Strategy Logic: Where are we?

Bastien Maubert[1], Aniello Murano[1]

[1]Università degli Studi di Napoli "Federico II", Italy

Abstract

We present Strategy Logic as a formalism for reasoning about strategies in multi-agent systems, and survey its main extensions, including imperfect information, epistemic reasoning, quantitative and probabilistic aspects, counting strategies, and hyperproperties.

Keywords

Logics, Multi-Agent Systems, Strategic Reasoning

1. Introduction

The study of logics for strategic reasoning in multi-agent systems started around a quarter of a century ago, with the introduction by Alur, Henzinger and Kupferman of Alternating-time Temporal Logic (ATL) [1]. Since then an active community has proposed and studied a wide array of logics to account for different aspects of multi-agent systems, such as imperfect information or probabilities, or express important concepts that were not captured by ATL, such as game equilibria. Among the most successful proposals is Strategy Logic, introduced in [2] and extended to the concurrent multi-agent setting in [3, 4]. Its syntax makes it very easy to express all sorts of complex game-theoretic solution concepts such as Nash equilibria or secure equilibria, its model-checking problem is decidable though non-elementary, and it contains many useful fragments with much better computational complexity [5, 6, 7, 8].

In this short paper we propose a brief overview of the different directions of extensions that have been studied so far for Strategy Logic, and discuss directions of interest that are yet to be explored.

2. Strategy Logic

We first recall briefly the syntax and semantics of classic Strategy Logic. Let us fix a finite set of *atomic propositions* AP, a finite set of *agents* or *players* Ag and a finite set of *variables* Var.

Definition 1. *The syntax of* SL *is defined by the following grammar:*

$$\varphi := p \mid \neg\varphi \mid \varphi \vee \varphi \mid \mathbf{X}\,\varphi \mid \varphi\,\mathbf{U}\,\varphi \mid \exists s.\varphi \mid (a, s)\varphi$$

where $p \in \mathrm{AP}$, $s \in \mathrm{Var}$, *and* $a \in \mathrm{Ag}$.

Boolean operators and temporal operators, $\mathbf{X}$ ("next") and $\mathbf{U}$ ("until"), have their usual meaning. The *strategy quantifier* $\exists s$ is a first-order-like quantification on strategies: $\exists s.\varphi$ reads as "there exists a strategy s such that φ holds", where s is a strategy variable. The *binding operator* (a, s) assigns a strategy to an agent: $(a, s)\varphi$ reads as "when agent a plays strategy s, φ holds".

We now provide the semantics of SL. We only sketch it and refer to [3] for details. Models of SL are *concurrent game structures*: at each turn each agent chooses an action, and the next state is determined by the current state and the joint action of all agents. A *strategy* maps *histories* (sequences of states) to actions, defining what to do in every situation, depending on the past and the current state. Strategies that have access to the whole history are called *perfect recall*. In contrast, when they can only see the

8th International Workshop on Artificial Intelligence and fOrmal VERification, Logic, Automata, and sYnthesis. OVERLAY 2026 @ FLoC, July 18–19, Lisbon, Portugal

bastien.maubert@gmail.com (B. Maubert); nello.murano@gmail.com (A. Murano)

0000-0002-9081-2920 (B. Maubert); 0000-0003-4876-3448 (A. Murano)

© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073

published 2026-11-09

CEUR-WS.org/Vol-4324/paper11.pdf

current state, i.e. they are functions from states to actions, they are called *memoryless*. The semantics for Boolean and temporal connectives is standard. Concerning strategy quantification, a formula $\exists s.\varphi$ holds if there exists a strategy σ such that φ holds when s evaluates to σ. Finally the binding $(a, s)\varphi$ holds if φ holds when agent a plays strategy s. For instance, if each agent $a_i \in \text{Ag} = \{a_1, \ldots, a_n\}$ has an LTL objective ψ_i, the following SL formula expresses the existence of a Nash equilibrium:

$$\exists s_1 \ldots \exists s_n (a_1, s_1) \ldots (a_n, s_n) \bigwedge_i \forall s_i' \left[((a_i, s_i')\psi_i) \rightarrow \psi_i \right]$$

It means that there exist strategies for each agent such that if some agent has a strategy to achieve its goal, then it is already achieved with its current strategy, so that they have no incentive to deviate.

The model-checking problem for SL is non-elementary decidable for perfect recall strategies. The complexity is very high, but it is important to stress the fact that it is non-elementary only in the size of the formula, which is typically small, and the number of exponentials depends on the number of nested strategic operators, which is usually low (rarely more than one or two). On the other hand the size of the model, which is typically big, contributes only polynomially to the complexity. When we restrict to memoryless strategies the problem becomes PSpace [9].

