

ATL* With Truly Perfect Recall: Expressivity and Validities

Nils Bulling¹ and Wojciech Jamroga² and Matei Popovici³

Abstract. In alternating-time temporal logic ATL*, agents with perfect recall assign choices to sequences of states, i.e., to possible finite histories of the game. However, when a nested strategic modality is interpreted, the new strategy does not take into account the previous sequence of events. It is as if agents collect their observations in the nested game again from scratch, thus effectively forgetting what they observed before. Intuitively, it does not fit the assumption of agents having perfect recall of the past.

Recently we have proposed a new semantics for ATL* where the past is not forgotten in nested games [8]. In this paper we give a formal treatment and show that the standard semantics of ATL* coincides with our new semantics in case of agents with perfect information. On the other hand, both kinds of semantics differ if agents have imperfect information about the state of the game. We compare the expressivity of the logics and their sets of validities. The latter characterize general properties of the underlying class of games.

1 Introduction

The *alternating-time temporal logic* ATL* and its fragment ATL [3] are modal logics that allow for reasoning about strategic interactions in multi-agent systems (MAS). ATL* extends the framework of temporal logic with the game-theoretic notion of *strategic ability*. Hence, ATL* is able to express statements about what agents (or groups of agents) can achieve. This is useful for specification, verification and reasoning about MAS, and is especially important because of the active development of algorithms and tools for ATL model checking [13, 10]. One challenge of MAS verification is to define desired properties in the *right way*. This means choosing the right language and the “right” semantics for a given setting, i.e., one which accurately captures characteristics of the scenario.

There are many semantic variants of ATL and ATL*. They are based on different assumptions about the capabilities of agents. For instance, agents may be able to observe the full state of the system or only parts of it (perfect or imperfect information), agents may recall the entire history of the game (perfect recall) or have no memory at all (memoryless or state-based strategies, sometimes referred to as imperfect recall) [16, 12]. Also, agents’ strategies can come with or without long-term commitment [1, 4], and so on.

More recently, we proposed a variant of the perfect recall semantics called *truly perfect recall* [8]. We argued that the standard perfect recall semantics of ATL and ATL* has counterintuitive effects:

agents may forget the past despite using perfect recall strategies. More precisely, agents *forget* their past observations once they proceed to realize a sub-goal in the game. For instance, consider the formula $\langle\langle a \rangle\rangle \bigcirc \langle\langle b \rangle\rangle \diamond \text{win}$ which says that agent *a* has a strategy ensuring that, from the next state, agent *b* can eventually win the game. Assuming that agents have perfect recall, the ability of agent *b* relies on its past observations of the game. However, the interpretation of the subformula $\langle\langle b \rangle\rangle \diamond \text{win}$ is done in the *original model*. Thus, when looking for its best strategy to win the game, agent *b* must ignore (or forget) all the observations it has made, when agent *a* was executing its strategy for $\langle\langle a \rangle\rangle \bigcirc \langle\langle b \rangle\rangle \diamond \text{win}$. This is also closely related to the way quantification in first-order logic works: variables are rebound in the scope of nested quantifiers. The semantics in [8] was proposed in order to overcome the forgetting phenomenon.

In this paper, we recall the new semantics, and argue in more detail that it offers a significantly different view of agents’ abilities than the original semantics of ATL*. More precisely, we show that if agents have imperfect information then ATL* with truly perfect recall differs from ATL* with standard perfect recall in terms of expressive power as well as valid sentences. Similar to [6], we conclude that the truly perfect recall semantics corresponds to a different *class of games*, and allows for expressing different properties of those games, than the standard variants of ATL* investigated, e.g., in [3, 16, 12, 6].

As said before, the forgetting aspect of ATL* and the idea of the no-forgetting semantics was introduced in the extended abstract [8]. This paper significantly extends that work by giving a complete, formal treatment as well as a detailed comparison in terms of expressive/distinguishing power and validity sets.

The rest of the paper is structured as follows. In Section 2 we recall the syntax of ATL* and the semantic variants which correspond to the assumptions of perfect recall and perfect/imperfect information. In Section 3 we present the no-forgetting semantics for ATL* from [8]. Then, in Section 5 and 6 we investigate the expressivity and sets of validities of ATL* with truly perfect recall, respectively. In Section 6, we conclude and discuss directions for future research.

Related work. An important strand in research on ATL* emerged in quest of the “right” semantics for strategic ability for a specific setting. We only mention some works due to lack of space. ATL was combined with epistemic logic [18, 12], and several semantic variants were defined for various assumptions about agents’ memory and available information [16, 12], cf. also [6]. Moreover, many conceptual extensions have been considered, e.g., with explicit reasoning about strategies [17, 14, 9], agents with bounded resources [2, 5] and reasoning about persistent strategies and commitment [1, 4].

The authors of [15] introduce memoryfull ATL* where the history is taken into account when evaluating cooperation modalities $\langle\langle A \rangle\rangle \psi$. More precisely, $\langle\langle A \rangle\rangle \psi$ is true in a current state *s* if, on all plays en-

¹ Clausthal University of Technology, Germany, Department of Informatics, email: bulling@in.tu-clausthal.de

² Computer Science and Communication and Interdisciplinary Centre on Security and Trust, University of Luxembourg, email: wojtek.jamroga@uni.lu

³ POLITEHNICA University of Bucharest, Romania, email: matei.popovici@cs.pub.ro

forced by A from s , ψ is true when evaluated from the beginning of the game. This is fundamentally different from our work. First of all, the authors do only consider a perfect information setting. Secondly, we do not use the history of events to evaluate a formula (our formulae are purely future-directed) but only to choose an appropriate strategy. Hence, the history only affects the agents' behavior by allowing them to learn from past events and to resolve observability limitations (which is, as we will show, only the case in imperfect information settings). Similarly, in [7] the history of events is included in the semantic relation to keep track of the agents satisfaction of goals. This is useful if agents need to plan which coalitions to join in order to satisfy their remaining (sub)goals.

2 ATL*: What Agents Can Achieve

In this section, we briefly recall the main concepts of ATL* and its variants. We introduce examples which will later be used to motivate the new no-forgetting semantics.

