

Strategic Abilities of Forgetful Agents in Stochastic Environments

Francesco Belardinelli¹, Wojciech Jamroga^{2,3}, Munyque Mittelmann⁴, Aniello Murano⁴

¹Imperial College, London, UK

²Interdisc. Centre on Security, Reliability and Trust, SnT, University of Luxembourg

³Institute of Computer Science, Polish Academy of Sciences, Warsaw, Poland

⁴University of Naples “Federico II”

francesco.belardinelli@imperial.ac.uk, wojciech.jamroga@uni.lu,

{munyque.mittelmann, aniello.murano}@unina.it

Abstract

In this paper, we investigate the probabilistic variants of the strategy logics ATL and ATL* under imperfect information. Specifically, we present novel decidability and complexity results when the model transitions are stochastic and agents play uniform strategies. That is, the semantics of the logics are based on multi-agent, stochastic transition systems with imperfect information, which combine two sources of uncertainty, namely, the partial observability agents have on the environment, and the likelihood of transitions to occur from a system state. Since the model checking problem is undecidable in general in this setting, we restrict our attention to agents with memoryless (positional) strategies. The resulting setting captures the situation in which agents have qualitative uncertainty of the local state and quantitative uncertainty about the occurrence of future events. We illustrate the usefulness of this setting with meaningful examples.

1 Introduction

Complex and interacting Multi-Agent Systems (MAS) often face different kinds of uncertainty. One of the sources of uncertainty is the inability to completely observe the current local situation (e.g., whether there is public transport available to the target destination). On the other hand, the occurrence of many natural events and the future behaviour of other agents, while it cannot be known with certainty, can be measured based on experiments or past observations. For instance, while we cannot know if the bus is going to arrive on time, we may have observed that this happens 0.7% of the time. Clearly, intelligent autonomous agents need to consider both the imperfect information about the local state and the likelihood of stochastic events when making strategic decisions and plans.

To see this, consider, for instance, the problem of online mechanisms, which are preference aggregation games in dynamic environments with multiple agents and private information. Many multi-agent problems are inherently dynamic rather than static. Practical examples include the problem of allocating computational resources (bandwidth, CPU, etc.) to processes arriving over time, selling items to a possibly changing group of buyers with uncertainty about the future supply, and selecting employees from a dynamically changing list of candidates (Nisan et al., 2007).

Probabilistic model-checking is a technique for the formal and automated analysis of probabilistic systems that can be modeled by stochastic state-transition models (Clarke et al., 2018a). Its aim is to establish the correctness of such systems against probabilistic specifications, which may describe, e.g., the probability of an unsafe event to occur, or the ability of a coalition to ensure the completeness of a task.

Logic-based approaches have been widely and successfully applied for probabilistic verification of MAS. For instance, probabilistic model-checking techniques have been used for verification of preference aggregation mechanisms (Mittelmann et al., 2023), negotiation games (Ballarini, Fisher, and Wooldridge, 2009), team formation protocols (Chen et al., 2011), and stochastic behaviors in dispersion games (Hao et al., 2012), to name a few. Gutierrez et al. (2021) investigates the problem of deciding whether the probability of satisfying a given temporal formula in a concurrent stochastic game is 1 or greater than 0. Kwiatkowska et al. (2022) details how verification techniques can be developed and implemented for concurrent stochastic games.

In this paper, we consider logics for reasoning about strategic abilities while taking into account both incomplete information and probabilistic behaviors of the environment and agents. We study the Probabilistic Alternating-time Temporal Logics PATL and PATL* (Chen and Lu, 2007; Hao et al., 2012) under imperfect information (II) for a classic type of agents (Fagin et al., 2004) called imperfect-recall (that is, agents who use memoryless strategies, also called Markovian strategies or policies). Model checking PATL* under II for agents with perfect-recall (who uses memoryful strategies) is known to be undecidable in general even for the fragment with a single-player (Hao et al., 2012). We introduce and motivate the problem of strategic reasoning under combined types of uncertainty and memoryless agents. We then provide results on the model-checking complexity for PATL with memoryless deterministic strategies for the coalition and point directions to challenging open questions.

Related Work. Recently, much work has been done on logics for strategic reasoning in Multi-Agent Systems, starting from the pioneering work on Alternating-time Temporal Logics ATL and ATL* (Alur, Henzinger, and Kupferman, 2002). These logics enable reasoning about the strategic abilities of agents in a cooperative or competitive sys-

tem. ATL has been extended in various directions, considering for instance strategy contexts (Laroussinie and Markey, 2015) or adding imperfect information (Jamroga and Bulling, 2011). Strategy Logic (SL) (Chatterjee, Henzinger, and Piterman, 2010; Mogavero et al., 2014) extends ATL to treat strategies as first-order variables.

