

Comparing Variants of Strategic Ability

Wojciech Jamroga

Computer Science and Communication,
University of Luxembourg
wojtekm.jamroga@uni.lu

Nils Bulling

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

Abstract

We show that different semantics of ability in **ATL** give rise to different validity sets. As a consequence, different notions of ability induce different strategic logics and different general properties of games. Moreover, the study can be seen as the first systematic step towards satisfiability-checking algorithms for **ATL** with imperfect information.

1 Introduction

Alternating-time temporal logic (ATL) [3] is a temporal logic that incorporates some basic game theoretical notions. Semantic variants of **ATL** are usually derived from different assumptions about agents' capabilities. Can the agents "see" the current state of the system, or only a part of it? Can they memorize the whole history of observations in the game? Is it enough that they have a way of enforcing the required temporal property "objectively", or should they be able to come up with the right strategy on their own? Different answers to these questions induce different semantics of strategic ability, and they clearly give rise to different analyses of a given problem domain. However, it is not entirely clear to what extent they give rise to different *logics*. One natural question that arises in this respect is whether these semantic variants generate different sets of valid (and, dually, satisfiable) sentences. In this paper, we settle the issue and show that most "classical" semantic variants of **ATL** are indeed different, and we characterize the relationship between their sets of validities.

The question is important for several reasons. First, by comparing validity sets we compare the respective logics in the traditional sense. Moreover, validities of **ATL** capture general properties of games under consideration: if, e.g., two variants of **ATL** generate the same valid sentences then the underlying notions of ability induce the same kind of games. All the variants studied here are defined over the same class of models that generalizes extensive games. The difference between games "induced" by different semantics lies in available strategies and the winning conditions for them.

Finally, the satisfiability problem for **ATL**, though far less studied than model checking, is not necessarily less important. While model checking **ATL** can be seen as the analogue of game solving, satisfiability corresponds naturally to mech-

anism design. A systematic study on the abstract level is the first step towards algorithms that solve the problem.

Ultimately, we show that what agents can achieve is more sensitive to the strategic model of an agent (and a precise notion of achievement) than it was generally realized. No less importantly, our study reveals that some natural properties – usually taken for granted when reasoning about action – may cease to be universally true if we change the strategic setting. Examples include fixpoint characterizations of temporal/strategic operators (that enable incremental synthesis and iterative execution of strategies) and the duality between necessary and obtainable outcomes in a game.

Related Work. **ATL** has been studied extensively; however, most of the research was focused on the way such logics can be used for specification and verification of multi-agent systems. Semantic variants were defined that match various interpretations of ability [7; 12; 9], and the complexity of model checking was investigated and compared for different settings and different variants of the logic [12; 13; 8]. Axiomatization and satisfiability were investigated in [5; 15; 11], and expressivity issues were raised in [10]. Surprisingly, relationships between the "classical" semantic variants (as defined in [12]) have not yet been studied, though analogous results exist for more sophisticated variations (cf. [1] for irrevocable strategies and [2] for agents with bounded memory). In particular, formal properties of **ATL** variants for imperfect information were largely left untouched.

2 Reasoning about Strategic Abilities

ATL [3] generalizes the branching time logic **CTL** by replacing path quantifiers E, A with *cooperation modalities* $\langle\langle A \rangle\rangle$. Informally, $\langle\langle A \rangle\rangle\gamma$ expresses that the group of agents A has a collective strategy to enforce temporal property γ . **ATL** formulae include temporal operators: " $\bigcirc$ " ("in the next state"), " $\Box$ " ("always from now on") and " $\mathcal{U}$ " ("until"). The additional operator " $\Diamond$ " ("now or sometime in the future") can be defined as $\Diamond\gamma \equiv \bigvee \mathcal{U}\gamma$. Formally, the language of **ATL*** is given by the grammar below, where A is a set of agents, and p is an atomic proposition:

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

The best known syntactic variant of alternating time temporal logic is "**ATL** without star" (or "vanilla" **ATL**) in which

every occurrence of a cooperation modality is uniquely coupled with a temporal operator. ATL^+ sits between ATL^* and “vanilla” ATL : it allows cooperation modalities to be followed by a *Boolean combination* of simple temporal subformulae. We will use the acronym ATL to refer to “ ATL without star” when no confusion can arise.