Syntax of Alternating-Time Temporal Logic ATL* [3] can be seen as a generalization of the branching time logic CTL* where path quantifiers E, A are replaced by *cooperation modalities* $\langle\langle A \rangle\rangle$. The formula $\langle\langle A \rangle\rangle\gamma$ expresses that group A has a *collective strategy* to enforce the temporal property γ where γ can include the temporal operators \bigcirc ("next"), and \mathcal{U} ("until"). Formally, let Π be a countable set of atomic propositions, and $\mathbb{A}gt$ be a finite nonempty set of agents. The language of ATL* is given by the following grammar:

$$\begin{aligned} \varphi &::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle A \rangle\rangle\gamma \\ \gamma &::= \varphi \mid \neg\gamma \mid \gamma \wedge \gamma \mid \bigcirc\gamma \mid \gamma\mathcal{U}\gamma \end{aligned}$$

where $A \subseteq \mathbb{A}gt$ and $p \in \Pi$. We define "sometime in the future" as $\diamond\gamma \equiv \bigcirc\mathcal{U}\gamma$ and "always in the future" as $\square\gamma \equiv \neg\diamond\neg\gamma$. Formulae φ and γ are called *state* and *path formulae* of ATL*, respectively. State formulae constitute the language of ATL*. By requiring that each temporal operator is immediately preceded by a strategic modality, we obtain the sub-language ATL; for example, $\langle\langle A \rangle\rangle\diamond p$ is an ATL formula but $\langle\langle A \rangle\rangle\diamond\square p$ and $\langle\langle A \rangle\rangle(\diamond p \wedge \diamond r)$ are not.

Models: Imperfect information concurrent game structures We interpret ATL* formulae over *imperfect information concurrent game structures* (iCGS) [18, 16]. An iCGS is given by $M = \langle \mathbb{A}gt, St, \Pi, \pi, Act, d, o, \{\sim_a \mid a \in \mathbb{A}gt\} \rangle$ consisting of a nonempty finite set of all agents $\mathbb{A}gt = \{1, \dots, k\}$, a nonempty set of states St , a set of atomic propositions Π and their valuation $\pi : \Pi \rightarrow 2^{St}$, and a nonempty finite set of (atomic) actions Act . Function $d : \mathbb{A}gt \times St \rightarrow 2^{Act}$ defines nonempty sets of actions available to agents at each state; we will usually write $d_a(q)$ instead of $d(a, q)$. Function o is a (deterministic) transition function that assigns the outcome state $q' = o(q, \alpha_1, \dots, \alpha_k)$ to each state q and tuple of actions $\langle \alpha_1, \dots, \alpha_k \rangle$ such that $\alpha_i \in d_i(q)$ for $1 \leq i \leq k$. Finally, each $\sim_a \subseteq St \times St$ is an equivalence relation that represents indistinguishability of states from agent a 's perspective.⁴ We assume that agents have identical choices in indistinguishable states ($d_a(q) = d_a(q')$ whenever $q \sim_a q'$). We also assume that collective knowledge is interpreted in the sense of "everybody knows", i.e., $\sim_A = \bigcup_{a \in A} \sim_a$. We will use $[q]_A = \{q' \mid q \sim_A q'\}$ to refer to A 's epistemic image of state q . Note that the perfect information models from [3] (concurrent game structures, CGS) can be modelled by iCGS by assuming each \sim_a to be the minimal reflexive relation.

⁴ It is important to note that these relations capture *observational* indistinguishability. The observations that agents collect during a course of action is not encoded in the model but in the semantics. This also relates to the difference between the computational and behavioral structure.




Figure 1. The iCGS M_1 describing the shell game. Tuples (α_1, α_2) represent the action profiles. α_1 denotes an action of player s —the shuffler—and action α_2 of player g —the guesser. The dotted line represents g 's indistinguishability relation; reflexive loops are omitted. State q_3 is labelled with the only proposition win . For example, when the guesser plays action pick_R in state q_2 the game proceeds to state q_3 . nop indicates the "do nothing" action.

Example 1 (Shell game) Consider model M_1 from Figure 1 that depicts a simple version of the shell game. There are two players: the shuffler s and the guesser g . Initially, the shuffler places a ball in one of two shells (the left or the right). The shells are open, and the guesser can see the location of the ball. Then the shuffler turns the shells over, so that the ball becomes hidden. The guesser wins if he picks up the shell containing the ball. Obviously, this is a very simplified version of the shell game as the shuffler does not even shuffle the shells; he simply places the ball in one of them and closes them. However, this example is rich enough to point out the limitations of the ATL*-semantics.

Two remarks are in order. First, the relation \sim_a encodes a 's (in)ability to distinguish pairs of states, based on the qualities encapsulated in those states. That is, $q \sim_a q'$ iff q and q' look the same to a , independent of the history of events that led to them. If one assumes that the agent has external memory that allows it to remember the history of past events, this must be represented by an indistinguishability relation on *histories*, see the next paragraph. Secondly, an iCGS is a template for possible actions and transitions. In order to turn it into a description of an actual game, we need also to fix the initial state. A pair (M, q) consisting of an iCGS M and a state of M is called a *pointed iCGS*.

A *history* h is a finite sequence of states $q_0q_1 \dots q_n \in St^+$ which results from the execution of subsequent transitions; that is, there must be an action profile connecting q_i with q_{i+1} for each $i = 0, \dots, n-1$. Two histories $h = q_0q_1 \dots q_n$ and $h' = q'_0q'_1 \dots q'_m$ are *indistinguishable for agent a* (denoted $h \approx_a h'$) iff $n = m$ and $q_i \sim_a q'_i$ for $i = 0 \dots n$. We use a synchronous semantics and assume that agents know how much time has passed. We also extend the indistinguishability relation over histories \approx_a , to groups: $\approx_A = \bigcup_{a \in A} \approx_a$. We write $h \circ h'$ to refer to the concatenation of the histories h and h' , $\text{fst}(h)$ and $\text{last}(h)$ to refer to the first and last state from history h , respectively. $\Lambda_M^{\text{fin}}(q)$ is the set of all histories of model M which start from state q , and $\Lambda_M^{\text{fin}} = \bigcup_{q \in St} \Lambda_M^{\text{fin}}(q)$ is the set of all histories in model M .