Contexts of imperfect information have been extensively considered in the literature on formal verification (see, for instance, (Dima and Tiplea, 2011; Kupferman and Vardi, 2000; Jamroga and Ågotnes, 2007; Reif, 1984; Bulling and Jamroga, 2014; Berthon et al., 2021; Belardinelli et al., 2020; Berwanger and Doyen, 2008)). Generally, imperfect information in MAS entails higher complexity, which may be even undecidable when considered in the context of memoryful strategies (Dima and Tiplea, 2011). In order to retrieve a decidable model-checking problem, it is interesting to study imperfect information MAS with memoryless agents (Cermák et al., 2018).

Several works consider the verification of systems against specifications given in probabilistic logics. In particular, Wan, Bentahar, and Hamza (2013) study the model-checking problem for Probabilistic Epistemic Computational Tree Logic with semantics based on probabilistic interpreted systems. In the context of MAS, (Huang and Luo, 2013) studies an ATL-like logic for stochastic MAS in a setting in which agents play deterministic strategies and have probabilistic knowledge about the system. (Fu et al., 2018) shows model-checking an epistemic logic with temporal operators under strategies that depend only on agents' observation history is undecidable.

Chen and Lu (2007) propose model-checking algorithms for Probabilistic ATL in the perfect information setting. Perfect information was also considered with specification in Probabilistic Alternating-Time μ -Calculus (Song et al., 2019) and Probabilistic Strategy Logic Aminof et al. (2019). ATL-based probabilistic logics were also considered for the verification of unbounded parameterized MAS (Lomuscio and Pirovano, 2020), for resource-bounded MAS (Nguyen and Rakib, 2019), and under assumptions over opponents' strategies (Bulling and Jamroga, 2009).

The closest related work is (Huang, Su, and Zhang, 2012), which considers the logic PATL* under incomplete information and synchronous perfect recall. The complexity results show that the model-checking problem is in general undecidable even for the single-agent fragment of the logic.

Also related are the works in (Gripon and Serre, 2009; Doyen and Raskin, 2011; Carayol, Löding, and Serre, 2018; Doyen, 2022), which consider algorithmic solutions for computing the existence of winning strategies and winning distributions for two-player stochastic games with imperfect information. Finally, Gurov, Goranko, and Lundberg (2022) investigate the problem of strategy synthesis for knowledge-based strategies against a non-deterministic environment.

2 Preliminaries

In this paper, we fix finite non-empty sets of agents Ag , actions Ac , atomic propositions AP . We write \mathbf{o} for a tuple of objects $(o_a)_{a \in \text{Ag}}$, one for each agent, and such tuples are called *profiles*. A *joint action* or *move* \mathbf{c} is an element of

Ac^{Ag} . Given a profile \mathbf{o} and $C \subseteq \text{Ag}$, we let \mathbf{o}_C be the components of agents in C , and \mathbf{o}_{-C} is $(o_b)_{b \notin C}$. Similarly, we let $\text{Ag}_{-C} = \text{Ag} \setminus C$.

Distributions. Let X be a finite non-empty set. A (*probability*) *distribution* over X is a function $\mathbf{d} : X \rightarrow [0, 1]$ such that $\sum_{x \in X} \mathbf{d}(x) = 1$, and $\text{Dist}(X)$ is the set of distributions over X . We write $x \in \mathbf{d}$ for $\mathbf{d}(x) > 0$. If $\mathbf{d}(x) = 1$ for some element $x \in X$, then \mathbf{d} is a *point* (a.k.a. *Dirac*) *distribution*. If, for $i \in I$, \mathbf{d}_i is a distribution over X_i , then, writing $X = \prod_{i \in I} X_i$, the *product distribution* of the \mathbf{d}_i is the distribution $\mathbf{d} : X \rightarrow [0, 1]$ defined by $\mathbf{d}(x) = \prod_{i \in I} \mathbf{d}_i(x_i)$.

Markov Chains. A *Markov chain* M is a tuple (St, p) where St is a set of states and $p \in \text{Dist}(St \times St)$ is a distribution. The values $p(s, t)$ are called *transition probabilities* of M . A *path* is an infinite sequence of states.