2.1 Basic Semantics of ATL

In [3], the semantics of alternating-time temporal logic is defined over a variant of transition systems where transitions are labeled with combinations of actions, one per agent. Formally, a *concurrent game structure* (CGS) is a tuple $M = \langle \text{Agt}, St, \Pi, \pi, Act, d, o \rangle$ which includes a nonempty finite set of all agents $\text{Agt} = \{1, \dots, k\}$, 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 $d : \text{Agt} \times St \rightarrow 2^{Act}$ defines nonempty sets of actions available to agents at each state, and o is a (deterministic) transition function that assigns the outcome state $q' = o(q, \alpha_1, \dots, \alpha_k)$ to state q and a tuple of actions $\alpha_i \in d(i, q)$ that can be executed by Agt in q .

A *path* $\lambda = q_0 q_1 q_2 \dots$ is an infinite sequence of states such that there is a transition between each q_i, q_{i+1} . We use $\lambda[i]$ to denote the i th position on path λ (starting from $i = 0$). The set of paths starting in q is denoted by $\Lambda_M(q)$. Moreover, we define $\Lambda_M^{fin}(q)$ as the set of all finite prefixes of $\Lambda_M(q)$.

In the standard version of ATL [3], strategies are represented by functions $s_a : St^+ \rightarrow Act$. A *collective strategy* for a group of agents $A = \{a_1, \dots, a_r\}$ is simply a tuple of individual strategies $s_A = \langle s_{a_1}, \dots, s_{a_r} \rangle$. The “outcome” function $out(q, s_A)$ returns the set of all paths that may occur when agents A execute strategy s_A from state q onward. Now, the semantics of ATL^* and its sublanguages can be defined by the standard clauses for Boolean and temporal operators, plus the following clause for $\langle\langle A \rangle\rangle$ (cf. [3] for details):

$M, q \models \langle\langle A \rangle\rangle \gamma$ iff there is a strategy s_A for agents A such that for each path $\lambda \in out(q, s_A)$, we have $M, \lambda \models \gamma$.

Note that the semantics does not address the issue of coordination [4]: if there exist several successful strategies for A , the agents in A will somehow choose between them.

We recall that the following fixpoint properties are valid in the original semantics of ATL [3]:

$$\begin{aligned} \langle\langle A \rangle\rangle \Box \varphi &\leftrightarrow \varphi \wedge \langle\langle A \rangle\rangle \Box \langle\langle A \rangle\rangle \Box \varphi \\ \langle\langle A \rangle\rangle \varphi_1 \mathcal{U} \varphi_2 &\leftrightarrow \varphi_2 \vee \varphi_1 \wedge \langle\langle A \rangle\rangle \Box \langle\langle A \rangle\rangle \varphi_1 \mathcal{U} \varphi_2. \end{aligned}$$

Moreover, the path quantifiers A, E of CTL can be expressed in ATL with $\langle\langle \emptyset \rangle\rangle, \langle\langle \text{Agt} \rangle\rangle$ respectively. As a consequence, the CTL duality axioms can be rewritten in ATL , and become validities in the basic semantics:

$$\begin{aligned} \neg \langle\langle \text{Agt} \rangle\rangle \Diamond \varphi &\leftrightarrow \langle\langle \emptyset \rangle\rangle \Box \neg \varphi, \\ \neg \langle\langle \emptyset \rangle\rangle \Diamond \varphi &\leftrightarrow \langle\langle \text{Agt} \rangle\rangle \Box \neg \varphi. \end{aligned}$$

2.2 Between Uncertainty and Recall

A number of semantic variations have been proposed for ATL , cf. e.g. [7; 12; 9; 1; 2]. In this paper, we study the “canonical” variants as proposed in [12]. There, a taxonomy of four strategy types was introduced and labeled as follows: I

(resp. i) stands for *perfect* (resp. *imperfect*) *information*, and R (resp. r) refers to *perfect recall* (resp. *no recall*). The semantics of ATL can be parameterized with the strategy type – yielding four different semantic variants of the logic, labeled accordingly ($\text{ATL}_{IR}, \text{ATL}_{IR}, \text{ATL}_{iR},$ and ATL_{ir}). In this paper, we extend the taxonomy with a distinction between *objective* and *subjective* abilities under imperfect information, denoted by i_o and i_s ; the distinction can be traced back to [7; 9].

Models, *imperfect information concurrent game structures* (iCGS) [14; 12], are CGS’s augmented with a family of indistinguishability relations $\sim_a \subseteq St \times St$, one per agent $a \in \text{Agt}$. The relations describe agents’ uncertainty: $q \sim_a q'$ means that agent a cannot distinguish between states q and q' of the system. Each $\sim_a$ is assumed to be an equivalence relation. It is also required that agents have the same choices in indistinguishable states: if $q \sim_a q'$ then $d(a, q) = d(a, q')$. Note that CGS’s can be seen as the subclass of iCGS’s where all $\sim_a$ are the minimal reflexive reflexive relations.

