Reasoning About Intuitionistic Alternating-Time Temporal Logic

Andrea Capone

University of Naples Federico II, Via Claudio 21, 80127, Naples (NA), Italy, IT

Abstract

Multi-Agent Systems (MAS) require formalisms capable of handling partial and evolving information. While imperfect-information extensions of Alternating-time Temporal Logic (ATL) are highly expressive, they typically make model checking undecidable. This paper summarises *Intuitionistic ATL* (IATL), an extension of ATL that replaces classical propositional truth with intuitionistic truth, capturing a novel, computationally tractable notion of imperfect information based on *information refinement*. We summarize IATL's syntax and semantics over *birelational concurrent game structures* (BCGS), the truth-monotonicity characterisation, and the PTIME-completeness of model checking, matching the complexity of classical ATL.

Keywords

Alternating-Time Temporal Logic, Intuitionistic Logic, Multi-Agent Systems, Model Checking

1. Introduction

Formal methods for strategic reasoning play a fundamental role in the specification and verification of Multi-Agent Systems (MAS) [1, 2, 3, 4, 5]. This line of research originated from the use of temporal logics for the specification of reactive systems [6, 7, 8], traditionally interpreted over Kripke structures. The need to reason about open systems with multiple autonomous agents led to the development of strategic formalisms [1, 9, 10, 11, 12], among which Alternating-time Temporal Logic (ATL) [1] stands as a central reference, enabling the specification of cooperative and competitive strategies among agents.

Imperfect information is pervasive in realistic MAS scenarios [13, 14, 15, 16], yet its integration with ATL is notoriously problematic: when agents employ strategies with perfect recall, model checking becomes undecidable [17]. Existing tractable alternatives, such as memoryless or bounded-memory strategies [18, 19, 20, 21, 22, 23, 24, 25, 26, 27], mitigate this at the cost of expressiveness.

Classical propositional logic assumes complete information: the law of the excluded middle $\varphi \vee \neg\varphi$ holds universally, so every formula is either true or false in a model. This assumption is inadequate for agents operating under partial or evolving knowledge [28, 29, 30]. Intuitionistic logic (IL) [31, 32, 33] drops this assumption, making truth procedural and dependent on the ability to verify a formula. The best known semantics for IL is based on Kripke models [34], where the accessibility relation is a partial order modelling progressive information accumulation: a formula is true at a state s only if it remains true in all more informative states reachable from s. This yields a natural *truth-monotonicity* property, making IL a fit base for reasoning about information refinement. Intuitionistic extensions of modal logics have been explored in [35, 36, 37] via birelational Kripke models, and more recently intuitionistic versions of temporal frameworks have been studied for LTL [38, 39] and CTL [40, 41]. Additionally, IL has played a role in non-monotonic reasoning via the Answer Set Programming paradigm [42] and its temporal extensions [43, 44]. The most recent development in this thread is the extension of IL to strategic reasoning in MAS, yielding IATL [45].

This paper summarises *Intuitionistic ATL* (IATL) [45], an extension of ATL that replaces classical propositional truth with intuitionistic truth to obtain a tractable forr of imperfect information. One of

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

andrea.capone@unina.it (A. Capone)

https://andrewbighead.github.io/andreacapone.github.io/ (A. Capone)

0009-0002-7708-9261 (A. Capone)



© 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/paper6.pdf

the core features of IATL is its ability to model *information refinement*: by equipping standard concurrent game structures with a partial order over states, agents can reason about strategies that are winning across all epistemic refinements of their current state, capturing the gradual accumulation of knowledge over time. This is formalised via *birelational concurrent game structures* (BCGS). Unlike imperfect-information ATL, where model checking is undecidable [17], IATL [45] achieves PTIME-completeness, matching classical ATL [1] and making it the first imperfect-information extension of ATL with tractable model checking. We recall IATL's syntax, semantics, and truth-monotonicity characterisation, and discuss an efficient implementation in the VITAMIN model checker [46], with benchmarks confirming practical scalability; we refer to [45] for full details.