Concurrent Game Structures. A *stochastic concurrent game structure with imperfect information* (or simply *CGS*) \mathcal{G} is a tuple $(St, L, \delta, \ell, \{\sim_a\}_{a \in \text{Ag}})$ where (i) St is a finite non-empty set of *states*; (ii) $L : St \times \text{Ag} \rightarrow 2^{\text{Ac}} \setminus \{\emptyset\}$ is a *legality function* defining the available actions for each agent in each state, we write $L(\mathbf{s})$ for the tuple $(L(s, a))_{a \in \text{Ag}}$; (iii) for each state $s \in St$ and each move $\mathbf{c} \in L(\mathbf{s})$, the *stochastic transition function* δ gives the (conditional) probability $\delta(s, \mathbf{c})$ of a transition from state s for all $s' \in St$ if each player $a \in \text{Ag}$ plays the action \mathbf{c}_a , we also write this probability as $\delta(s, \mathbf{c})(s')$, to emphasize that $\delta(s, \mathbf{c})$ is a probability distribution on St ; (iv) $\ell : St \rightarrow 2^{\text{AP}}$ is a *labelling function*; (v) $\sim_a \subseteq St \times St$ is an equivalence relation called the *observation relation* of agent a .

Throughout this paper, we assume that the CGS is uniform, that is, if two states are indistinguishable for an agent a , then a has the same available actions in both states. Formally, if $s \sim_a s'$ then $L(s, a) = L(s', a)$, for any $s, s' \in St$ and $a \in \text{Ag}$. For each state $s \in St$ and joint action $\mathbf{c} \in \prod_{a \in \text{Ag}} L(s, a)$, we also assume that there is a state $s' \in St$ such that $\delta(s, \mathbf{c})(s')$ is non-zero, that is, every state has a successive state from a legal move.

We say that \mathcal{G} is *deterministic* (instead of stochastic) if every $\delta(s, \mathbf{c})$ is a point distribution.

Plays. A *play* or *path* in a CGS \mathcal{G} is an infinite sequence $\pi = s_0 s_1 \dots$ of states such that there exists a sequence $\mathbf{c}_0 \mathbf{c}_1 \dots$ of joint-actions such that $\mathbf{c}_i \in L(s_i)$ and $s_{i+1} \in \delta(s_i, \mathbf{c}_i)$ (i.e., $\delta(s_i, \mathbf{c}_i)(s_{i+1}) > 0$) for every $i \geq 0$. We write π_i for s_i , $\pi_{\geq i}$ for the suffix of π starting at position i . Finite paths are called *histories*, and the set of all histories is denoted Hist . Write $\text{last}(h)$ for the last state of a history h .

Strategies. A (general) *probabilistic strategy* is a function $\sigma : \text{Hist} \rightarrow \text{Dist}(\text{Ac})$ that maps each history to a distribution of actions. We let Str be the set of all strategies. A *memoryless uniform probabilistic strategy* for an agent a is a function $\sigma_a : St \rightarrow \text{Dist}(\text{Ac})$ in which for all positions s, s' such that $s \sim_a s'$, we have $\sigma(s) = \sigma(s')$. We let Str_a^r be the set of uniform strategies for agent a . A deterministic (or *pure*) strategy σ is a strategy in which $\sigma(s)$ is a point distribution for any s . A *strategy profile* is a tuple σ of strategies,

one for each agent. We write σ_a for the strategy of a in the strategy profile σ . For a strategy σ_a for agent a , we assume that $\sigma(h)(c) = 0$ if $c \notin L(\text{last}(h), a)$.

3 Probabilistic ATL and ATL*

We begin by introducing the Probabilistic Alternating-Time Temporal Logics PATL* and PATL.

The syntax of PATL* is defined by the grammar

$$\varphi ::= p \mid \varphi \vee \varphi \mid \neg\varphi \mid \mathbf{X}\varphi \mid \varphi \mathbf{U}\varphi \mid \langle\langle C \rangle\rangle^{\bowtie d} \varphi$$

where $p \in \text{AP}$, $C \subseteq \text{Ag}$, d is a rational constant in $[0, 1]$, and $\bowtie \in \{\leq, <, >, \geq\}$.

The intuitive reading of the operators is as follows: $\langle\langle C \rangle\rangle^{\bowtie d} \varphi$ means that there exists a strategy for the coalition C to collaboratively enforce φ with a probability in relation \bowtie with constant d , “next” \mathbf{X} and “until” \mathbf{U} are the standard temporal operators. We make use of the usual syntactic sugar $\mathbf{F}\varphi := \top \mathbf{U}\varphi$ and $\mathbf{G}\varphi := \neg \mathbf{F}\neg\varphi$ for temporal operators. Finally, we use $\llbracket C \rrbracket^{\bowtie d} \varphi := \neg \langle\langle C \rangle\rangle^{\bowtie d} \neg\varphi$ to express that no strategy of C can prevent φ with a probability in relation \bowtie with constant d .