A *history* is a finite sequence of states. Two histories $h = q_0 q_1 \dots q_n$ and $h' = q'_0 q'_1 \dots q'_n$ are *indistinguishable for agent a* ($h \approx_a h'$) iff $n = n'$ and $q_i \sim_a q'_i$ for $i = 1, \dots, n$. Additionally, for any equivalence relation $\mathcal{R}$ over a set X we use $[x]_{\mathcal{R}}$ to denote the equivalence class of x . Moreover, we use the abbreviations $\sim_A := \bigcup_{a \in A} \sim_a$ and $\approx_A := \bigcup_{a \in A} \approx_a$. Note that relations $\sim_A$ and $\approx_A$ implement the “everybody knows” type of collective knowledge (i.e., q and q' are indistinguishable for group A iff there is at least one agent in A for whom q and q' look the same).

The following types of strategies are used in the respective semantic variants:

- I_r : $s_a : St \rightarrow Act$ such that $s_a(q) \in d(a, q)$ for all q ;
- IR : $s_a : St^+ \rightarrow Act$ such that $s_a(q_0 \dots q_n) \in d(a, q_n)$ for all $q_0, \dots, q_n$;
- i_r (i.e., $i_s r$ or $i_o r$): like I_r , with the additional constraint that $q \sim_a q'$ implies $s_a(q) = s_a(q')$;
- iR (i.e., $i_s R$ or $i_o R$): like IR , with the additional constraint that $h \approx_a h'$ implies $s_a(h) = s_a(h')$.

That is, strategy s_a is a conditional plan that specifies a ’s action in each state of the system (for memoryless agents) or for every possible history of the system evolution (for agents with perfect recall). Moreover, imperfect information strategies specify the same choices for indistinguishable states (resp. histories). As before, collective xy -strategies s_A are tuples of individual xy -strategies s_a , one per $a \in A$.

The set of possible outcomes of a strategy is defined as:

- $out^{xy}(q, s_A) = out(q, s_A)$ for $x \in \{I, i_o\}$ and $y \in \{r, R\}$;
- $out^{xy}(q, s_A) = \bigcup_{q \sim_a q'} out(q', s_A)$ for $x = i_s$ and $y \in \{r, R\}$.

We obtain the semantics for ATL_{xy} by changing the clause for $\langle\langle A \rangle\rangle \gamma$ from Section 2.1 in the following way:

$M, q \models_{xy} \langle\langle A \rangle\rangle \gamma$ iff there is an xy -strategy s_A such that for each $\lambda \in out^{xy}(q, s_A)$, we have $M, \lambda \models_{xy} \gamma$.

Note that the I and i_o semantics of ATL look only at outcome paths starting from the current global state of the system. In other words, they formalize the properties which

agents can enforce *objectively* (but, in case of uncertainty about the current state, they may be unaware of the fact). In contrast, the i_s semantics of $\langle\langle A \rangle\rangle\gamma$ refers to all outcome paths starting from states that look the same as the current state for coalition A . Hence, it formalizes the notion of A *knowing how to play* in the sense that A can identify a single strategy that succeeds from all the states they consider possible. We follow [12] by taking the “everybody knows” interpretation of collective uncertainty. More general settings were proposed in [9]; we believe that the results in this paper carry over to the other cases of “knowing how to play”, too.

We observe that the basic semantics of ATL from [3] corresponds exactly to ATL_{IR} . Moreover, in “vanilla” ATL both semantics for perfect information coincide:

Proposition 1 ([3; 12]) *For every iCGS M , state q , and ATL formula φ , we have that $M, q \models_{\text{IR}} \varphi$ iff $M, q \models_{\text{IR}} \varphi$.*

3 Comparing Validities for Variants of ATL

In this section we present a formal comparison of the semantic variants defined in Section 2. As stated in the introduction, we compare the variants on the level of their validity sets (or, equivalently, satisfiable sentences). In most cases, they turn out to be different. Also, we can usually show that one variant is a refinement of the other in the sense that its set of validities strictly subsumes the validities induced by the other variant.