A *path* $\lambda = q_0q_1q_2 \dots$ is an infinite sequence of states such that there is a transition between each q_i and q_{i+1} . We write $h \circ \lambda$, where $h = q'_0q'_1 \dots q'_n$ to refer to the path $q'_0q'_1 \dots q'_nq_0q_1q_2 \dots$ obtained by concatenating h and λ . We use $\Lambda_M(q)$ to refer to the set of paths in M that start in state q and define $\Lambda_M := \bigcup_{q \in St_M} \Lambda_M(q)$ to be the

set of paths in M , respectively. We use $\lambda[i]$ to denote the i th position on path λ (starting from $i = 0$) and $\lambda[i, \infty]$ to denote the subpath of λ starting from i . Whenever the model is clear from context, we shall omit the subscript.

Strategies and Their Outcomes A strategy of agent a is a conditional plan that specifies what a is going to do in each situation. It makes sense, from a conceptual and computational point of view, to distinguish between two types of strategies: an agent may base its decision on the current state or on the whole history of events that have happened. In this paper, we consider only the latter case. A *perfect information* strategy (I -strategy for short) is a function $s_a : St^+ \rightarrow Act$ such that $s_a(q_0 \dots q_n) \in d_a(q_n)$ for all $q_0 \dots q_n \in St^+$. An *imperfect information strategy* (i -strategy) must be additionally *uniform*, in the sense that $h \approx_a h'$ implies $s_a(h) = s_a(h')$. A *collective x -strategy* s_A with $x \in \{I, i\}$, is a tuple of x -strategies, one per agent in A . Moreover, $s_A|_a$ denotes agent a 's part of the collective strategy s_A and s_\emptyset is the empty profile which is the only strategy of the empty coalition.

Function $out_M(h, s_A)$ returns the set of all paths that can occur when s_A is executed after an initial history h took place. Function $plays_M^x(h, s_A)$ returns the set of relevant paths for strategy s_A executed from h on. For perfect information, $plays_M^I(h, s_A) = out_M(h, s_A)$. For imperfect information, $plays_M^i(h, s_A)$ includes also the paths that A think might occur, i.e., ones starting from histories indistinguishable for A :

$$out_M(h, s_A) = \{h \circ \lambda = q_0 q_1 q_2 \dots \mid \text{such that for each } i \geq |h| \text{ there exists } \langle \alpha_{a_1}^{i-1}, \dots, \alpha_{a_k}^{i-1} \rangle \text{ such that } \alpha_a^{i-1} \in d_a(q_{i-1}) \text{ for every } a \in \text{Agt}, \alpha_a^{i-1} = s_A|_a(q_0 q_1 \dots q_{i-1}) \text{ for every } a \in A, \text{ and } \langle \alpha_{a_1}^{i-1}, \dots, \alpha_{a_k}^{i-1} \rangle = q_i\}.$$

$$plays_M^x(h, s_A) = \begin{cases} \bigcup_{h \approx_A h'} out_M(h', s_A) & \text{for } x = i \\ out_M(h, s_A) & \text{for } x = I \end{cases}$$

Note that the above definitions of functions out and $plays$ are slightly more general than the ones from [3, 16, 6]: outcome paths are constructed given an initial *sequence of states* rather than a single state. This will prove convenient when we define the truly perfect recall semantics of ATL^* in Section 3.

Standard Perfect Recall Semantics Let M be an iCGS and $\lambda \in \Lambda_M$. The (standard perfect recall) semantics of ATL^* , parameterized with $x \in \{i, I\}$, can be defined as follows:

$$\begin{aligned} M, \lambda \models_x p & \text{ iff } \lambda[0] \in \pi(p) \quad (\text{where } p \in \Pi); \\ M, \lambda \models_x \neg\varphi & \text{ iff } M, \lambda \not\models_x \varphi; \\ M, \lambda \models_x \varphi_1 \wedge \varphi_2 & \text{ iff } M, \lambda \models_x \varphi_1 \text{ and } M, \lambda \models_x \varphi_2; \\ M, \lambda \models_x \langle\langle A \rangle\rangle\varphi & \text{ iff there is a collective } x\text{-strategy } s_A \text{ such that,} \\ & \text{for each } \lambda' \in plays_M^x(\lambda[0], s_A), M, \lambda' \models_x \varphi; \\ M, \lambda \models_x \bigcirc\varphi & \text{ iff } M, \lambda[1, \infty] \models_x \varphi; \\ M, \lambda \models_x \varphi_1 \mathcal{U} \varphi_2 & \text{ iff there is } i \in \mathbb{N}_0 \text{ such that } M, \lambda[i, \infty] \models_x \\ & \varphi_2 \text{ and for all } 0 \leq j < i, \text{ we have that } M, \lambda[j, \infty] \models_x \varphi_1. \end{aligned}$$

Also, for a state q and a state formula φ , we define $M, q \models_x \varphi$ iff $M, \lambda \models_x \varphi$ for any $\lambda \in \Lambda_M(q)$.⁵ We refer to the logic obtained by combining \models_x with the language of ATL^* , i.e. all state formulae, as ATL_x^* . A state formula φ is *valid* in ATL_x^* iff $M, q \models_x \varphi$ for all M and states q in M .

Example 2 (Shell game ctd.) Consider the iCGS M_1 from Figure 1, and assume q_2 is the initial state of the game. It is easy

⁵ We observe that $M, \lambda \models_x \varphi$ for some $\lambda \in \Lambda_M(q)$ and state formula φ iff $M, \lambda \models_x \varphi$ for all $\lambda \in \Lambda_M(q)$.

to see that $M_1, q_2 \models_I \langle\langle g \rangle\rangle \diamond \text{win}$: under perfect information, the guesser can win by choosing the left shell in q_2 . On the other hand, $M_1, q_2 \not\models_i \langle\langle g \rangle\rangle \diamond \text{win}$: under imperfect information, the guesser has no uniform strategy that succeeds from both q_2 and q_2' . Finally, if the game begins in q_0 then the guesser can win ($M_1, q_0 \models_i \langle\langle g \rangle\rangle \diamond \text{win}$) by using the i -strategy s_g : “play $pick_L$ (resp. $pick_R$) after history $q_0 q_1 q_2$ (resp. $q_0 q_1' q_2'$)”. The strategy is uniform as both histories are distinguishable for the guesser⁶.