An PATL* formula of the form $\langle\langle C \rangle\rangle^{\bowtie d} \varphi$ or $\llbracket C \rrbracket^{\bowtie d} \varphi$ is also called state formula. An important syntactic restriction of PATL*, namely PATL, is defined as follows.

The syntax of PATL is defined by the grammar

$$\varphi ::= p \mid \varphi \vee \varphi \mid \neg\varphi \mid \langle\langle C \rangle\rangle^{\bowtie d} \mathbf{X}\varphi \mid \langle\langle C \rangle\rangle^{\bowtie d} (\varphi \mathbf{U}\varphi)$$

where $p \in \text{AP}$, $C \subseteq \text{Ag}$, and $\bowtie \in \{\leq, <, >, \geq\}$.

Formulas of PATL and PATL* are interpreted over CGSs.

Probability Space on Outcomes. An *outcome* of a strategy profile σ and a state s is a play π that starts with s and is extended by σ , i.e., $\pi_0 = s$, and for every $k \geq 0$ there exists $c_k \in \sigma(\pi_k)$ such that $\pi_{k+1} \in \delta(\pi_k, c_k)$. The set of outcomes of a strategy profile σ and state s is denoted $\text{Out}(\sigma, s)$. A given system \mathcal{G} , strategy profile σ , and state s induce an infinite-state Markov chain $M_{\sigma, s}$ whose states are the finite prefixes of plays in $\text{Out}(\sigma, s)$. Such finite prefixes of plays are called *histories* and written h , and we let $\text{last}(h)$ denote the last state in h . Transition probabilities in $M_{\sigma, s}$ are defined as $p(h, hs') = \sum_{c \in \text{Ac}^{\text{Ag}}} \sigma(h)(c) \times \delta(\text{last}(h), c)(s')$. The Markov chain $M_{\sigma, s}$ induces a canonical probability space on its set of infinite paths (Kemeny, Snell, and Knapp, 1976), which can be identified with the set of plays in $\text{Out}(\sigma, s)$ and the corresponding measure is denoted $\text{out}(\sigma, s)$.¹

Given a coalition strategy $\sigma_C \in \prod_{a \in C} \text{Str}_a^r$, we let $n = |\text{Ag} \setminus \{C\}|$ and define the set of possible outcomes of σ_C from a state $s \in \text{St}$ to be the set $\text{out}_C(\sigma_C, s) = \{\text{out}((\sigma_C, \sigma_{-C}), s) : \sigma_{-C} \in \text{Str}^n\}$ of probability measures that the players in C enforce when they follow the strategy σ_C , namely, for each $a \in \text{Ag}$, player a follows strategy σ_a . We use $\mu_s^{\sigma_C}$ to range over $\text{out}_C(\sigma_C, s)$.

PATL and PATL* Semantics PATL and PATL* formulas are interpreted in a transition system \mathcal{G} and a path π ,

$$\begin{aligned} \mathcal{G}, \pi &\models p && \text{iff } p \in \ell(\pi_0) \\ \mathcal{G}, \pi &\models \neg\varphi && \text{iff } \mathcal{G}, \pi \not\models \varphi \\ \mathcal{G}, \pi &\models \varphi_1 \vee \varphi_2 && \text{iff } \mathcal{G}, \pi \models \varphi_1 \text{ or } \mathcal{G}, \pi \models \varphi_2 \\ \mathcal{G}, \pi &\models \langle\langle C \rangle\rangle^{\bowtie d} \varphi && \text{iff } \exists \sigma_C \in \prod_{a \in C} \text{Str}_a^r \text{ such that} \\ &&& \forall \mu_{\pi_0}^{\sigma_C} \in \text{out}_C(\sigma_C, \pi_0), \\ &&& \mu_{\pi_0}^{\sigma_C}(\{\pi' : \mathcal{G}, \pi' \models \varphi\}) \bowtie d \\ \mathcal{G}, \pi &\models \mathbf{X}\varphi && \text{iff } \mathcal{G}, \pi_{\geq 1} \models \varphi \\ \mathcal{G}, \pi &\models \psi_1 \mathbf{U}\psi_2 && \text{iff } \exists k \geq 0 \text{ s.t. } \mathcal{G}, \pi_{\geq k} \models \psi_2 \text{ and} \\ &&& \forall j \in [i, k]. \mathcal{G}, \pi_{\geq j} \models \psi_1 \end{aligned}$$