In what follows, we write $\text{Val}(\text{ATL}_{\text{sem}})$ to denote the set of ATL validities under semantics sem . Likewise, we write $\text{Sat}(\text{ATL}_{\text{sem}})$ for the set of ATL formulae satisfiable in sem . The conceptual reading of $\text{Val}(\text{ATL}_{\text{sem1}}) \subsetneq \text{Val}(\text{ATL}_{\text{sem2}})$ can be as follows: for “game boards” given by iCGS’s, we have that the “game rules” in semantics ATL_{sem1} strictly refine the rules in ATL_{sem2} .

3.1 Perfect vs. Imperfect Information

We begin by comparing perfect and imperfect information scenarios. That is, in the first class (I), agents recognize the current global state of the system by definition. In the latter (i), uncertainty of agents about states constraints their choices.

Comparing ATL_{IR} vs. ATL_{IR}

First, we observe that perfect information can be seen as a special case of imperfect information.

Proposition 2 $\text{Val}(\text{ATL}_{i_s r}) \subseteq \text{Val}(\text{ATL}_{\text{IR}})$ and $\text{Val}(\text{ATL}_{i_o r}) \subseteq \text{Val}(\text{ATL}_{\text{IR}})$.

Proof. Since perfect information of agents can be explicitly represented in iCGS by fixing all relations $\sim_a$ as the minimal reflexive relations ($q \sim_a q'$ iff $q = q'$), we have that $\varphi \in \text{Sat}(\text{ATL}_{\text{IR}})$ implies $\varphi \in \text{Sat}(\text{ATL}_{i_s r})$ and $\varphi \in \text{Sat}(\text{ATL}_{i_o r})$. Thus, dually, $\text{Val}(\text{ATL}_{i_s r}) \subseteq \text{Val}(\text{ATL}_{\text{IR}})$ and $\text{Val}(\text{ATL}_{i_o r}) \subseteq \text{Val}(\text{ATL}_{\text{IR}})$. ■

Proposition 3 $\text{Val}(\text{ATL}_{\text{IR}}) \not\subseteq \text{Val}(\text{ATL}_{i_s r})$.

Proof. We show that by presenting a validity for ATL_{IR} which is not valid in $\text{ATL}_{i_s r}$. Consider the formula that captures




Figure 1: “Poor duck model” M_1 with one player (a) and transitions labeled with a ’s actions. Dotted lines depict the indistinguishability relations.

the right-to-left direction in the fixpoint characterization of $\langle\langle a \rangle\rangle\Diamond$:

$$\Phi_1 \equiv (p \vee \langle\langle a \rangle\rangle\Diamond \langle\langle a \rangle\rangle\Diamond p) \rightarrow \langle\langle a \rangle\rangle\Diamond p$$

Φ_1 is IR-valid (cf. Section 2). To see its invalidity in the $i_s r$ semantics, consider model M_1 from Figure 1.¹ Indeed, for $p \equiv \text{shot}$, we get $M_1, q_0 \models_{i_s r} p \vee \langle\langle a \rangle\rangle\Diamond \langle\langle a \rangle\rangle\Diamond p$ and $M_1, q_0 \not\models_{i_s r} \langle\langle a \rangle\rangle\Diamond p$, which formally concludes our proof. ■

Proposition 4 $\text{Val}(\text{ATL}_{\text{IR}}) \not\subseteq \text{Val}(\text{ATL}_{i_o r})$.

Proof. It is sufficient to show that $\Phi_1 \equiv (p \vee \langle\langle a \rangle\rangle\Diamond \langle\langle a \rangle\rangle\Diamond p) \rightarrow \langle\langle a \rangle\rangle\Diamond p$ is invalid in the $i_o r$ semantics. Take model M_2 in Figure 2 and $p \equiv \text{shot}$. Now we have that $M_2, q'_0 \models_{i_o r} p \vee \langle\langle a \rangle\rangle\Diamond \langle\langle a \rangle\rangle\Diamond p$ because a has a uniform strategy to achieve $\Diamond p$ in q_0 ($s_a(q) = \text{shoot}_L$ for every q) and another uniform strategy in q_1 ($s'_a(q) = \text{shoot}_R$ for every q). However, s_a and s'_a cannot be merged into a single uniform strategy, and indeed $M_2, q'_0 \not\models_{i_o r} \langle\langle a \rangle\rangle\Diamond p$. ■

Corollary 1 $\text{Val}(\text{ATL}_{i_s r}) \subsetneq \text{Val}(\text{ATL}_{\text{IR}})$ and $\text{Val}(\text{ATL}_{i_o r}) \subsetneq \text{Val}(\text{ATL}_{\text{IR}})$.