In the next sections we discuss the main existing extensions of Strategy Logic.

3. Imperfect information

The first extension of importance that we mention concerns imperfect information. A system has imperfect information if the agents only have a partial view of the system's state, for instance they may observe only some of its variables but ignore the values of some "hidden" variables. In game theory such situations are modeled by *games with imperfect information*, where agents' strategies must only depend on the available information. This constraint makes these games harder to solve than games with perfect information, and in particular the existence of a distributed winning strategy for a team of players in a game with imperfect information is undecidable already for reachability objectives [10]. However decidability can be recovered with two main types of restrictions, one being when there is a hierarchy among the different agents' knowledge, where each one knows more than the next (see e.g. [11]), the other being when all new information is public (as in [12]). These results have been extended to Strategy Logic: in [13] SL is extended to the imperfect-information setting, and the notion of hierarchical information gives rise to the notion of *hierarchical instances*, where the knowledge must increase as one goes down the syntactic tree of the formula. It is proven that for such instances model checking Strategy Logic with imperfect information is decidable. Concerning the second type of restriction that makes imperfect information manageable, it is proven in [14] that model checking Strategy Logic with imperfect information is decidable when all actions are public.

4. Epistemic reasoning

When agents do not know everything about the system, i.e. in the context of imperfect information, it becomes meaningful and useful to model and reason about what agents do or do not know. This is the purpose of epistemic logics, which are characterized by knowledge operators of the form $K_a\varphi$, meaning "agent a knows that φ holds" [15]. Since this is an important aspect in multi-agent systems, epistemic extensions of logics for strategic reasoning such as ATL have been studied, and Strategy Logic has been extended with epistemic operators.

As we said in the previous section, distributed strategy synthesis and thus logics for strategic reasoning are in general undecidable when agents have partial knowledge/imperfect information, and it remains so when we add knowledge operators. It was shown that the same kind of assumptions that make SL with imperfect information decidable can also yield decidability results when we add knowledge operators, but how one defines the semantics of these operators is crucial here.

As pointed out in [16], there exist in the literature two distinct semantics for knowledge operators in strategic contexts: one where agents know all agents' strategies, and one where they do not know anything about anyone's strategy. These semantics are respectively called *informed* and *uninformed*. The latter turns out to be computationally easier, and it is shown in [16] that model checking SL with knowledge operators and uninformed semantics is decidable for a class of hierarchical instances, while [14] proved that it is also decidable with public actions. With the informed semantics on the other hand the logic is undecidable even when information is hierarchical [16]. For public actions it was proved in [17] that decidability can be recovered, but only for a fragment of SL called "Boolean Goals" [18, 8].

5. Quantitative reasoning

Strategy Logic has also been extended to capture quantitative aspects of multi-agent systems.

5.1. Fuzzy Strategy Logic

A quantitative version of Strategy Logic that allows reasoning about the *quality* of systems is defined and studied in [19]. This is done by assigning to formulas not Boolean truth values, but rational values that represent quantitative aspects about the system, such as the distance between an agent and its target or an agent's reward. To do so atomic propositions are given rational values, and Boolean operations are given fuzzy logic semantics such as the Zadeh semantics. The temporal operators' semantics is taken from Fuzzy LTL [20, 21], and strategy quantification maximizes or minimizes over all possible strategies: the semantics of $\exists s.\varphi$ is the supremum of the semantics of φ over all strategies s. In addition arbitrary computable functions can be used to combine the values of subformulas, as in LTL$[\mathcal{F}]$ [22]. Examples of problems that can be formalized and solved in this very expressive setting include synthesis with quantitative objectives, or computing ε-equilibria in weighted games. It can even be used to reason about social welfare, for instance, by computing the Price of Anarchy or Price of Stability of a system. The model-checking problem is shown to be non-elementary decidable for perfect-recall strategies. In [23] the setting has been further investigated using a *discounting operator* to study economic scenarios in which distance in the future is an important parameter.