4 Strategic Reasoning under Uncertainty

Many real-life scenarios require agents to interact in partially observable environments with stochastic phenomena. A natural application of strategic reasoning over both of these sources of uncertainty is card games, as the distribution of cards is a stochastic event and the hand of each agent is kept secret from the other players. —Let us see a more detailed example based on online mechanism design² and, in particular, elections. While the majority of elections have a static set of candidates which is known upfront, there are contexts where candidates appear over time. A classic example is hiring a committee: the candidates that will appear the next day to pass an interview are unknown, and the voters must decide immediately whether to hire one of the current candidates or not (Do et al., 2022).

In online approval-based election (Do et al., 2022), there is a non-empty set of candidates $C = \{1, \dots, m\}$ and the goal is to select $k \leq 1$ candidates for a committee. In each state, an unseen candidate j is presented and the agents vote on whether to include the current candidate in the committee or not. The election continue until the committee is completed or all candidates have been rejected. For a candidate j , we let the propositions *rejected_j*, *selected_j*, *interview_j*, denote whether candidate j was already rejected, whether she was selected to the committee, and whether she is been currently interviewed, resp. For each agent a , *likes_{a,j}* denotes whether a is currently willing to approve the candidate j .

Agents know their own preferences, that is, the candidates they like but are uncertain about others’ preferences. Voters can distinguish the candidate currently interviewed, but are unaware of the next candidate to be presented (i.e., whether $\mathbf{Xinterview}_j$ holds in any given state).

In each state s , agents can either accept or reject the current candidate (actions y and n , resp.). The probability of selecting candidate j being selected is determined by the transition function $\delta(s, c)$, according to the actions in c . If all agents accept (similarly, reject) a candidate, the system transitions to a state in which the candidate is selected (resp.

¹This is a classic construction, see for instance (Clarke et al., 2018b; Berthon et al., 2020).

²Previous work (Maubert et al., 2021; Mittelmann et al., 2022, 2023) have shown how to encode notions from Mechanism Design (e.g., strategyproofness) using logics for strategic reasoning.

rejected) with a probability equal to one. If there is no consensus on whether to accept the candidate, the probability to transition to a state in which the candidate is selected is given by a rational constant $p_{j,c} \in (0, 1)$. Similarly, the probability of moving to a state where she is rejected is $1 - p_{j,c}$.

The PATL formula

$$\text{rejected}_j \rightarrow \neg \langle\langle C \rangle\rangle^{\geq 1} \mathbf{F} \text{selected}_j$$

represents that the coalition C cannot select a candidate that was already rejected.

The PATL* formula

$$\langle\langle C \rangle\rangle^{\geq \frac{1}{2}} \bigwedge_{a \in C} \bigvee_{j \in C} \text{likes}_{a,j} \wedge \mathbf{F} \text{selected}_j$$

represents that the coalition C can ensure, with probability greater or equal to $\frac{1}{2}$ to select in the future at least one candidate liked by each agent in a , while

$$\langle\langle C \rangle\rangle^{\geq \frac{1}{2}} \bigwedge_{a \in C} \bigwedge_{j \in C} \text{likes}_{a,j} \wedge \mathbf{F} \text{selected}_j$$

states that they can ensure, with probability greater or equal to $\frac{1}{2}$, all their liked candidates are eventually selected.

The formula

$$\text{interview}_j \rightarrow \langle\langle C \rangle\rangle^{\leq \frac{1}{4}} \mathbf{X} \text{selected}_j$$

says that the probability the coalition C ensures the currently interviewed candidate is selected in the next state is at most $\frac{1}{4}$.

5 Model Checking Complexity

In this section, we look at the complexity of model-checking for PATL. In particular, we show that the problem for *memoryless deterministic strategies of the coalition* against probabilistic play of the other agents and a stochastic environment is no more complex than in standard (non-probabilistic) case. The settings introduced in this paper include both deterministic and probabilistic memoryless strategies for the coalition and deterministic and stochastic CGSs. This gives 4 semantic variants in total, but the case of deterministic strategies and deterministic CGSs consists of the standard setting for ATL, whose complexity results are well-established.

The main technical result of this paper is as follows.

Theorem 1. Model checking PATL_{ir} ³ with deterministic strategies for the coalition is Δ_2^{P} -complete.