Comparing ATL_{IR} vs. ATL_{IR}

First, we observe that for $\text{ATL}_{i_o r}$ vs. ATL_{IR} we can employ the same reasoning as for $\text{ATL}_{i_o r}$ vs. ATL_{IR} . Abilities under perfect information can be still seen as a special case of imperfect information abilities, and we can use the same model M_2 to invalidate the same formula Φ_1 in $\text{ATL}_{i_o r}$. Thus, analogously to Corollary 1 we get:

¹The story behind Figure 1 is as follows. A man wants to shoot down a yellow rubber duck in a shooting gallery. The man knows that the duck is in one of the two cells in front of him, but he does not know in which one. Moreover, this has been a long party, and he is very tired, so he is only capable of using memoryless strategies at the moment. Does he have a memoryless strategy which he knows will achieve the goal? No. He can either decide to shoot to the left, or to the right, or reach out to the cells and look what is in (note also that the cells close in the moment after being opened). In each of these cases the man risks that he will fail (at least from his subjective point of view). Does he have an opening strategy that he knows will guarantee his knowing how to shoot the duck in the next moment? Yes. The opening strategy is to look; if the system proceeds to q_4 then the second strategy is to shoot to the left, otherwise the second strategy is to shoot to the right.




Figure 2: “Modified poor duck model” M_2 with 2 agents a, b . This time, we explicitly represent the agent (b) who puts the duck in one of the cells.

Corollary 2 $Val(ATL_{i_oR}) \subsetneq Val(ATL_{IR})$.

By the same reasoning as above, $Val(ATL_{i_sR}) \subseteq Val(ATL_{IR})$. To settle the other direction, we need to use another counterexample, though.

Proposition 5 $Val(ATL_{IR}) \not\subseteq Val(ATL_{i_sR})$.

Proof. This time we consider the other direction of the fix-point characterization for $\langle\langle a \rangle\rangle\Diamond$:

$$\Phi_2 \equiv \langle\langle a \rangle\rangle\Diamond p \rightarrow (p \vee \langle\langle a \rangle\rangle\Diamond \langle\langle a \rangle\rangle\Diamond p).$$

Φ_2 is IR-valid, but it is not valid in i_sR . Consider the “poor duck model” M_1 from Figure 1 but without the transitions corresponding to *look* leading from state q_0 to q_4 and from q_1 to q_5 . We call this model M'_1 and take $p \equiv \text{shot}$. Then, we have that $M'_1, q_4 \models_{i_sR} \langle\langle a \rangle\rangle\Diamond p$, but $M'_1, q_4 \not\models_{i_sR} p \vee \langle\langle a \rangle\rangle\Diamond \langle\langle a \rangle\rangle\Diamond p$, which concludes the proof. ■

Corollary 3 $Val(ATL_{i_sR}) \subsetneq Val(ATL_{IR})$ and $Val(ATL_{i_oR}) \subsetneq Val(ATL_{IR})$.

3.2 Perfect Recall vs. Memoryless Strategies

Now we proceed to examine the impact of perfect vs. no recall on the general strategic properties of agent systems.

Comparing ATL_{IR} vs. ATL_{IR}

We have already mentioned that, in “vanilla” ATL , the Ir- and IR-semantics coincide (Proposition 1). As a consequence, they induce the same validities: $Val(ATL_{IR}) = Val(ATL_{IR})$. Thus, regardless of the type of their recall, perfect information agents possess the same abilities with respect to winning conditions that can be specified in “vanilla” ATL . An interesting question is: does it carry over to more general classes of winning conditions, or are there (broader) languages that can discern between the two types of ability? The answer is: yes, there are. The Ir- and IR-semantics induce different validity sets for ATL^* , and in fact the distinction is already present in ATL^+ . Moreover, it turns out that perfect recall can be seen as a special case of memoryless play in the sense of their general properties.

Our proof of Proposition 6 draws inspiration from the proof of [1, Theorem 8.3]. We start with some additional notions and two useful lemmata.




Figure 3: Model M_3 : robot with multiple tasks

Definition 1 (Tree-like CGS) Let M be a CGS and q be a state in it. M is called tree-like iff there is a state q (the root) such that for every q' there is a unique finite sequence of states leading from q to q' .