5.2. Quantifying bounds in Strategy Logic

Another work that introduces quantitative aspects in SL is [24]. It does so by introducing two bounded operators: $\mathbf{F}^N \psi$, which means "ψ is satisfied in at most N steps", and $\mathbf{A}^N \psi$, which means "there are at most N outcomes of the current strategies that satisfy ψ". It also allows one to quantify over this bound, and thus one can write $\exists s_1(a_1, s_1)\forall s_2(a_2, s_2)\exists N\, \mathbf{A}\,\mathbf{G}\,\mathbf{F}^{\leq N}\, p$, which expresses that there exists a strategy for agent a_1 such that for all strategies of agent a_2 there is a bound N that can depend on the strategy for a_2, such that in all outcomes p holds at least once every N steps. The bounded outcome quantifier can be used to write, for example, that there exists a strategy for an agent such that p holds infinitely often in all outcomes but finitely many: $\exists s(a, s)\exists N\, \mathbf{A}^{\leq N}\, \mathbf{G}\,\mathbf{F}\, p$. Here again the model-checking problem is shown to be non-elementary decidable for perfect-recall strategies.

6. Probabilistic reasoning

Yet another fundamental aspect of multi-agent systems is that the agents' actions or the outcomes of those actions may not be deterministic (think of flipping a coin). To model this several probabilistic extensions of logics for strategic reasoning have been defined, most of them based on ATL. A common shortcoming of all these proposals was that they lacked the expressiveness to reason about important game-theoretic solution concepts such as Nash equilibria, subgame perfect equilibria or secure equilibria. The only proposals that could express some of these concepts did so through dedicated operators, which requires introducing one operator for each concept.

In [25] a probabilistic extension of Strategy Logic is introduced. Unlike previous proposals for probabilistic strategic reasoning, it can naturally express with a fixed set of fundamental operators a wide array of concepts such as expected outcome, Nash equilibria or social-welfare maximization. It does so by allowing one to write arithmetic terms and to compare their values in formulas of the form $\tau \leq \tau'$, where the terms are formed using addition, product and their inverse operations on constants and atomic terms of the form $\mathbb{P}_\beta(\psi)$, which designates the probability that ψ holds when the agents follow the strategies assigned by $\beta : \mathcal{A}g \to \mathit{Str}$. The model-checking problem for this logic is proved to be undecidable for perfect-recall strategies, decidable in 3ExpSpace for memoryless strategies, and in ExpSpace for a fixed-size formula.

7. Counting strategies

Strategy Logic has also been extended to allow for counting strategies and strategy tuples that satisfy a given property [26, 27, 28, 29]. The main purpose in this endeavor is to express the uniqueness of Nash equilibria or other solution concepts expressible in SL. This is done by defining a counting strategic operator: $\exists^{\geq k}(s_1, \ldots, s_n)\varphi$ means that there exist at least k distinct tuples of strategies $(\sigma_1, \ldots, \sigma_n)$ that satisfy φ. So, for instance, if the formula $\Phi_{\text{NE}}(s_1, \ldots, s_n)$ expresses that $(s_1, \ldots, s_n)$ forms a Nash equilibrium for some given objectives, then the graded SL formula

$$\exists^{\geq 1}(s_1, \ldots, s_n)\Phi_{\text{NE}} \wedge \neg\exists^{\geq 2}(s_1, \ldots, s_n)\Phi_{\text{NE}}$$

expresses that there exists a unique Nash equilibrium.

It was proved in [29] that model checking graded Strategy Logic is decidable, and that it can be used to decide the existence of unique Nash equilibria in games with LTL objectives in 2ExpTime, which is no more than the complexity of deciding the existence of a solution.

8. Hyperproperties

While usual temporal logics such as LTL express properties of individual executions of systems, *hyperproperties* are properties of sets of executions. There is a trend of research into extending classical temporal logics to express such properties and evaluate them [30]. Recently Strategy Logic was extended to allow reasoning about properties of the outcomes of multiple strategy profiles [31]. The resulting logic, called Hyper Strategy Logic, can express properties that cannot be expressed in standard SL, such as non-interference or quantitative Nash equilibria. Model checking Hyper Strategy Logic is undecidable for perfect recall, but a useful fragment is identified for which the problem is non-elementary decidable.

There is a strong connection between Hyper Strategy Logic and Strategy Logic with imperfect information, and it was proven in [32] that forbidding strategic sub-formulas in path formulas of both logics leads to equiexpressive fragments. This resonates with the known connection between epistemic temporal logics and hyper temporal logics [33].

Conclusion and Future work

Strategy Logic has been extensively studied and extended in many directions during the past fifteen years. We have described most of them, but there are others, such as the works [34, 35], which consider nondeterministic, revocable, or refinable strategies, and the works [36, 37], which study natural or human-friendly strategies for agents with bounded memory. We would also like to mention the recent works [38, 39, 40, 41, 42].