Proof. The lower bound follows by a reduction of ATL_{ir} model checking, which is Δ_2^{P} -hard (Jamroga and Dix, 2006). Given are: a pointed CGS (M, q) and a formula $\langle\langle C \rangle\rangle \varphi$ of ATL_{ir} . Note that M can be seen as stochastic CGS with only Dirac probability distributions for transitions. Recall that, in finite games, the opponents always have a deterministic best-response strategy to any given strategy σ_C . Thus, $M, q \models_{\text{ATL}_{\text{ir}}} \langle\langle C \rangle\rangle \varphi$ iff the agents in C have a uniform deterministic memoryless strategy to enforce φ on all paths

³As usual in the verification process, we denote no recall with r and imperfect information with i .

iff they have such a strategy against all the probabilistic responses from \bar{C} . Since the set of best responses includes deterministic strategies of \bar{C} played against deterministic strategy σ_C in the deterministic CGS M , this is equivalent to saying that $M, q \models_{\text{PATL}_{\text{ir}}} \langle\langle C \rangle\rangle^{\geq 1} \varphi$, which completes the reduction.

For the upper bound, we apply a similar procedure to that of ATL_{ir} (Schobbens, 2004). For formulas of type $\langle\langle C \rangle\rangle^{\bowtie d} \varphi$ without nested strategic modalities, we guess a strategy σ_C , prune the model accordingly, and merge the remaining agents (\bar{C}) into a single opponent. This yields a single-agent Markov Decision Process with full observability. Then, we check the Probabilistic Computation Tree Logic formula $A^{\bowtie d} \varphi$, which can be done in time $\text{NP} \cap \text{co-NP}$ (Chen and Lu, 2007).

For nested strategic modalities, we proceed recursively (bottom up), which runs in time $\text{P}^{\text{NP} \cap \text{co-NP}} = \Delta_2^{\text{P}}$. \square

6 Discussion

This paper analyses the verification of the strategic abilities of autonomous agents in MAS while accounting for both incomplete information and probabilistic behaviours of the environment and agents. The setting considered in this paper is significant as MAS are often set in partially observable environments, whose evolution might not be known with certainty, but can be measured based on experiments and past observations. Verification of strategic abilities in the general setting with perfect recall is known to be undecidable, but the restriction to memoryless strategies is meaningful. We provided complexity results for deterministic strategies for the proponent coalition and point out different settings that are currently challenging open questions, based on probabilistic strategies for the proponent coalition.

For solving the model checking problem w.r.t. probabilistic strategies for the proponent coalition, notice that it is not possible to exploit the technique used in Section 5 for deterministic strategies, i.e., calling an oracle that guesses the successful memoryless strategy. This is because there are *infinitely many* probabilistic memoryless strategies, and hence the oracle Turing machine would either have to run in unbounded time, or allow for infinite branching. In fact, the synthesis of optimal probabilistic strategies is a special case of *jointly constrained bilinear optimization*, which is a notoriously hard problem (Al-Khayyal, 1990). Additionally, techniques employed for partially observable Markov decision processes (see for instance (Vlassis, Littman, and Barber, 2012)) can not be easily adapted as they refer to single-agent abilities. Moreover, the work on Probabilistic Alternating μ -calculus (Song et al., 2019) seems unhelpful in our case. First, it is known that Probabilistic Alternating μ -calculus and PATL are incomparable (Bulling and Jamroga, 2011; Song et al., 2019). Second, the work (Song et al., 2019) only considers perfect information strategies. Finally, using the work on PSL (Aminof et al., 2019) does not seem the right direction either. Indeed, it only considers perfect information strategies. Additionally, the model checking problem for PSL is 3-EXPTIME-complete, while we expect a much lower complexity in our setting.

Acknowledgments

This research has been supported by the PRIN project RIPER (No. 20203FFYLK), the PNRR MUR project PE0000013-FAIR, the InDAM project “Strategic Reasoning in Mechanism Design”, the EU ICT-48 2020 project TAILOR (No. 952215), and the NCBR Poland/FNR Luxembourg projects STV (POLLUX-VII/1/2019) and SpaceVote (POLLUX-XI/14/SpaceVote/2023).