Definition 2 (Tree unfolding) Let $M = (Agt, St, \Pi, \pi, Act, d, o)$ be a CGS and q be a state in it. The tree-unfolding of M starting from state q denoted $T(M, q)$ is defined as $(Agt, St', \Pi, \pi', Act, d', o')$ where $St' := \Lambda_M^{fin}(q)$, $d'(a, h) := d(a, last(h))$, $o'(h, \vec{\alpha}) := h \circ o(last(h), \vec{\alpha})$, and $\pi'(h) := \pi(last(h))$.

Lemma 1 For every tree-like CGS M , state q in M , and ATL^* formula φ , we have: $M, q \models_{Ir} \varphi$ iff $M, q \models_{IR} \varphi$.

Proof sketch. Induction over the structure of φ . The main case is $\varphi \equiv \langle\langle A \rangle\rangle\gamma$; to see that the proof goes through, observe that the subtree of M starting from q is also a tree-like CGS, and on tree-like CGS’s Ir and IR strategies coincide. ■

Lemma 2 For every CGS M , state q in M , and ATL^* formula φ , $M, q \models_{IR} \varphi$ iff $T(M, q), q \models_{IR} \varphi$.

Proof sketch. Induction over the structure of φ . The main case is again $\varphi \equiv \langle\langle A \rangle\rangle\gamma$ for which it is sufficient to observe that (i) IR-strategies in M, q uniquely correspond to IR-strategies in $T(M, q), q$; (ii) $out_M(q, s_A) = out_{T(M, q)}(q, s_A)$ for every IR strategy s_A . ■

Proposition 6 $Val(ATL_{IR}^*) \subseteq Val(ATL_{IR}^*)$

Proof. Let an ATL^* formula φ be Ir-valid in iCGS’s, then it is also Ir-valid in tree-like CGS’s, and by Lemma 1 also IR-valid in tree-like CGS’s. Thus, by Lemma 2, it is IR-valid in arbitrary CGS’s. Since indistinguishability relations do not influence $\models_{IR}$, we get that φ is IR-valid in iCGS’s. ■

In particular, the subsumption holds for formulae of ATL^+ . Moreover:

Proposition 7 $Val(ATL_{IR}^+) \not\subseteq Val(ATL_{IR}^+)$.

Proof. Consider formula

$$\Phi_3 \equiv \langle\langle a \rangle\rangle(\Diamond p_1 \wedge \Diamond p_2) \leftrightarrow \langle\langle a \rangle\rangle\Diamond(p_1 \wedge \langle\langle a \rangle\rangle\Diamond p_2 \vee p_2 \wedge \langle\langle a \rangle\rangle\Diamond p_1).$$

The formula is valid in ATL_{IR}^+ [6]. On the other hand, its right-to-left part is not valid in ATL_{IR}^+ . To see this, we take the single-agent CGS M_3 from Figure 3 where agent a (the robot) can either do the cleaning or the delivery of a package. Then, for $p_1 \equiv \text{clean}$, $p_2 \equiv \text{delivered}$, we have $M_3, q_0 \models_{Ir} \langle\langle a \rangle\rangle\Diamond(p_1 \wedge \langle\langle a \rangle\rangle\Diamond p_2 \vee p_2 \wedge \langle\langle a \rangle\rangle\Diamond p_1)$ but also $M_3, q_0 \not\models_{Ir} \langle\langle a \rangle\rangle(\Diamond p_1 \wedge \Diamond p_2)$. ■

Corollary 4 $Val(ATL_{IR}^+) \subsetneq Val(ATL_{IR}^+)$.




Figure 4: Model M_4 with $\text{Agt} = \{a\}$: dangers of marital life

Comparing ATL_{i_oR} vs. ATL_{i_oR} and ATL_{i_sR} vs. ATL_{i_sR}

Now we compare the memoryless and perfect recall semantics under uncertainty.

Proposition 8 $\text{Val}(\text{ATL}_{i_oR}) \subseteq \text{Val}(\text{ATL}_{i_oR})$.

Proof idea. The basic idea is similar to the one behind Proposition 6. First, we define i_oR -tree unfoldings of $iCGS$'s – similar to unfoldings of CGS 's but extended with indistinguishability relations over histories (assuming perfect recall of agents). Then we use structural induction to prove that the truth of ATL formulae coincides in a model and its unfolding (the tricky part is to show that, for nested temporal subformulae, agents in the unfolding do not have more abilities despite having more precise knowledge). Thus, i_oR -satisfiability in $iCGS$'s implies i_oR -satisfiability in tree unfoldings, but on tree unfoldings i_oR and i_oR strategies coincide. ■

The converse does not hold:

Proposition 9 $\text{Val}(\text{ATL}_{i_oR}) \not\subseteq \text{Val}(\text{ATL}_{i_oR})$

Proof. To show this, we take the ATL embedding of the CTL duality between combinators $E\Box$ and $A\Diamond$. In fact, only one direction of the equivalence is important here:

$$\Phi_4 \equiv \neg\langle\langle\emptyset\rangle\rangle\Diamond\neg p \rightarrow \langle\langle\text{Agt}\rangle\rangle\Box p.$$

First, we observe that: (i) $\neg\langle\langle\emptyset\rangle\rangle\Diamond\neg p$ expresses (regardless of the actual type of ability being considered) that there is a path in the system on which p always holds; (ii) in the “objective” semantics the set $\text{out}(q, s_{\text{Agt}})$ always consists of exactly one path; (iii) for every path λ starting from q , there is an i_oR -strategy s_{Agt} such that $\text{out}(q, s_{\text{Agt}}) = \{\lambda\}$. From these, it is easy to see that Φ_4 is valid in ATL_{i_oR} .

Second, we consider model M_4 in Figure 4.² Let us take $p \equiv \neg\text{angry} \wedge \neg\text{suspicious}$. Then, we have $M_4, q_0 \models_{i_oR} \neg\langle\langle\emptyset\rangle\rangle\Diamond\neg p$ but also $M_4, q_0 \not\models_{i_oR} \langle\langle\text{Agt}\rangle\rangle\Box p$, which demonstrates that Φ_4 is not valid in ATL_{i_oR} . ■

Corollary 5 $\text{Val}(\text{ATL}_{i_oR}) \subsetneq \text{Val}(\text{ATL}_{i_oR})$.

²The example depicts some simple traps that await a married man if he happens to be absent-minded. If he doesn't kiss his wife in the morning, he is likely to make her angry. However, if he kisses her more than once, she might get suspicious. It is easy to see that the absent-minded (i.e., memoryless) husband does not have a strategy to survive safely through the morning, though a “safe” path through the model does exist ($\lambda = q_0q_1q_1\dots$).

Proposition 10 $\text{Val}(\text{ATL}_{i_sR}) \subseteq \text{Val}(\text{ATL}_{i_sR})$.

Proof idea. The proof is analogous to Proposition 8, but the unfolding is now a forest that takes into account trees that are subjectively possible from the agents' point of view (also in the subtrees – for nested strategic subformulae). ■

Finally, we consider the reverse direction.

Proposition 11 $\text{Val}(\text{ATL}_{i_sR}) \not\subseteq \text{Val}(\text{ATL}_{i_sR})$.

Proof. We take

$$\Phi_5 \equiv \langle\langle a \rangle\rangle\Diamond\langle\langle a \rangle\rangle\Diamond p \rightarrow \langle\langle a \rangle\rangle\Diamond p.$$

The formula states that if a has an opening move and a follow-up strategy to achieve p eventually, then these can be integrated into a single strategy achieving p already from the initial state. It is easy to see that Φ_5 is valid in ATL_{i_sR} , and that the single strategy is just a concatenation of the two strategies that we get on the left hand side of the implication. On the other hand, for the “poor duck model” M_1 and $p \equiv \text{shot}$, we get that $M_1, q_0 \models_{i_sR} \langle\langle a \rangle\rangle\Diamond\langle\langle a \rangle\rangle\Diamond p$ but also $M_1, q_0 \not\models_{i_sR} \langle\langle a \rangle\rangle\Diamond p$, so Φ_5 is not valid in ATL_{i_sR} . ■

Corollary 6 $\text{Val}(\text{ATL}_{i_sR}) \subsetneq \text{Val}(\text{ATL}_{i_sR})$.

3.3 Between Subjective and Objective Ability

Finally, we compare validity sets for the semantic variants of ATL that differ on the outcome paths which are taken into account, i.e., whether only the paths representing the “objectively” possible courses of action are considered, or all the executions that are “subjectively” possible from the agents' perspective.

Proposition 12 Formula $\Phi_2 \equiv \langle\langle a \rangle\rangle\Diamond p \rightarrow p \vee \langle\langle a \rangle\rangle\Diamond\langle\langle a \rangle\rangle\Diamond p$ is valid in ATL_{i_oR} and ATL_{i_oR} , but invalid in ATL_{i_sR} and ATL_{i_sR} .