IATL exploits intuitionistic truth to represent states where a fact may be unknown, yet sufficient information may still be available to guarantee a winning strategy. Unlike imperfect-information ATL, the law of the excluded middle does not hold in IATL, and model checking remains PTIME-complete. The two frameworks are expressively incomparable: imperfect-information ATL can explicitly model information unavailable to agents, making it suitable for information-flow security analysis, whereas IATL prioritises tractability and the modelling of incremental knowledge acquisition.

2. Preliminaries

We fix a finite set AP of atomic propositions, a finite set Ag of agents, and an action set Act. We assume familiarity with standard ATL [1] and recall only what is needed for IATL.

Intuitionistic Propositional Logic. An *Intuitionistic Kripke Structure* (IKS) is a tuple $\mathcal{K} = \langle S, S_I, \preceq, \mathrm{V}\rangle$, where $\preceq$ is a partial order over S modelling information refinement and $\mathrm{V} : S \to 2^{\mathsf{AP}}$ is monotone w.r.t. $\preceq$. The satisfaction relation for IPL formulas is standard except for implication:

$$\mathcal{K}, s \models \varphi_1 \to \varphi_2 \;\Leftrightarrow\; \forall\, t \geq s :\; \mathcal{K}, t \models \varphi_1 \Rightarrow \mathcal{K}, t \models \varphi_2.$$

This yields *truth monotonicity* ($s \preceq t$ and $\mathcal{K}, s \models \varphi$ imply $\mathcal{K}, t \models \varphi$) and refutes the excluded middle $\varphi \vee \neg\varphi$.

Concurrent Game Structures and Strategies. A CGS $\mathcal{G} = \langle S, S_I, \delta, \mathrm{V}\rangle$ is a standard multi-agent game structure [1] with transition function $\delta : S \times \mathrm{Act}^{\mathsf{Ag}} \to S$. Strategies, plays, histories, and outcome function $\mathrm{out}(s, f_A)$ are defined as in [1]; we consider both perfect-recall and memoryless strategies.

3. Intuitionistic ATL

Syntax. Syntax of IATL is as for standard ATL, with the addition of the intuitionistic implication $\to$:

$$\varphi ::= \bot \mid p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \to \varphi \mid \langle\!\langle A \rangle\!\rangle \psi \mid [\![A]\!]\, \psi \qquad \psi ::= \bigcirc \varphi \mid \varphi \mathsf{U} \varphi \mid \varphi \mathsf{R} \varphi$$

where $p \in \mathsf{AP}$, $A \subseteq \mathsf{Ag}$, and $\langle\!\langle A \rangle\!\rangle / [\![A]\!]$ are the existential/universal strategic quantifiers. Unlike classical ATL, $\langle\!\langle A \rangle\!\rangle$ and $[\![A]\!]$ are *not* interdefinable in IATL, nor are U and R.

Birelational Concurrent Game Structures. IATL formulas are interpreted over *Birelational CGS* (BCGS), which extend CGS with a partial order $\preceq$ over states modelling information refinement. Formally, a BCGS is a tuple $\mathcal{G} = \langle S, S_I, \delta, \preceq, \mathrm{V}\rangle$ where $\langle S, S_I, \delta\rangle$ is a CGF, $\preceq$ is a partial order over S, and $\mathrm{V} : S \to 2^{\mathsf{AP}}$ satisfies the *monotonicity condition*: $s \preceq t \Rightarrow \mathrm{V}(s) \subseteq \mathrm{V}(t)$.

A BCGS is *well-behaved* if $\preceq$ satisfies two conditions regulating the interplay between information refinement and agent actions. Let $\mathrm{Dec}(s, A)$ denote the A-decisions available at state s:

(C_1) For all $s \preceq s'$ and each $d_A \in \mathrm{Dec}(s, A)$, there exists $d'_A \in \mathrm{Dec}(s', A)$ such that for each full decision d' extending d'_A, there is a full decision d extending d_A with $\delta(s, d) \preceq \delta(s', d')$.

(C_2) The symmetric condition, with $s \preceq s'$ replaced by $s' \preceq s$ and the successor ordering reversed.