References

- Al-Khayyal, F. 1990. Jointly constrained bilinear programs and related problems: An overview. *Computers & Mathematics with Applications* 19(11):53–62.
- Alur, R.; Henzinger, T.; and Kupferman, O. 2002. Alternating-time temporal logic. *J. ACM* 49(5):672–713.
- Aminof, B.; Kwiatkowska, M.; Maubert, B.; Murano, A.; and Rubin, S. 2019. Probabilistic strategy logic. In *Proc. of IJCAI 2019*, 32–38. ijcai.org.
- Ballarini, P.; Fisher, M.; and Wooldridge, M. 2009. Uncertain agent verification through probabilistic model-checking. In *Safety and Security in Multiagent Systems*.
- Belardinelli, F.; Lomuscio, A.; Murano, A.; and Rubin, S. 2020. Verification of multi-agent systems with public actions against strategy logic. *Artif. Intell.* 285.
- Berthon, R.; Fijalkow, N.; Filiot, E.; Guha, S.; Maubert, B.; Murano, A.; Pinault, L.; Pinchinat, S.; Rubin, S.; and Serre, O. 2020. Alternating tree automata with qualitative semantics. *ACM Trans. Comput. Logic* 22(1):1–24.
- Berthon, R.; Maubert, B.; Murano, A.; Rubin, S.; and Vardi, M. Y. 2021. Strategy logic with imperfect information. *ACM Trans. Comput. Log.* 22(1):1–51.
- Berwanger, D., and Doyen, L. 2008. On the power of imperfect information. In Hariharan, R.; Mukund, M.; and Vinay, V., eds., *Proc. of FSTTCS 2008*, volume 2 of *LIPICs*, 73–82.
- Bulling, N., and Jamroga, W. 2009. What agents can probabilistically enforce. *Fundam. Informaticae* 93(1-3):81–96.
- Bulling, N., and Jamroga, W. 2011. Alternating epistemic mu-calculus. In *Proc. of IJCAI 2011*, 109–114. IJCAI/AAAI.
- Bulling, N., and Jamroga, W. 2014. Comparing variants of strategic ability: how uncertainty and memory influence general properties of games. *Journal of Autonomous Agents and Multi-Agent Systems* 28(3):474–518.
- Carayol, A.; Löding, C.; and Serre, O. 2018. Pure strategies in imperfect information stochastic games. *Fundamenta Informaticae* 160(4):361–384.
- Cermák, P.; Lomuscio, A.; Mogavero, F.; and Murano, A. 2018. Practical verification of multi-agent systems against SLK specifications. *Inf. Comput.* 261:588–614.
- Chatterjee, K.; Henzinger, T. A.; and Piterman, N. 2010. Strategy Logic. *Inf. Comput.* 208(6):677–693.
- Chen, T., and Lu, J. 2007. Probabilistic alternating-time temporal logic and model checking algorithm. In *Proc. of FSKD*, 35–39.
- Chen, T.; Kwiatkowska, M.; Parker, D.; and Simaitis, A. 2011. Verifying team formation protocols with probabilistic model checking. In *Proc. of CLIMA 2011*, LNCS 6814, 190–207. Springer.
- Clarke, E. M.; Henzinger, T. A.; Veith, H.; and Bloem, R. 2018a. *Handbook of Model Checking*. Springer Publishing Company, Incorporated, 1st edition.
- Clarke, E.; Grumberg, O.; Kroening, D.; Peled, D.; and Veith, H. 2018b. *Model checking*. MIT press.
- Dima, C., and Tiplea, F. 2011. Model-checking ATL under imperfect information and perfect recall semantics is undecidable. *CoRR* abs/1102.4225.
- Do, V.; Hervouin, M.; Lang, J.; and Skowron, P. 2022. Online approval committee elections. In *Proc. of IJCAI 2022*, 251–257. ijcai.org.
- Doyen, L., and Raskin, J.-F. 2011. Games with imperfect information: theory and algorithms. *Lectures in Game Theory for Computer Scientists* 10.
- Doyen, L. 2022. Stochastic games with synchronizing objectives. In *Proc. of LICS, LICS ’22*. New York, NY, USA: Association for Computing Machinery.
- Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. 2004. *Reasoning about knowledge*. MIT press.
- Fu, C.; Turrini, A.; Huang, X.; Song, L.; Feng, Y.; and Zhang, L. 2018. Model checking probabilistic epistemic logic for probabilistic multiagent systems. In *Proc. of IJCAI 2018*, 4757–4763.
- Gripon, V., and Serre, O. 2009. Qualitative concurrent stochastic games with imperfect information. In *Proc. of ICALP 2009*, 200–211. Springer.
- Gurov, D.; Goranko, V.; and Lundberg, E. 2022. Knowledge-based strategies for multi-agent teams playing against nature. *Artificial Intelligence* 309:103728.
- Gutierrez, J.; Hammond, L.; Lin, A. W.; Najib, M.; and Wooldridge, M. 2021. Rational Verification for Probabilistic Systems. In *Proc. of KR 2021*, 312–322.
- Hansson, H., and Jonsson, B. 1994. A logic for reasoning about time and reliability. *Formal Aspects of Computing* 6(5):512–535.
- Hao, J.; Song, S.; Liu, Y.; Sun, J.; Gui, L.; Dong, J. S.; and Leung, H.-f. 2012. Probabilistic model checking multi-agent behaviors in dispersion games using counter abstraction. In *Proc. of PRIMA 2012*, LNCS 7455, 16–30. Springer.
- Huang, X., and Luo, C. 2013. A logic of probabilistic knowledge and strategy. In *Proc. of AAMAS 2013*, 845–852.
- Huang, X.; Su, K.; and Zhang, C. 2012. Probabilistic alternating-time temporal logic of incomplete information and synchronous perfect recall. In *Proc. of AAAI 2012*, 765–771.