Much remains to explore though: for instance other quantitative aspects have been studied in the settings of multiplayer games and ATL, such as resource-boundedness, but have not yet been studied in conjunction with SL. It would also be good to find a useful fragment of Probabilistic Strategy Logic for which the model checking problem is decidable with perfect recall strategies. Regarding

implementations, some prototypes exist, such as MCMAS-SLK [9, 43, 44] or Vitamin [45], but much work remains to be done to obtain mature tools that can be used for practical real-world purposes.

Acknowledgments

This work has been partially funded by the Italian MASE Project GAMEL (CUP FS3C25001420001) and the Horizon project COSMI ALGORITMICI (CUP E63C26000840001).

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

[1] R. Alur, T. A. Henzinger, O. Kupferman, Alternating-time temporal logic, Journal of the ACM 49 (2002) 672–713.
[2] K. Chatterjee, T. A. Henzinger, N. Piterman, Strategy logic, Information and Computation 208 (2010).
[3] F. Mogavero, A. Murano, G. Perelli, M. Y. Vardi, Reasoning about strategies: On the model-checking problem, ACM Transactions on Computational Logic 15 (2014) 34:1–34:47.
[4] F. Mogavero, A. Murano, G. Perelli, M. Y. Vardi, Reasoning about strategies: on the satisfiability problem, Log. Methods Comput. Sci. 13 (2017).
[5] F. Belardinelli, W. Jamroga, V. Malvone, A. Murano, Strategy logic with simple goals: Tractable reasoning about strategies, in: Proc. of IJCAI'19, 2019, pp. 88–94.
[6] D. Catta, R. Galimullin, A. Murano, First-order coalition logic, in: Proc. of IJCAI'25, ijcai.org, 2025, pp. 4410–4418.
[7] P. Gardy, P. Bouyer, N. Markey, Dependences in strategy logic, Theory of Computing Systems 64 (2020) 467–507.
[8] F. Mogavero, A. Murano, G. Perelli, M. Y. Vardi, What makes ATL* decidable? A decidable fragment of strategy logic, in: Proc. of CONCUR'12, Springer, 2012, pp. 193–208.
[9] P. Čermák, A. Lomuscio, F. Mogavero, A. Murano, MCMAS-SLK: A model checker for the verification of strategy logic specifications, in: Proc. of CAV'14, Springer, 2014, pp. 525–532.
[10] G. Peterson, J. H. Reif, S. Azhar, Lower bounds for multiplayer noncooperative games of incomplete information, Computers & Mathematics with Applications 41 (2001) 957–992.
[11] G. Peterson, J. Reif, S. Azhar, Decision algorithms for multiplayer noncooperative games of incomplete information, Computers & Mathematics with Applications 43 (2002) 179–206.
[12] R. van der Meyden, T. Wilke, Synthesis of distributed systems from knowledge-based specifications, in: Proc. of CONCUR'05, 2005, pp. 562–576.
[13] R. Berthon, B. Maubert, A. Murano, S. Rubin, M. Y. Vardi, Strategy logic with imperfect information, ACM Trans. on Comp. Logic (TOCL) 22 (2021) 1–51.
[14] F. Belardinelli, A. Lomuscio, A. Murano, S. Rubin, Verification of multi-agent systems with public actions against strategy logic, Artificial Intelligence 285 (2020) 103302.
[15] R. Fagin, J. Y. Halpern, Y. Moses, M. Y. Vardi, Reasoning about knowledge, volume 4, MIT press Cambridge, 1995.
[16] B. Maubert, A. Murano, Reasoning about knowledge and strategies under hierarchical information, in: Proc. of KR'18, 2018, pp. 530–540.
[17] F. Belardinelli, S. Knight, A. Lomuscio, B. Maubert, A. Murano, S. Rubin, Reasoning about agents that may know other agents' strategies, in: Proc. of IJCAI'21, IJCAI/AAAI Press, 2021.
[18] F. Mogavero, A. Murano, L. Sauro, On the boundary of behavioral strategies, in: Proc. of LICS'13, IEEE, 2013, pp. 263–272.
[19] P. Bouyer, O. Kupferman, N. Markey, B. Maubert, A. Murano, G. Perelli, Reasoning about quality and fuzziness of strategic behaviors, ACM Transactions on Computational Logic 24 (2023) 1–38.
[20] K. B. Lamine, F. Kabanza, Using fuzzy temporal logic for monitoring behavior-based mobile robots, in: Proc. of IASTED Int. Conf. on Robotics and Applications, 2000, pp. 116–121.