Intuitively, C_1 ensures that $\preceq$ behaves as an alternating-time simulation [1]: if s' is more informative than s, then every A-decision at s has a counterpart at s' whose outcomes are at least as informative. C_2 is the dual condition. Well-behavedness can be checked in polynomial time on finite structures [1, 45]. The semantics of IATL extends that of ATL by interpreting implication intuitionistically.

Theorem 1. *Every IATL-satisfiable formula is also ATL-satisfiable, but not vice versa. Adding the law of the excluded middle $\varphi \vee \neg\varphi$ causes IATL to collapse to standard ATL.*

Theorem 2 (Truth Monotonicity)**.** *The key semantic property of IATL is that truth is preserved under information refinement. A BCGS is* IATL-monotonic *if $\mathcal{G}, s \models \varphi$ and $s \preceq t$ imply $\mathcal{G}, t \models \varphi$ for all states s, t and formulas φ.*

The following characterisation shows that well-behavedness is precisely the right structural condition:

Theorem 3 (IATL-monotinicity)**.** *Let $\mathcal{F}$ be a BCGF. Then:* (i) *$\mathcal{F}$ satisfies C_1 iff it is IATL$_\exists$-monotonic;* (ii) *$\mathcal{F}$ satisfies C_2 iff it is IATL$_\forall$-monotonic;* (iii) *$\mathcal{F}$ is well-behaved iff it is IATL-monotonic.*

Theorem 4. *IATL$_\exists$ and IATL$_\forall$ are expressively incomparable, and IATL$_\mathsf{U}$ and IATL$_\mathsf{R}$ are expressively incomparable, independently of the well-behaved assumption.*

Theorem 5. *For any finitely-branching BCGS $\mathcal{G}$, state s of $\mathcal{G}$, and IATL formula φ: $\mathcal{G}, s \models \varphi \Leftrightarrow \mathcal{G}, s \models_m \varphi$, independently of the well-behaved assumption.*

For a formula φ and a BCGS $\mathcal{G}$, we denote by $[\![\varphi]\!]_\mathcal{G}$ the set of states satisfying φ. Given $X \subseteq S$, we write $X^{\mathbf{c}}$ for its complement and $X^{\uparrow}$ for the upward closure $\{s \in S \mid \forall s' \geq s,\ s' \in X\}$. Moreover, for a coalition A, we write $\mathrm{Pre}^{\exists}(A, X)$ for the set of states s such that for *some* A-decision d_A available at s, *all the successors* of s consistent with d_A are in X. Finally, we denote by $\mathrm{Pre}^{\forall}(A, X)$, the set of states s such that for *all* A-decisions d_A available at s, *some successor* of s consistent with d_A is in X. The following proposition easily follows from the semantics of IATL.

Proposition 1. *Given a BCGS $\mathcal{G}$ with set of states S and a coalition A, the following holds:*

- $[\![\varphi_1 \to \varphi_2]\!] = ([\![\varphi_1]\!]^{\mathbf{c}} \cup [\![\varphi_2]\!])^{\uparrow}$.
- $[\![\langle\!\langle A \rangle\!\rangle \bigcirc \varphi]\!] = \mathit{Pre}^{\exists}(A, [\![\varphi]\!])$.
- $[\![[\![A]\!] \bigcirc \varphi]\!] = \mathit{Pre}^{\forall}(A, [\![\varphi]\!])$.
- $(\mathit{Pre}^{\forall}(A, [\![X]\!]))^{\mathbf{c}} = \mathit{Pre}^{\exists}(A, [\![X]\!]^{\mathbf{c}})$ *for each* $X \subseteq S$.

The above characterisations are specific to IATL; denotational and fixpoint identities are as in ATL [1].

4. IATL Model Checking

The *model-checking problem* for IATL asks, given a finite BCGS $\mathcal{G}$ and an IATL formula Φ, whether $\mathcal{G}, s \models \Phi$ for each initial state s.

