless, the computational treatment of imperfect-recall games is largely unexplored so far, and no efficient strategy representation for this setting is known. In this paper, we focus on general imperfect-recall games without *absentmindedness*, and we study how to produce a perfect-recall representation of these games using *personalities*. In particular, a valid personality assignment is a decomposition of an imperfect-recall player such that she does not exhibit memory losses within the same personality. Given a valid personality assignment, we can build an *auxiliary team game* where a team of perfect-recall players—sharing the same objectives—replaces a player with imperfect recall. Our primary goal is the construction of a compact representation in terms of number of personalities. We study the (iterated) *inflation* operation as a way to simplify the information structure of a game with imperfect recall. We show that the *complete* (i.e., maximal) *inflation* of a game can be found in polynomial time. We also show that finding the valid personality assignment minimizing the number of personalities is NP-hard, and also hard to approximate, unless P = NP, even in a completely inflated game.

KEYWORDS

Equilibrium computation; imperfect-recall games

ACM Reference Format:
Andrea Celli, Giulia Romano, and Nicola Gatti. 2019. Personality-Based Representations of Imperfect-Recall Games. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION

The treatment of imperfect-recall games is much more involved than that of games with perfect recall. First, the equivalence between *behavioral* and *mixed* strategies no longer holds [16, 18]. Moreover, even solving two-player zero-sum games is exponential in the worst case, unless P = NP [12]. In this setting, the sequence form is not well defined [19] and other approaches for computing equilibria in the perfect-recall setting, e.g. regret-based algorithms [9, 22], lose their theoretical guarantees.

Despite these difficulties, imperfect-recall games recently became widely adopted in solving large imperfect-information games with imperfect-recall abstractions of the original game (see, e.g., [10, 21]). However, most of the algorithmic techniques known so far for imperfect-recall games require strong assumptions on the structure of the game [4, 13–15].

One of the crucial issues related to games with imperfect recall is the choice of the most appropriate strategy representation. Mixed strategies are more expressive than behavioral strategies [11], which may lead to arbitrarily inefficient (in the number of terminal nodes) solutions [5]. [1] Hence, such inefficiencies pushes for the adoption of mixed strategies. On the other hand, mixed strategies are often impractical since their size is exponential, while behavioral strategies represent a compact (linear) alternative. In practice, the empirical inefficiency of behavioral strategies in team games has been shown to be negligible in many settings [1–3, 5].

We start by providing the first polynomial-time algorithm to compute the *complete inflation* of an imperfect-recall game in an iterative fashion. Next, we introduce the notion of *personality* and define a *valid personality assignment.* A personality of an imperfect-recall player is a subset of her information sets where she does not exhibit memory losses. An imperfect-recall player can be seen as a *team*—a set of perfect-recall players with the same objectives [20]—where each personality corresponds to a team member. Given an imperfect-recall game, we build an equivalent *auxiliary team game* with perfect recall with the following features: (i) mixed strategies have the same expressive power in the former and in the latter, (ii) it reduces the inefficiencies of employing behavioral strategies. The inefficiency of behavioral strategies increases exponentially in the number of team members [5]. Equivalently, in imperfect-recall games, the inefficiency increases in the number of personalities used to represent an imperfect-recall player. Then, in the perspective of computing an approximate solution (i.e., an approximate Nash equilibrium) through behavioral strategies, it is crucial to minimize the number of personalities required to build the auxiliary game. We show that the problem of computing a valid personality assignment with the smallest number of personalities is APX-hard even on completely inflated trees.

2 PRELIMINARIES

An extensive-form game Γ with imperfect information has a finite set of players $\mathcal{P}$, a finite set of actions A, and a finite set of histories H. Each history identifies a node of the game tree. History h is a prefix of h' ($h \sqsubseteq h'$) if h' begins with h. Each player $i \in \mathcal{P}$ has a set $\mathcal{I}_i$ of information sets, which form a partition over her decision nodes. For any $I \in \mathcal{I}_i$, all $h, h' \in I$ are indistinguishable to player i. The set of actions available at I is denoted by $A(I)$. Given a history h, define $X_i(h)$ to be the sequence of pairs composed of information

[1] In particular, [5] show that, in team games, the inefficiency may increase exponentially in the number of team members.