- Jamroga, W., and Ågotnes, T. 2007. Constructive knowledge: what agents can achieve under imperfect information. *J. Applied Non-Classical Logics* 17(4):423–475.
- Jamroga, W., and Bulling, N. 2011. Comparing variants of strategic ability. In *Proc. of IJCAI 2011*, 252–257. IJCAI/AAAI.
- Jamroga, W., and Dix, J. 2006. Model checking abilities under incomplete information is indeed Δ_2 -complete. In *Proc. of EUMAS 2006*, CEUR 223. CEUR-WS.org.
- Kemeny, J. G.; Snell, J. L.; and Knapp, A. W. 1976. Stochastic processes. In *Denumerable Markov Chains*. Springer. 40–57.
- Kupferman, O., and Vardi, M. Y. 2000. Synthesis with incomplete information. In *Advances in Temporal Logic*. Berlin: Springer. 109–127.
- Kwiatkowska, M.; Norman, G.; Parker, D.; Santos, G.; and Yan, R. 2022. Probabilistic model checking for strategic equilibria-based decision making: Advances and challenges. In *Proc. of MFCS 2022*, 4–22.
- Laroussinie, F., and Markey, N. 2015. Augmenting ATL with strategy contexts. *Inf. Comput.* 245:98–123.
- Lomuscio, A., and Pirovano, E. 2020. Parameterised verification of strategic properties in probabilistic multi-agent systems. In *Proc. of AAMAS 2020*, 762–770.
- Maubert, B.; Mittelmann, M.; Murano, A.; and Perrussel, L. 2021. Strategic reasoning in automated mechanism design. In *KR-21*.
- Mittelmann, M.; Maubert, B.; Murano, A.; and Perrussel, L. 2022. Automated synthesis of mechanisms. In *IJCAI-22*.
- Mittelmann, M.; Maubert, B.; Murano, A.; and Perrussel, L. 2023. Formal verification of bayesian mechanisms. In *Proc. of AAAI 2023*.
- Mogavero, F.; Murano, A.; Perelli, G.; and Vardi, M. Y. 2014. Reasoning about strategies: On the model-checking problem. *ACM Trans. Comput. Log.* 15(4).
- Nguyen, H. N., and Rakib, A. 2019. A probabilistic logic for resource-bounded multi-agent systems. In *Proc. of IJCAI 2019*, 521–527.
- Nisan, N.; Roughgarden, T.; Tardos, É.; and Vazirani, V. 2007. *Algorithmic Game Theory*. Cambridge Univ. Press.
- Reif, J. H. 1984. The complexity of two-player games of incomplete information. *Journal of Computer and System Sciences* 29(2):274–301.
- Schobbens, P. 2004. Alternating-time logic with imperfect recall. *Electr. Notes Theor. Comput. Sci.* 85(2):82–93.
- Song, F.; Zhang, Y.; Chen, T.; Tang, Y.; and Xu, Z. 2019. Probabilistic alternating-time μ -calculus. In *Proc. of AAAI 2019*, 6179–6186.
- Vlassis, N.; Littman, M. L.; and Barber, D. 2012. On the computational complexity of stochastic controller optimization in POMDPs. *ACM Trans. Comput. Theory* 4(4):12:1–12:8.
- Wan, W.; Bentahar, J.; and Hamza, A. B. 2013. Model checking epistemic–probabilistic logic using probabilistic interpreted systems. *Knowledge-Based Systems* 50:279–295.