Theorem 6 (IATL Model Checking Complexity)**.** *IATL model checking is* Ptime*-complete and solvable in time $O(|\mathcal{G}|^2 \cdot |\Phi|)$.*

Ptime-hardness follows from Ptime-hardness of ATL model checking [1]: any CGS $\mathcal{G}$ can be viewed as a BCGS equipped with the identity partial order, reducing ATL model checking to IATL model checking. For the upper bound, the algorithm computes the denotation $[\![\varphi]\!]$ of each subformula φ of Φ by a bottom-up traversal of its parse tree. All cases coincide with standard ATL model checking, except for intuitionistic implication, handled via upward closure:

$$[\![\varphi_1 \to \varphi_2]\!] \;:=\; ([\![\varphi_1]\!]^{\complement} \cup [\![\varphi_2]\!])^{\uparrow}.$$

Each subformula is processed once; computing the upward closure $X^{\uparrow}$ takes $O(|S|^2)$, yielding the overall $O(|\mathcal{G}|^2 \cdot |\Phi|)$ bound.

Implementation. The algorithm has been implemented by extending the VITAMIN model checker [46] to support IATL semantics. Benchmarks on the muddy-children scenario [47] confirm practical scalability as the number of agents grows; reference to [45] for full experimental details.

5. Conclusion and Future Work

This paper summarises IATL [45], an extension of ATL that replaces classical propositional truth with intuitionistic truth to model information refinement in MAS. Unlike classical ATL, which assumes complete knowledge, IATL allows facts to be neither true nor false at a state, yet still guarantees the existence of winning strategies across all epistemic refinements. This makes IATL suitable for realistic scenarios such as sensing, incremental knowledge acquisition, and distributed protocol verification. Crucially, model checking remains PTIME-complete, matching classical ATL while gaining expressive power for imperfect-information settings. By providing a tractable logical foundation for reasoning under partial knowledge, IATL contributes to the growing agenda of formal approaches to AI verification [2, 3], particularly in safety-critical MAS where incomplete information is the rule.

Several directions are open. A complete axiomatization would unlock satisfiability and validity checking, while extending the intuitionistic approach to more expressive logics such as ATL* or Strategy Logic [3] is another natural avenue. Combining intuitionistic truth with the fuzzy semantics and resource-bounded natural strategies of HumanATL[$\mathcal{F}$] [48] would yield a framework where strategic truth is simultaneously graded and monotone with respect to information refinement, unifying epistemic and quantitative uncertainty. On the applied side, recent work [49] has deployed ATL model checking within game engines for real-time NPC decision-making: extending this in-engine framework with IATL semantics would let NPCs reason strategically under the asymmetric information that classical ATL cannot handle. The neurosymbolic architecture of [50], where a logical component progressively eliminates hypotheses incompatible with evidence, is conceptually aligned with the information refinement paradigm of IATL, which could replace hypothesis elimination with epistemic refinement to keep diagnostic strategies robust to future observations. Finally, pairing IATL with *neuro-fuzzy* architectures such as [51] would feed a quantitative, fuzzy variant of the logic in the spirit of quantitative ATL [52], with the neural component emitting graded evidence and fuzzy IATL combining monotonicity under refinement with graded goal satisfaction in a single learnable framework.

Acknowledgments

This work has been partially funded by the following projects: "Soluzioni Hardware e software per l'INgegnerizzazionE di una piattaforma Quantistica di Computazione – SHINE-QC" (Prog. n. F/350362/03/X60, CUP B69J25000260005, COR 24282027); "Game Elettrico – GAMEL" (Prog. n. MI_DDR_00423, CUP F53C25001420001); and "Cosmi Algoritmici" (CUP E63C26000840001).

Declaration on Generative AI

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

References

[1] R. Alur, T. Henzinger, O. Kupferman, Alternating-time temporal logic, Journal of the ACM 49 (2002) 672–713.

[2] O. Kupferman, M. Vardi, P. Wolper, Module checking, Information and Computation 164 (2001) 322–344.

[3] F. Mogavero, A. Murano, G. Perelli, M. Vardi, Reasoning about strategies: On the model-checking problem, ACM Transactions on Computational Logic 15 (2014) 1–47.