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada. © 2019 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

set and action s.t. $(I, a) \in X_i(h)$ if $I \in \mathcal{I}_i$ and there exists $h' \sqsubseteq h$ such that $h' \in I$ and $h'a \sqsubseteq h$. The order of the pairs in $X_i(h)$ is the same in which they occur in h. Define $X(h)$ to be the sequence pairs belonging to all players, and $X_{-i}(h)$ similarly, by removing player i's information set, action pairs from $X(h)$.

We say that a player i has perfect recall if, $\forall I \in \mathcal{I}_i$, and $\forall h, h' \in I$: $X_i(h) = X_i(h')$. Otherwise, the player has imperfect recall. An imperfect-recall game exhibits *absentmindedness* if a player forgets previous moves she made [18]. In the following, we focus on games without absentmindedness.

A *pure normal-form plan* for player i is a tuple $s \in S_i = \times_{I \in \mathcal{I}_i} A(I)$. We denote by $s(I)$ the action selected in s at information set I. A *mixed strategy* σ_i for player i is defined as $\sigma_i : S_i \rightarrow \Delta^{|S_i|}$. A *behavioral strategy* $\pi_i \in \Pi_i$ associates each $I \in \mathcal{I}_i$ with a probability vector over $A(I)$. Two strategies of player i are *realization equivalent* if, given any strategy of the other players, they force the same distribution over the outcomes of the game.

3 INFLATION

The basic idea behind *inflation* [7, 11, 17] is splitting information sets that cause memory losses while guaranteeing that, $\forall i \in \mathcal{P}$, each mixed strategy of the original game has a realization equivalent mixed strategy in the new game, and vice versa.

Definition 3.1 (Immediate Inflation). Let $\mathcal{I}_i$ and $\mathcal{I}_i'$ be two possible information partitions of $i \in \mathcal{P}$. $\mathcal{I}_i'$ is an immediate inflation of $\mathcal{I}_i$ iff there exist $I \in \mathcal{I}_i$ and $I_1, I_2 \in \mathcal{I}_i'$ such that: (i) $I = I_1 \cup I_2$ and $\mathcal{I}_i \setminus \{I\} = \mathcal{I}_i' \setminus \{I_1, I_2\}$, and (ii) $\forall h_1 \in I_1, h_2 \in I_2$, there exists $\bar{I} \in \mathcal{I}_i \cap \mathcal{I}_i'$ such that $(\bar{I}, a) \in X_i(h_1)$, $(\bar{I}, b) \in X_i(h_2)$ and $a \neq b$.

Definition 3.2 (Inflation). Given a player $i \in \mathcal{P}$, an information partition $\mathcal{I}_i'$ is called an inflation of $\mathcal{I}_i$ iff $\mathcal{I}_i'$ is obtained by successive applications of immediate inflations to $\mathcal{I}_i$.

When an inflation of $\mathcal{I}_i$ has no further immediate inflations, it is called *complete inflation* of $\mathcal{I}_i$ and it is denoted by $\check{\mathcal{I}}_i$.

Theorem 3.3. *Given Γ and an imperfect-recall player $i \in \mathcal{P}$, the completely inflated game $\check{\Gamma}_i$ can be computed in polynomial time in the size of the game tree.*

Intuitively, the polynomial-time algorithm has to compare the histories associated with each pair of decision nodes in the same information set. However, the inflation operation does not entirely solve the representation problem, since a completely inflated game may still have imperfect recall.

4 MULTIPLE PERSONALITY APPROACH

In this section, we describe the relationship between games with imperfect recall and team games (see [2, 5, 20]) by introducing the notion of *personality*. A *team* is a set of players that have the same preferences but different information about the state of the game. A player with imperfect recall (not absentminded) can always be seen as a collection of hypothetical team members (personalities). Farina et al. [2018a] show that the problem of computing *ex ante coordinated* strategies for a team of players can be addressed by working on an imperfect-recall *meta-player*. We show that the converse is also true by building, given a valid assignment of personalities, an auxiliary game where the player with imperfect recall is decomposed in a team of players with perfect recall.