Note that $M, q \models_i \langle\langle A \rangle\rangle \varphi$ requires A to have a single strategy that is successful in all states indistinguishable from q for any member of the coalition. Note also that standard epistemic operators can be expressed in ATL^* . Let $\mathcal{N}\varphi \equiv \varphi \mathcal{U} \varphi$ be the “now” operator. Then, $K_a \varphi \equiv \langle\langle a \rangle\rangle \mathcal{N}\varphi$ (“ a knows that φ ”), and $E_A \equiv \langle\langle A \rangle\rangle \mathcal{N}\varphi$ (“everybody in A knows that φ ”).

3 ATL^* with Truly Perfect Recall

In the standard semantics of ATL^* agents “forget” some information about the past, even if they are assumed to have perfect recall. This can be seen in the case of nested cooperation modalities such as in $\langle\langle a \rangle\rangle \diamond \langle\langle b \rangle\rangle \square p$: b has to start collecting observations from scratch when executing its strategy for the subgoal $\square p$. This leads to counterintuitive effects, as the following example shows.

Example 3 (Forgetting in perfect recall) Recall that, on one hand, $M_1, q_0 \models_i \langle\langle g \rangle\rangle \diamond \text{win}$, that is, the guesser has a uniform strategy to win the shell game starting in q_0 . On the other hand, $M_1, q_2 \models_i \neg \langle\langle g \rangle\rangle \diamond \text{win}$. As the shuffler in q_0 can easily enforce the future state to be q_2 , we get that $M_1, q_0 \models_i \langle\langle s \rangle\rangle \diamond \neg \langle\langle g \rangle\rangle \diamond \text{win}$. Thus, in (M_1, q_0) , the guesser has the ability to win no matter what the shuffler does, and at the same time the shuffler has a strategy to deprive the guesser of the ability no matter what the guesser does!

This counterintuitive behavior is our motivation for proposing a new perfect recall semantics for ATL^* which really deserves the attribute of perfect recall.

The *no-forgetting semantics* [8] is captured by the relation \models_x^{nf} , $x \in \{i, I\}$ for the language of ATL^* , again for the perfect (I) and imperfect information (i) cases. Formulae are interpreted over triples consisting of a model, a path and an index $k \in \mathbb{N}_0$ which indicates the current state on the path. Intuitively, the subhistory of the path up to k encodes the past, and the subpath starting after k , the future. The crucial part of this semantics is that the agents always remember the sequence of past events—and they can learn from these past events.

Definition 1 (No-forgetting semantics for ATL^*) Let M be an iCGS, $\lambda \in \Lambda_M$, $k \in \mathbb{N}_0$, and $x \in \{i, I\}$. Relation \models_x^{nf} is defined as follows:

$$\begin{aligned} M, \lambda, k \models_x^{nf} p & \text{ iff } \lambda[k] \in \Pi(p) \text{ for } p \in \Pi; \\ M, \lambda, k \models_x^{nf} \neg\varphi & \text{ iff } M, \lambda, k \not\models_x^{nf} \varphi; \\ M, \lambda, k \models_x^{nf} \varphi_1 \wedge \varphi_2 & \text{ iff } M, \lambda, k \models_x^{nf} \varphi_1 \text{ and } M, \lambda, k \models_x^{nf} \varphi_2; \\ M, \lambda, k \models_x^{nf} \langle\langle A \rangle\rangle\varphi & \text{ iff there exists an } x\text{-strategy } s_A \text{ such that, for} \\ & \text{all paths } \lambda' \in plays_M^x(\lambda[0, k], s_A), M, \lambda', k \models_x^{nf} \varphi; \\ M, \lambda, k \models_x^{nf} \bigcirc\varphi & \text{ iff } M, \lambda, k+1 \models_x^{nf} \varphi; \\ M, \lambda, k \models_x^{nf} \varphi_1 \mathcal{U} \varphi_2 & \text{ iff there exists } i \geq k \text{ such that } M, \lambda, i \models_x^{nf} \varphi_2 \\ & \text{and } M, \lambda, j \models_x^{nf} \varphi_1 \text{ for all } k \leq j < i. \end{aligned}$$

⁶ We note that the guesser has no memoryless strategy (i.e. a strategy that assigns actions to states only) to win, as such a strategy had to assign the same choices to q_2 and q_2' .

We use $\text{ATL}_{\text{nf},x}^*$ to refer to the logic that combines the syntax of ATL^* with the semantic relation \models_x^{nf} . Given a state formula φ and a history h , we define $M, h \models_x^{\text{nf}} \varphi$ iff $M, \lambda, k \models_x^{\text{nf}} \varphi$ for any $\lambda \in \Lambda$ such that $\lambda[0, k] = h$. A state formula φ is *valid* in $\text{ATL}_{\text{nf},x}^*$ iff $M, q \models_x^{\text{nf}} \varphi$ for all models M and states q (note that states can be seen as a special kind of histories); and *satisfiable* if such a pair (M, q) exists.

Our new semantics differs from the standard semantics of ATL^* only in that it keeps track of the history by incorporating it into λ and *plays*^x. This affects the set of paths that are relevant when evaluating a strategy: Instead of starting with the current state of the game (as in the standard semantics) we look at paths λ that describe the play from the very beginning. $\lambda[0, k-1]$ represents the sequence of past states (excluding the current one), $\lambda[k]$ is the current state, and $\lambda[k+1, \infty]$ is the future part of the play.

We illustrate the semantics by the following example.

Example 4 (Shell game ctd.) Consider the pointed iCGS (M_1, q_0) again. Whatever the shuffler does in the first two steps, g can adapt its choice (in q_2 and q_2') to win the game. In particular, the i -strategy s_g from Example 2 can be used to demonstrate that for all $\lambda \in \text{plays}^i(q_0, s_g)$ —for every strategy of s —we have $M_1, \lambda, 0 \models_i^{\text{nf}} \diamond \langle\langle g \rangle\rangle \diamond \text{win}$. As a consequence, $M_1, q_0 \models_i^{\text{nf}} \neg \langle\langle s \rangle\rangle \diamond \neg \langle\langle g \rangle\rangle \diamond \text{win}$.