[4] A. Pnueli, R. Rosner, On the synthesis of a reactive module, in: Proc. 16th POPL, ACM Press, 1989, pp. 179–190.
[5] M. Wooldridge, An introduction to multiagent systems, John wiley & sons, 2009.
[6] E. Clarke, E. Emerson, Design and synthesis of synchronization skeletons using branching time temporal logic, in: Workshop on Logics of Programs, LNCS 131, Springer, 1981, pp. 52–71.
[7] E. Emerson, J. Halpern, "Sometimes" and "Not Never" revisited: on branching versus linear time temporal logic, Journal of the ACM 33 (1986) 151–178.
[8] A. Pnueli, The Temporal Logic of Programs, in: Proc. 18th FOCS, IEEE Computer Society, 1977, pp. 46–57.
[9] F. Laroussinie, N. Markey, Augmenting ATL with strategy contexts, Information and Computation 245 (2015) 98–123.
[10] K. Chatterjee, T. A. Henzinger, N. Piterman, Strategy logic, Information and Computation 208 (2010) 677–693.
[11] F. Mogavero, A. Murano, G. Perelli, M. Vardi, What makes ATL* decidable? A decidable fragment of strategy logic, in: Proc. 23rd CONCUR, LNCS 7454, Springer, 2012, pp. 193–208.
[12] F. Mogavero, A. Murano, L. Sauro, On the boundary of behavioral strategies, in: 2013 28th Annual ACM/IEEE Symposium on Logic in Computer Science, IEEE, 2013, pp. 263–272.
[13] F. Belardinelli, A. Lomuscio, A. Murano, S. Rubin, Verification of multi-agent systems with public actions against strategy logic, Artificial Intelligence 285 (2020) 103302.
[14] R. Berthon, B. Maubert, A. Murano, S. Rubin, M. Vardi, Strategy logic with imperfect information, ACM Transactions on Computational Logic 22 (2021) 1–51.
[15] W. Jamroga, N. Bulling, Comparing variants of strategic ability, in: Proc. 22nd IJCAI, IJCAI/AAAI, 2011, pp. 252–257.
[16] J. Reif, The complexity of two-player games of incomplete information, Journal of Computer and System Sciences 29 (1984) 274–301.
[17] C. Dima, F. Tiplea, Model-checking ATL under imperfect information and perfect recall semantics is undecidable, CoRR abs/1102.4225 (2011).
[18] T. Ågotnes, V. Goranko, W. Jamroga, M. Wooldridge, Knowledge and ability, in: Handbook of Epistemic Logic, College Publications, 2015, pp. 543–589.
[19] N. Bulling, W. Jamroga, Comparing variants of strategic ability: how uncertainty and memory influence general properties of games, Autonomous agents and multi-agent systems 28 (2014) 474–518.
[20] A. Lomuscio, F. Raimondi, Model checking knowledge, strategies, and games in multi-agent systems, in: Proc. 5th AAMAS, ACM, 2006, pp. 161–168.
[21] W. van der Hoek, M. Wooldridge, Tractable multiagent planning for epistemic goals, in: Proc. 1st AAMAS, ACM, 2002, pp. 1167–1174.
[22] B. Maubert, A. Murano, Reasoning about knowledge and strategies under hierarchical information, in: Proc. 16th KR, AAAI Press, 2018, pp. 530–540.
[23] P. Čermák, A. Lomuscio, F. Mogavero, A. Murano, Practical verification of multi-agent systems against SLK specifications, Information and Computation 261 (2018) 588–614.
[24] W. Jamroga, V. Malvone, A. Murano, Natural strategic ability, Artificial Intelligence 277 (2019).
[25] W. Jamroga, V. Malvone, A. Murano, Natural strategic ability under imperfect information, in: Proc. of 18th AAMAS'19, IFAAMAS, 2019, pp. 962–970.
[26] W. Jamroga, W. Van Der Hoek, Agents that know how to play, Fundamenta Informaticae 63 (2004) 185–219.
[27] W. Jamroga, T. Ågotnes, Constructive knowledge: what agents can achieve under imperfect information, Journal of Applied Non-Classical Logics 17 (2007) 423–475.
[28] M. Dummett, Elements of intuitionism, volume 39, Oxford University Press, 2000.
[29] J. Van Benthem, The information in intuitionistic logic, Synthese 167 (2009) 251–270.
[30] D. van Dalen, A. Troelstra, Constructivism in Mathematics, North-Holland, Amsterdam, 1988.
[31] S. Bauer-Mangelberg, J. van Heijenoort, S. Bauer-Mengelberg, On the significance of the principle of excluded middle in mathematics, especially in function theory, Journal of Symbolic Logic 35