[21] A. Frigeri, L. Pasquale, P. Spoletini, Fuzzy time in linear temporal logic, ACM Transactions on Computational Logic (TOCL) 15 (2014) 1–22.
[22] S. Almagor, U. Boker, O. Kupferman, Formally reasoning about quality, Journal of the ACM (JACM) 63 (2016) 1–56.
[23] M. Mittelmann, A. Murano, L. Perrussel, Discounting in strategy logic, in: Proc. of IJCAI'23, ijcai.org, 2023, pp. 225–233.
[24] N. Fijalkow, B. Maubert, A. Murano, S. Rubin, Quantifying bounds in strategy logic, in: CSL, volume 23, 2018.
[25] B. Aminof, M. Kwiatkowska, B. Maubert, A. Murano, S. Rubin, Probabilistic strategy logic, in: Proc. of IJCAI'19, 2019, pp. 32–38.
[26] B. Aminof, V. Malvone, A. Murano, S. Rubin, Graded strategy logic: Reasoning about uniqueness of nash equilibria, in: Proc. of AAMAS'16, ACM, 2016, pp. 698–706.
[27] V. Malvone, F. Mogavero, A. Murano, L. Sorrentino, Reasoning about graded strategy quantifiers, Inf. Comput. 259 (2018) 390–411.
[28] V. Malvone, F. Mogavero, A. Murano, L. Sorrentino, On the counting of strategies, in: Proc. of TIME 2015, IEEE, 2015, pp. 170–179.
[29] B. Aminof, V. Malvone, A. Murano, S. Rubin, Graded modalities in strategy logic, Information and Computation 261 (2018) 634–649.
[30] M. R. Clarkson, B. Finkbeiner, M. Koleini, K. K. Micinski, M. N. Rabe, C. Sánchez, Temporal logics for hyperproperties, in: Proc. of POST 2014, Springer, 2014, pp. 265–284.
[31] R. Beutner, B. Finkbeiner, Hyper strategy logic, in: Proc. of AAMAS'24, IFAAMAS, 2024, p. 189–197.
[32] R. Beutner, B. Finkbeiner, Strategy Logic, Imperfect Information, and Hyperproperties, in: Proc. of KR'25, 2025, pp. 148–153.
[33] L. Bozzelli, B. Maubert, S. Pinchinat, Unifying hyper and epistemic temporal logics, in: Proc. of FoSSaCS'15, Springer, 2015, pp. 167–182.
[34] G. De Giacomo, B. Maubert, A. Murano, et al., Nondeterministic strategies and their refinement in strategy logic, in: Proc. of KR'20, 2020, pp. 294–303.
[35] C. Chareton, J. Brunel, D. Chemouil, A logic with revocable and refinable strategies, Information and Computation 242 (2015) 157–182.
[36] W. Jamroga, V. Malvone, A. Murano, Natural strategic ability under imperfect information, in: Proc. of AAMAS'19, volume 2, 2019, pp. 962–970.
[37] F. Belardinelli, W. Jamroga, V. Malvone, M. Mittelmann, A. Murano, L. Perrussel, Reasoning about human-friendly strategies in repeated keyword auctions, in: Proc. of AAMAS'22, 2022, pp. 62–71.
[38] R. Jin, S. Li, Y. Liu, Enhancing strategy logic with procedural rationality, in: Proc. of AAAI'26, volume 40, 2026, pp. 19169–19177.
[39] F. Belardinelli, A. Ferrando, W. Jamroga, V. Malvone, A. Murano, Scalable verification of strategy logic through three-valued abstraction, in: Proc. of IJCAI'23, 2023, pp. 46–54.
[40] M. Mittelmann, B. Maubert, A. Murano, L. Perrussel, Formal verification and synthesis of mechanisms for social choice, Artif. Intell. 339 (2025) 104272.
[41] Y. He, M. Mittelmann, A. Murano, A. Saffidine, M. Thielscher, Verification of general games with imperfect information using strategy logic, in: Proc. of KR'24, 2024.
[42] C. Mu, N. Motamed, N. Alechina, B. Logan, Probabilistic strategy logic with degrees of observability, in: Proc. of AAAI'25, volume 39, 2025, pp. 15092–15099.
[43] P. Čermák, A. Lomuscio, F. Mogavero, A. Murano, Practical verification of multi-agent systems against SLK specifications, Information and Computation 261 (2018) 588–614.
[44] P. Čermák, A. Lomuscio, A. Murano, Verifying and synthesising multi-agent systems against one-goal strategy logic specifications, in: Proc. of AAAI'15, volume 29, 2015.
[45] A. Ferrando, V. Malvone, Vitamin: A compositional framework for model checking of multi-agent systems, arXiv preprint arXiv:2403.02170 (2024).