4 Truly Perfect Recall: Expressivity

We now proceed to show that the seemingly small change in semantics has important consequences for the resulting logics. We prove that the forgetting and no-forgetting variants of ATL^* differ in the properties they allow to express. We will look at which properties of iCGS can be expressed in ATL_x^* and $\text{ATL}_{\text{nf},x}^*$, respectively. To do this, we briefly recall the notions of distinguishing power and expressive power (cf. e.g. [11]).

Definition 2 (Distinguishing and expressive power) Let $\mathcal{L}_1 = (\mathcal{L}_1, \models_1)$ and $\mathcal{L}_2 = (\mathcal{L}_2, \models_2)$ be two logical systems over the same class of models \mathcal{M} —the class of iCGS in our case. By $\llbracket \varphi \rrbracket_{\models} = \{(M, q) \mid M, q \models \varphi\}$, we denote the class of pointed models that satisfy φ according to \models . Likewise, $\llbracket \varphi, M \rrbracket_{\models} = \{q \mid M, q \models \varphi\}$ is the set of states (or, equivalently, pointed models) that satisfy φ in a given structure M .

\mathcal{L}_2 is at least as expressive as \mathcal{L}_1 (written: $\mathcal{L}_1 \preceq_e \mathcal{L}_2$) iff for every formula $\varphi_1 \in \mathcal{L}_1$ there exists $\varphi_2 \in \mathcal{L}_2$ such that $\llbracket \varphi_1 \rrbracket_{\models_1} = \llbracket \varphi_2 \rrbracket_{\models_2}$. \mathcal{L}_2 is at least as distinguishing as \mathcal{L}_1 (written: $\mathcal{L}_1 \preceq_d \mathcal{L}_2$) iff for every model M and formula $\varphi_1 \in \mathcal{L}_1$ there exists $\varphi_2 \in \mathcal{L}_2$ such that $\llbracket \varphi_1, M \rrbracket_{\models_1} = \llbracket \varphi_2, M \rrbracket_{\models_2}$. Finally, we say that \mathcal{L}_1 and \mathcal{L}_2 are equally expressive (resp. distinguishing) iff $\mathcal{L}_2 \preceq_x \mathcal{L}_1$ and $\mathcal{L}_1 \preceq_x \mathcal{L}_2$ where $x = e$ (resp. $x = d$).

Note that $\mathcal{L}_1 \preceq_e \mathcal{L}_2$ implies $\mathcal{L}_1 \preceq_d \mathcal{L}_2$ but the converse is not true. For example, it is known that CTL has the same distinguishing power as CTL^* , but strictly less expressive power [11].

Perfect Information Since the difference between ATL_x^* and $\text{ATL}_{\text{nf},x}^*$ lies in the “forgetting” of past observations when evaluating nested formulae, it comes as no real surprise that the two semantics coincide for perfect information. Due to the perfect knowledge agents cannot learn anything new; and thus, they can also not forget.

Proposition 1 For all iCGSs M , $\lambda \in \Lambda_M$, and ATL^* formulae φ we have that $M, \lambda, 0 \models_i^{\text{nf}} \varphi$ iff $M, \lambda \models_I \varphi$.

Proof. The proof is done by structural induction over the formula structure. The base case ($\varphi = p$) is omitted. We formulate the induction hypothesis as $M, h \circ \lambda, k \models_i^{\text{nf}} \varphi$ iff $M, \text{last}(h) \circ \lambda \models_I \varphi$, where $k = |h| - 1$ and $|h| \geq 1$, for all path formulae φ . Then, the claim follows for $|h| = 1$. We only prove the interesting case where $\varphi \equiv \langle\langle A \rangle\rangle \gamma$.

Firstly, for a given I -strategy s_A and history h with $|h| \geq 1$ we define $s_A^{\oplus h-1}$ as follows: $s_A^{\oplus h-1}(h') = s_A(h[0, |h| - 2] \circ h')$ if $h[0, |h| - 2] \circ h'$ is a valid history in the model and arbitrarily otherwise. That is, the new strategy always simulates s_A , thus assuming that history h (without the current state) took place.

(\Rightarrow) Assume $M, h \circ \lambda, k \models_i^{\text{nf}} \langle\langle A \rangle\rangle \gamma$. Thus, there is an I -strategy s_A such that for all paths $\lambda' \in \text{out}(h, s_A)$, we have that $M, \lambda', k \models_i^{\text{nf}} \gamma$. We note that λ' is of the form $h \circ \lambda''$, where $\text{last}(h) \circ \lambda'' \in \text{out}(\text{last}(h), \hat{s}_A)$ and \hat{s}_A is the strategy such that $\hat{s}_A^{\oplus h-1} = s_A$.

By applying the induction hypothesis, we obtain that $M, \text{last}(h) \circ \lambda'' \models_I \gamma$, for all paths $\text{last}(h) \circ \lambda'' \in \text{out}(\text{last}(h), \hat{s}_A)$. Hence, $M, \text{last}(h) \circ \lambda \models_I \langle\langle A \rangle\rangle \gamma$.

(\Leftarrow) Let h be a valid history of M such that $M, \text{last}(h) \circ \lambda \models_I \langle\langle A \rangle\rangle \gamma$ and $|h| \geq 1$. Thus, there exists a strategy s_A , such that for all $\lambda' \in \text{out}(\text{last}(h), s_A)$, $M, \lambda' \models_I \gamma$. Note that λ' is of the form $\text{last}(h) \circ \lambda''$, thus, by applying the induction hypothesis, we have $M, h \circ \lambda'', k \models_i^{\text{nf}} \gamma$ for all $h \circ \lambda'' \in \text{out}(h, s_A^{\oplus h-1})$ where $k = |h| - 1$. It follows that $M, h \circ \lambda, k \models_i^{\text{nf}} \langle\langle A \rangle\rangle \gamma$. ■

The result below (and also Theorem 3) is an immediate consequence. It shows that the logics ATL_I^* and $\text{ATL}_{\text{nf},I}^*$ for perfect information are essentially equivalent.