(1970) 332–333.
[32] L. Brouwer, Intuitionism and formalism, in: Philosophy and Foundations of Mathematics, Elsevier, 1975, pp. 123–138.
[33] J. Moschovakis, Intuitionistic Logic, in: The Stanford Encyclopedia of Philosophy, Summer 2023 ed., Metaphysics Research Lab, Stanford University, 2023.
[34] S. Kripke, Semantical analysis of intuitionistic logic. I, in: Studies in Logic and the Foundations of Mathematics, volume 40, Elsevier, 1965, pp. 92–130.
[35] G. Fischer Servi, On modal logic with an intuitionistic base, Studia Logica 36 (1977) 141–149.
[36] G. Plotkin, C. Stirling, A framework for intuitionistic modal logics, in: Proc. 1st TARK, Morgan Kaufmann, 1986, pp. 399–406.
[37] A. Simpson, The proof theory and semantics of intuitionistic modal logic, Ph.D. thesis, University of Edinburgh, UK, 1994.
[38] P. Balbiani, M. Diéguez, Temporal Here and There, in: Proc. 15th JELIA, volume 10021 of *LNCS*, 2016, pp. 81–96.
[39] P. Balbiani, J. Boudou, M. Diéguez, D. Fernández-Duque, Intuitionistic linear temporal logics, ACM Transactions on Computational Logic 21 (2019) 1–32.
[40] D. Catta, V. Malvone, A. Murano, Reasoning about intuitionistic computation tree logic, in: Proc. 3rd AREA@ECAI 2023, EPTCS 391, 2023, pp. 42–48.
[41] A. Capone, D. Catta, A. Murano, et al., An intuitionistic version of computation tree logic, in: Proc. of 22nd EUMAS 2025, LNCS 16259, Springer, 2026.
[42] G. Brewka, T. Eiter, M. Truszczynski, Answer set programming at a glance, Communications of the ACM 54 (2011) 92–103.
[43] P. Cabalar, G. P. Vega, Temporal equilibrium logic: A first approach, in: Proc. 11th EUROCAST, LNCS 4739, Springer, 2007, pp. 241–248.
[44] L. Bozzelli, D. Pearce, On the complexity of temporal equilibrium logic, in: Proc. 30th LICS, IEEE Computer Society, 2015, pp. 645–656.
[45] L. Bozzelli, A. Capone, D. Catta, A. Murano, An Intuitionistic Version of Alternating-Time Temporal Logic, in: Proc. 22nd KR, 2025, pp. 185–195.
[46] A. Ferrando, V. Malvone, VITAMIN: A compositional framework for model checking of multi-agent systems, in: Proc. 17th ICAART 2025 - Vol. 1, SCITEPRESS, 2025, pp. 648–655.
[47] J. Barwise, Scenes and other situations, The journal of Philosophy 78 (1981) 369–397.
[48] M. Aruta, F. Improta, V. Malvone, A. Murano, When natural strategies meet fuzziness and resource-bounded actions, in: Proc. of 40th AAAI 2026, volume 40, 2026, pp. 29305–29313.
[49] M. Aruta, A. Murano, S. Romano, ATL for dynamic gaming environments, in: European Conference on Multi-Agent Systems, Springer, 2024, pp. 129–137.
[50] C. Listone, V. Malvone, A. Murano, Toward explainable diagnosis: A neurosymbolic approach, in: Proc. 18th ICAART, Vol. 4, INSTICC, SciTePress, 2026, pp. 3190–3197.
[51] A. Ciaramella, R. Tagliaferri, W. Pedrycz, A. Di Nola, Fuzzy relational neural network, International Journal of Approximate Reasoning 41 (2006) 146–163.
[52] A. Ferrando, G. Luongo, V. Malvone, A. Murano, Theory and practice of quantitative ATL, in: Proc. of PRIMA'24, Springer, 2024, p. 231–247.