Representation. We define a personality to be a subset of information sets of a player such that, for each information set belonging to it, the player does not exhibit memory losses with respect to what he knew at past information sets of the same personality.

Definition 4.1 (Personality). Given $i \in \mathcal{P}$ with information partition $\mathcal{I}_i$, a personality $\tilde{\mathcal{I}}_i^k \subseteq \mathcal{I}_i$ is a subset of information sets such that a hypothetical player j with $\mathcal{I}_j = \tilde{\mathcal{I}}_i^k$ would have perfect recall.

We say that $\tilde{\mathcal{I}}_i = \{\tilde{\mathcal{I}}_i^1, \tilde{\mathcal{I}}_i^2, \ldots\}$ is a valid personality assignment for $i \in \mathcal{P}$ iff $\tilde{\mathcal{I}}_i$ is a partition of $\mathcal{I}_i$ and each element of $\tilde{\mathcal{I}}_i$ is a personality of player i. In games without absentmindedness, it is always possible to find a (trivial) valid personality assignment (i.e., assigning one personality $\forall I \in \mathcal{I}_i$). The notion of valid assignment allows us to define an equivalent *auxiliary game* where a player with imperfect recall is substituted with a team of perfect recall players, sharing the same utility function. Definition 4.1 implies that each new game has perfect recall. For each $i \in \mathcal{P}$, and each $\sigma_i \in \Sigma_i$, there exists a realization equivalent *ex ante coordinated normal-form strategy* for the team[2]. In this setting, just before the beginning of the game, a coordination device randomly draws a joint normal-form plan from a known joint probability distribution, and the team members act as specified in the selected plan. We can conclude that an imperfect recall can be equivalently tackled as a team game with perfect recall.

Minimizing the Number of Personalities. When one looks for an approximate NE in behavioral strategies, the minimization of the number of personalities is crucial to reduce both the inefficiencies and the complexity of the problem. The problem of computing the assignment with the minimum possible cardinality is denoted by MIN-P. It is possible to show that MIN-P is NP-hard for the general class of imperfect-recall games, and that it remains hard on the subclass of completely inflated games. In order to prove these results, it is possible to provide a reduction from 3-SAT. For a given boolean formula in conjunctive normal form, we can map the satisfiability problem to a personality assignment problem by embedding in a suitable tree an OR-gadget for each clause. Finally, it can be shown that MIN-P does not admit a PTAS, unless P = NP. Specifically, MIN-P is APX-hard, and no better approximation than 4/3 is possible in polynomial time, unless P = NP.

5 DISCUSSION

The duality between imperfect-recall games and team games raises a number of interesting questions. First, it would be interesting to evaluate, in a practical scenario, the impact of inflation in terms of *split* information sets. Moreover, our negative results suggest that, when working with mixed strategies, one should look for techniques for equilibrium computation in team games that are robust with respect to the number of team members. It would be interesting to confirm this idea with an experimental evaluation of state-of-the-art techniques from the team game domain. Finally, providing scalable algorithms to compute valid personality assignments—other than the trivial one—is a problem we are planning to address. This could be done, for example, through local search techniques as already done in other areas of the algorithmic game theory field [6].

[2] Some works call these strategies *ex ante correlated strategies* [5]. We adopt the terminology introduced by Farina et al. [8]

REFERENCES

[1] Nicola Basilico, Andrea Celli, Giuseppe De Nittis, and Nicola Gatti. 2017. Computing the team–maxmin equilibrium in single–team single–adversary team games. *Intelligenza Artificiale* 11, 1 (2017), 67–79.

[2] Nicola Basilico, Andrea Celli, Giuseppe De Nittis, and Nicola Gatti. 2017. Team-maxmin equilibrium: efficiency bounds and algorithms. In *AAAI Conference on Artificial Intelligence (AAAI)*. 356–362.