Theorem 1 ATL_I^* and $\text{ATL}_{\text{nf},I}^*$ are equally expressive and have the same distinguishing power.

Imperfect Information In what follows, we compare the expressiveness of our no-forgetting logic $\text{ATL}_{\text{nf},i}^*$ with that of its forgetting counterpart ATL_i^* . First, we show that the two semantics differ. We consider model M_1 and state q_0 from Example 3. Let $\varphi \equiv \langle\langle s \rangle\rangle \diamond \neg \langle\langle g \rangle\rangle \diamond \text{win}$. In Examples 3 and 4 we have shown that $M_1, q_0 \models_i \varphi$ but $M_1, q_0 \not\models_i^{\text{nf}} \varphi$. Thus, we have:

Proposition 2 There is an iCGS M , a state q in M , and an ATL^* -formula φ such that $M, q \models_i \varphi$ and $M, q \not\models_i^{\text{nf}} \varphi$.

Example 5 We consider the models in Figure 2. We have that $M_2, a_0 \models_i^{\text{nf}} \langle\langle 1 \rangle\rangle \circ \langle\langle 2 \rangle\rangle \circ \text{win}$ but $M_2', a_0 \not\models_i^{\text{nf}} \langle\langle 1 \rangle\rangle \circ \langle\langle 2 \rangle\rangle \circ \text{win}$. In model M_2 player 2 can learn in which state the game is after the first move (1 plays α in a_0); this is not the case in M_2' . Therefore, under the no-forgetting semantics the two models are distinguishable.

Proposition 3 There are pointed iCGSs which satisfy the same ATL_i^* -formulae, but can be distinguished in $\text{ATL}_{\text{nf},i}^*$: $\text{ATL}_{\text{nf},i}^* \not\preceq_d \text{ATL}_i^*$.

Proof. [sketch] Let M_2 and M_2' be the iCGSs shown in Figure 2 and φ be any ATL^* -formula. Clearly, we have: (1) $M_2, x_j \models_i \varphi$ iff $M_2, x_j' \models_i \varphi$ for $x \in \{a, b\}$ and $j = 1, 2$, and analogously for M_2' . Moreover, we have (2) $M_2, x_j \models_i \varphi$ iff $M_2', x_j \models_i \varphi$ for $x \in \{a, b, a', b'\}$ and $j = 1, 2$. Now, we prove (\star) $M_2, a_0 \models_i \varphi$ iff $M_2', a_0 \models_i \varphi$ by structural induction on φ . *Basis case:* The case for atomic propositions is clear. Next, we consider $\varphi = \langle\langle A \rangle\rangle \gamma$ where γ contains no strategic quantifiers. The cases $A \in \{\emptyset, \{1\}\}$ follow trivially as each strategy of A generates the same outcome set in both models.

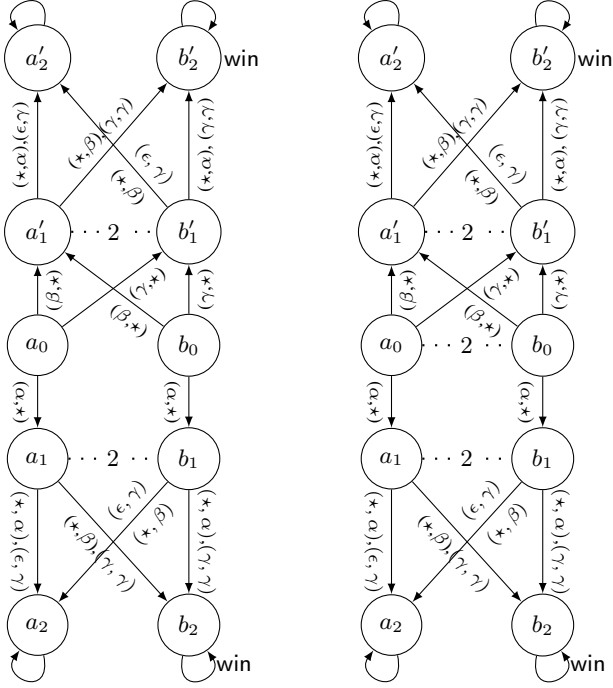


Figure 2. Models M_2 (left-hand-side) and M'_2 (right-hand side). Both models consist of two players 1 and 2. Action tuples (α_1, α_2) give the action of player 1 (α_1) and of player 2 (α_2). The only difference between both models is that in model M_2 player 2 can also not distinguish a_0 and b_0 .

Case $A = \{2\}$. The direction “ \Leftarrow ” in (\star) is clear as any (uniform) strategy in M'_2 generates the same outcome set as in M_2 . “ \Rightarrow ” Let s_1 be an arbitrary strategy in M_2 . We investigate $plays_{M_2}^i(a_0, s_1)$. First we note that $plays_{M_2}^i(a_0, s_1)$ includes either $\{a_0 a'_1 (a'_2)^\omega, a_0 b'_1 (b'_2)^\omega\}$, or $\{a_0 a'_1 (b'_2)^\omega, a_0 b'_1 (a'_2)^\omega\}$. What is essential is that the outcome set contains a path on which win holds and one where win never holds. The same is true for $plays_{M_2}^i(a_0, s'_1)$ for any strategy s'_1 in M'_2 . Now, it is easy to see that there cannot be any formula γ which distinguishes both models.

Case $A = \{1, 2\}$. The reasoning is similar to the previous case. However, we need to make sure that $\{2\}$ has a uniform strategy in $a_0 a_1$ and $b_0 b_1$ in M'_2 which ensures winning (or not winning) when cooperating with player 1. This requires the additional action γ in states a_1 and b_1 in both models. Without it we would, e.g. have that $M_2, a_0 \models_i \langle\langle 1, 2 \rangle\rangle \diamond \text{win}$ but $M'_2, a_0 \not\models_i \langle\langle 1, 2 \rangle\rangle \diamond \text{win}$.

Induction step: The cases for negation and conjunction are as usual. It remains to consider $\varphi = \langle\langle A \rangle\rangle \gamma$ where γ contains strategic quantifiers. From the previous considerations, we know that the claim follows immediately for $A \in \{\emptyset, \{1\}\}$ (as the outcome sets are equivalent) and from (2). The remaining two cases follow from (1), (2) and the specific structure of the models. Due to space limitation we skip the formal details. This concludes this part of the proof.

In Ex. 5 we have shown that both pointed models can be distinguished in $ATL_{nf,i}^*$. For every ATL_i^* -formula φ we have $a_0 \in \llbracket \varphi, M_2 \rrbracket_{=i}$ iff $a_0 \in \llbracket \varphi, M'_2 \rrbracket_{=i}$ but $a_0 \in \llbracket \varphi', M_2 \rrbracket_{\neq i}$ and $a_0 \notin \llbracket \varphi', M'_2 \rrbracket_{\neq i}$ for some φ' . Thus, we have that $ATL_{nf,i}^* \not\subseteq_d ATL_i^*$. ■

Next, we investigate whether $ATL_{nf,i}^*$ is at least as distinguishing as ATL_i^* .

Example 6 Let us consider the two iCGSs M_3 and M'_3 shown in




Figure 3. M_3 (resp. M'_3) is the iCGS shown on the left-hand-side (resp. right-hand-side) of the Figure. For an explanation, please consider Fig. 2.

Figure 3. There is an ATL_i^* -formula that can distinguish both models: $M_3, a_0 \models_i \langle\langle 1 \rangle\rangle \langle\langle 2 \rangle\rangle \bigcirc \text{win}$ and $M'_3, a_0 \not\models_i \langle\langle 1 \rangle\rangle \langle\langle 2 \rangle\rangle \bigcirc \text{win}$. In the latter case player 2 “forgets” that the game has started in state a_0 . Thus, in M_2 the player cannot distinguish the states a_1 and a_2 when evaluating the nested formula. It is easy to see that there is no uniform winning strategy from a_1 and a_2 in M'_2 , respectively.

Proposition 4 There are pointed iCGSs which satisfy the same $ATL_{nf,i}^*$ formulae, but can be distinguished in ATL_i^* : $ATL_i^* \not\subseteq_d ATL_{nf,i}^*$.

Proof. [sketch] We consider models M_3 and M'_3 from Figure 3. We only give an informal argument that there is no $ATL_{nf,i}^*$ -formula that can distinguish (M_3, a_0) from (M'_3, a_0) . Clearly, the only way to distinguish both pointed models is that some state formula is evaluated in a_1 or b_1 . The paths that start in (M_3, a_0) are isomorphic to those that start in (M'_3, a_0) . Moreover, in both models all histories that pass through a_1 are distinguishable from those that pass through b_1 , because the former start in a_0 while the latter start in b_0 . Thus, there is no way that a formula can distinguish the pointed models under the no-forgetting semantics. But both models can be distinguished in ATL_i^* as shown in Example 6. ■

Theorem 2 The logics ATL_i^* and $ATL_{nf,i}^*$ have incomparable distinguishing and expressive powers.

5 Comparing Validities

Another way of comparing two logics is to compare the sets of validities that they induce, that is, the general properties of games that can be specified and studied (cf. [6]). Intuitively, each formula can be interpreted as a property of interaction between agents in a iCGS. While expressiveness concerns the ability to capture such properties, validities are properties that universally hold. Thus, by comparing validity sets of different semantics, one is able to compare the general properties of games induced by the semantics (cf. [6]).

Perfect Information The following result is a direct corollary of Proposition 1. The result is not surprising as agents with perfect information cannot learn from past events: they have perfect information about the current state and the truth of temporal formulae does only depend on the current and future states. Thus, under perfect information both semantics yield the same logics.

Theorem 3 $Val(ATL_i^*) = Val(ATL_{nf,i}^*)$.

Imperfect Information Now we will compare the validity sets of $ATL_{nf,i}^*$ and ATL_i^* . Due to the lack of space we can only give proof sketches.

Proposition 5 $Val(ATL_i^*) \subseteq Val(ATL_{nf,i}^*)$

Proof. [sketch] We show that $Sat(ATL_{nf,i}^*) \subseteq Sat(ATL_i^*)$. Suppose $\varphi \in Sat(ATL_{nf,i}^*)$. Then, there is an iCGS M and a state q such that $M, q \models_i^{nf} \varphi$. For the moment suppose that all states can be distinguished from q . Then, the model M can be unfolded from state q to an infinite tree $T(M, q)$ the states of which correspond to histories in M . Two nodes h and h' in these trees are linked by an epistemic relation for an agent a , $h \sim_a h'$, iff $h \approx_a h'$ in M . Actually, these models correspond to the objective epistemic tree unfoldings proposed in [6]. It is easy to see that the epistemic relation in the tree already encodes no-forgetting; thus, both semantics \models_i and \models_i^{nf} coincide over them.

Now, it might be the case that there are states indistinguishable from q . They have to be considered as well. Let $Q' = \{q' \mid q \sim_{\text{Agt}} q'\}$ be the set of all states indistinguishable from q for some agent from Agt . For each state $\hat{q} \in Q'$ we construct the unfolding $T(M, \hat{q})$ as described above. Moreover, we introduce epistemic links between these trees. For any two histories h and h' in any of these trees we define $h \sim_a h'$, iff $h \approx_a h'$ in M . Let the resulting model—the collection of all these trees plus the inter-tree epistemic relations—be denoted by $\hat{T}(M, q)$. Actually, this unfolding was considered as a naive epistemic unfolding in [6, Example 7] which was shown to be insufficient for ATL_i^* . In our setting however, we can show that $M, q \models_i^{nf} \varphi$ iff $\hat{T}(M, q), q \models_i^{nf} \varphi$ iff $\hat{T}(M, q), q \models_i \varphi$ which shows that $\varphi \in Sat(ATL_i^*)$. ■

Proposition 6 $Val(ATL_{nf,i}^*) \not\subseteq Val(ATL_i^*)$

Proof. [sketch] We consider the formula $\varphi \equiv \langle\langle a \rangle\rangle \Box p \rightarrow E \circ \langle\langle a \rangle\rangle \Box p$ where $E\varphi \equiv \neg(\langle\langle \emptyset \rangle\rangle \circ \neg\varphi)$. Essentially, we can use model M'_3 from Figure 3 to show that $\varphi \notin Val(ATL_i^*)$. If we interpret $\{1, 2\}$ as a single player we have $M'_3, a_0 \models_i \langle\langle \{1, 2\} \rangle\rangle \Box \neg \text{win}$ but $M'_3, a_0 \not\models_i E \circ \langle\langle \{1, 2\} \rangle\rangle \Box \neg \text{win}$, which concludes this part.