[3] Nicola Basilico, Andrea Celli, Giuseppe De Nittis, and Nicola Gatti. 2017. Coordinating Multiple Defensive Resources in Patrolling Games with Alarm Systems. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 678–686.

[4] Branislav Bosanskỳ, Albert Xin Jiang, Milind Tambe, and Christopher Kiekintveld. 2015. Combining Compact Representation and Incremental Generation in Large Games with Sequential Strategies.. In *AAAI Conference on Artificial Intelligence (AAAI)*. 812–818.

[5] Andrea Celli and Nicola Gatti. 2018. Computational Results for Extensive-Form Adversarial Team Games. In *AAAI Conference on Artificial Intelligence (AAAI)*.

[6] Sofia Ceppi, Nicola Gatti, Giorgio Patrini, and Marco Rocco. 2010. Local search techniques for computing equilibria in two-player general-sum strategic-form games. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 1469–1470.

[7] Norman Dalkey. 1953. Equivalence of Information Patterns and Essentially Determinate Games. *Contributions to the Theory of Games* 2 (1953), 217.

[8] Gabriele Farina, Andrea Celli, Nicola Gatti, and Tuomas Sandholm. 2018. Ex ante coordination and collusion in zero-sum multi-player extensive-form games. In *Advances in Neural Information Processing Systems (NeurIPS)*.

[9] Gabriele Farina, Christian Kroer, and Tuomas Sandholm. 2018. Online convex optimization for sequential decision processes and extensive-form games. *arXiv preprint arXiv:1809.03075* (2018).

[10] Michael Johanson, Neil Burch, Richard Valenzano, and Michael Bowling. 2013. Evaluating state-space abstractions in extensive-form games. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 271–278.

[11] Mamoru Kaneko and J. Jude Kline. 1995. Behavior strategies, mixed strategies and perfect recall. *International Journal of Game Theory* 24, 2 (01 Jun 1995), 127–145.

[12] Daphne Koller, Nimrod Megiddo, and Bernhard Von Stengel. 1994. Fast algorithms for finding randomized strategies in game trees. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*. ACM, 750–759.

[13] Christian Kroer and Tuomas Sandholm. 2016. Imperfect-Recall Abstractions with Bounds in Games. In *Proceedings of the ACM Conference on Economics and Computation (EC)*. 459–476.

[14] Marc Lanctot, Richard Gibson, Neil Burch, Martin Zinkevich, and Michael Bowling. 2012. No-regret learning in extensive-form games with imperfect recall. In *Proceedings of the International Conference on Machine Learning (ICML)*. 1035–1042.

[15] Viliam Lisý, Trevor Davis, and Michael H. Bowling. 2016. Counterfactual Regret Minimization in Sequential Security Games. In *AAAI Conference on Artificial Intelligence (AAAI)*. 544–550.

[16] Michael Maschler, Eilon Solan, and Shmuel Zamir. 2013. *Game Theory*. Cambridge University Press.

[17] Akira Okada. 1987. Complete inflation and perfect recall in extensive games. *International Journal of Game Theory* 16, 2 (1987), 85–91.

[18] Michele Piccione and Ariel Rubinstein. 1997. On the interpretation of decision problems with imperfect recall. *Games and Economic Behavior* 20, 1 (1997), 3–24.

[19] Bernhard Von Stengel. 1996. Efficient computation of behavior strategies. *Games and Economic Behavior* 14, 2 (1996), 220–246.

[20] Bernhard von Stengel and Daphne Koller. 1997. Team-maxmin equilibria. *Games and Economic Behavior* 21, 1-2 (1997), 309–321.

[21] Kevin Waugh, Martin Zinkevich, Michael Johanson, Morgan Kan, David Schnizlein, and Michael H. Bowling. 2009. A Practical Use of Imperfect Recall. In *Eighth Symposium on Abstraction, Reformulation, and Approximation, (SARA)*.

[22] Martin Zinkevich, Michael Johanson, Michael Bowling, and Carmelo Piccione. 2008. Regret minimization in games with incomplete information. In *Advances in Neural Information Processing Systems (NeurIPS)*. 1729–1736.