Now, suppose that $M, q \models_i^{nf} \langle\langle a \rangle\rangle \Box p$ and let s_a be a witnessing strategy and $qq_1q_2 \dots \in \text{plays}^i(q, s_a)$. Then, we have $M, qq_1 \dots, 1 \models_i^{nf} \langle\langle a \rangle\rangle \Box p$ because $\text{plays}^i(qq_1, s_a) \subseteq \text{plays}^i(q, s_a)$; so, $M, q, 0 \models_i^{nf} E \circ \langle\langle a \rangle\rangle \Box p$ which concludes the proof. ■

With these propositions it is immediate that $ATL_{nf,i}^*$ describes a more specific class of games than ATL_i^* —games in which players do not forget past events:

Theorem 4 $Val(ATL_i^*) \subsetneq Val(ATL_{nf,i}^*)$

6 Conclusion

In this paper, we formally study the semantics for ATL^* which, unlike the standard semantics, assumes that agents forget none of their past observations. In particular, we investigate the relation between the standard perfect recall semantics and the new semantics of truly perfect recall (or no-forgetting). In the case of perfect information the no-forgetting semantics turns out to be equivalent to the standard one—due to the perfect knowledge agents cannot learn anything new and thus they also cannot forget. In the case of incomplete information, however, we show that the new semantics is incomparable to the standard one with respect to the expressive as well as distinguishing

power. Equally interesting is the comparison of general properties of games induced by the different semantics. Formally, we compare the sets of validities (similarly to [6]), and show that the truly perfect recall semantics captures a more specific class of games than the standard semantics of ATL_i^* does.

In our future work, we plan to study how strategy commitment is affected by our new semantics. Therefore, we plan to investigate the two strategy commitment logics presented in [4] and [1]. Also, we plan to investigate the complexity of model checking and satisfiability checking for the alternating-time logics with truly perfect recall.

REFERENCES

- [1] T. Ågotnes, V. Goranko, and W. Jamroga, ‘Alternating-time temporal logics with irrevocable strategies’, in *Proceedings of TARK XI*, ed., D. Samet, pp. 15–24, (2007).
- [2] N. Alechina, B. Logan, H.N. Nguyen, and A. Rakib, ‘Resource-bounded alternating-time temporal logic’, in *Proceedings of AAMAS*, pp. 481–488, (2010).
- [3] R. Alur, T. A. Henzinger, and O. Kupferman, ‘Alternating-time Temporal Logic’, *Journal of the ACM*, **49**, 672–713, (2002).
- [4] T. Brihaye, A. Da Costa Lopes, F. Laroussinie, and N. Markey, ‘ATL with strategy contexts and bounded memory’, in *Proceedings of LFCS*, volume 5407 of *Lecture Notes in Computer Science*, pp. 92–106. Springer, (2009).
- [5] N. Bulling and B. Farwer, ‘On the (Un-)Decidability of Model-Checking Resource-Bounded Agents’, in *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI 2010)*, eds., Helder Coelho and Michael Wooldridge, pp. 567–572, Lisbon, Portugal, (August 16-20 2010).
- [6] N. Bulling and W. Jamroga, ‘Comparing variants of strategic ability’, *Journal of Autonomous Agents and Multi-Agent Systems*, (2013). Springer OnlineFirst, DOI:10.1007/s10458-013-9231-3.
- [7] N. Bulling and J. Dix, ‘Modelling and verifying coalitions using argumentation and ATL’, *Inteligencia Artificial*, **14**(46), 45–73, (March 2010).
- [8] N. Bulling, W. Jamroga, and M. Popovici, ‘Agents with truly perfect recall in alternating-time temporal logic (extended abstract)’, in *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2014)*, pp. 1561–1562, Paris, France, (May 2014). ACM Press.
- [9] N. Bulling, W. Jamroga, and J. Dix, ‘Reasoning about temporal properties of rational play’, *Annals of Mathematics and Artificial Intelligence*, **53**(1-4), 51–114, (2009).
- [10] T. Chen, V. Forejt, M. Kwiatkowska, D. Parker, and A. Simaitis, ‘PRISM-games: A model checker for stochastic multi-player games’, in *Proceedings of TACAS*, volume 7795 of *LNCS*, pp. 185–191. Springer, (2013).
- [11] E.M. Clarke and B.-H. Schlingloff, ‘Model checking’, in *Handbook of Automated Reasoning*, eds., A. Robinson and A. Voronkov, 1635–1790, Elsevier, (2001).
- [12] W. Jamroga and W. van der Hoek, ‘Agents that know how to play’, *Fundamenta Informaticae*, **63**(2–3), 185–219, (2004).
- [13] A. Lomuscio and F. Raimondi, ‘MCMAS : A model checker for multi-agent systems’, in *Proceedings of TACAS*, volume 4314 of *Lecture Notes in Computer Science*, pp. 450–454, (2006).
- [14] F. Mogavero, A. Murano, and M.Y. Vardi, ‘Reasoning about strategies’, in *Proceedings of FSTTCS*, pp. 133–144, (2010).
- [15] F. Mogavero, A. Murano, and M.Y. Vardi, ‘Relentful strategic reasoning in alternating-time temporal logic’, in *Proceedings of the 16th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning, LPAR’10*, pp. 371–386, Berlin, Heidelberg, (2010). Springer-Verlag.
- [16] P. Y. Schobbens, ‘Alternating-time logic with imperfect recall’, *Electronic Notes in Theoretical Computer Science*, **85**(2), 82–93, (2004).
- [17] W. van der Hoek, W. Jamroga, and M. Wooldridge, ‘A logic for strategic reasoning’, in *Proceedings of AAMAS’05*, pp. 157–164, (2005).
- [18] W. van der Hoek and M. Wooldridge, ‘Cooperation, knowledge and time: Alternating-time Temporal Epistemic Logic and its applications’, *Studia Logica*, **75**(1), 125–157, (2003).