Proof. We first prove validity of Φ_2 in ATL_{i_oR} , which implies also validity in ATL_{i_oR} by Proposition 8. Suppose that $M, q \models_{i_oR} \langle\langle a \rangle\rangle\Diamond p$, then there must be an iR -strategy s_A that enforces $\Diamond p$ for every execution starting from q . But then, if p is not the case right at the beginning, s_A must lead to a next state from which it enforces $\Diamond p$.

For the second part, invalidity of Φ_2 in ATL_{i_sR} was already proved in Proposition 5. Thus, by Proposition 10, Φ_2 is not valid in ATL_{i_sR} , too. ■

Proposition 13 Let us define an additional operator N (“now”) as $N\varphi \equiv \varphi \mathcal{U} \varphi$. Formula

$$\Phi_6 \equiv \langle\langle a \rangle\rangle N\langle\langle c \rangle\rangle\Diamond\langle\langle a \rangle\rangle\Diamond p \rightarrow \langle\langle a, c \rangle\rangle\Diamond p$$

is valid in ATL_{i_sR} and ATL_{i_sR} , but invalid in ATL_{i_oR} and ATL_{i_oR} .

Proof. Analogously to Proposition 12, we will prove the validity of Φ_6 in ATL_{i_sR} , and its invalidity in ATL_{i_oR} .

First, let $M, q \models_{i_sR} \langle\langle a \rangle\rangle N\langle\langle c \rangle\rangle\Diamond\langle\langle a \rangle\rangle\Diamond p$. Then, for every state $q' \in [q]_{\sim_a}$, c has an action α_c that enforces $\langle\langle a \rangle\rangle\Diamond p$ from $[q']_{\sim_c}$. By collecting these actions into an iR -strategy s_c (we can do it since single actions are successful for whole




Figure 5: Summary: comparison of validity sets induced by various semantics of ATL. The arrows denote strict subsumption, i.e. $L_1 \rightarrow L_2$ means that $Val(L_1) \subsetneq Val(L_2)$. Unconnected pairs of nodes correspond to logics with incomparable validity sets. The double line indicates that the IR and Ir semantics induce the same validities in “vanilla” ATL, but different in the broader languages ATL^+ and ATL^* .

indistinguishability classes of c), we have that s_c enforces $\bigcirc \langle\langle a \rangle\rangle \bigcirc p$ from every state in $[q]_{\sim_{\{a,c\}}}$, regardless of what the other players do (in particular, regardless of what a does). But then, for every execution λ of s_c from $[q]_{\sim_{\{a,c\}}}$, a will have a choice to enforce $\bigcirc p$ from $\lambda[1]$. Again, collecting these choices together yields an ir-strategy s_a (we can fix the remaining choices arbitrarily). By taking $s_{\{a,c\}} = (s_a, s_c)$, we get a strategy for $\{a, c\}$ that enforces that p will be true in two steps, from every state in $[q]_{\sim_{\{a,c\}}}$. Hence, also $M, q \models_{i_{sr}} \langle\langle a, c \rangle\rangle \bigcirc p$.

For the invalidity, consider the “modified poor duck model” M_2 augmented with additional agent c that has no choice (i.e., at each state, it has only a single irrelevant action *wait* available). Let us denote the new iCGS by M'_2 . If we identify p with *shot*, it is easy to see that $M'_2, q'_0 \models_{i_{or}} \langle\langle c \rangle\rangle \bigcirc \langle\langle a \rangle\rangle \bigcirc p$, and hence also $M'_2, q'_0 \models_{i_{or}} \langle\langle a \rangle\rangle N \langle\langle c \rangle\rangle \bigcirc \langle\langle a \rangle\rangle \bigcirc p$. On the other hand, $M'_2, q'_0 \not\models_{i_{or}} \langle\langle a, c \rangle\rangle \bigcirc p$, which concludes the proof. ■

Corollary 7 *For every $y, z \in \{R, r\}$, the sets $Val(ATL_{i_{sy}})$ and $Val(ATL_{i_{oz}})$ are incomparable.*

4 Summary and Conclusions

In this paper, we compare validity sets for different semantic variants of alternating-time temporal logic. In other words, we compare the general properties of games induced by different notions of ability. It is clear that changing the notions of strategy and success in a game leads to a different game. The issue considered here is whether, given a *class* of games, such a change leads to a different class of games, too. And, if so, what is the precise relationship between the two classes.

A summary of the results is presented in Figure 5. The first, and most important, conclusion is that all the semantic

variants discussed here are *different* on the level of general properties they induce; before our study, it was by no means obvious. Moreover, our results show a very strong pattern: perfect information is a special case of imperfect information, perfect recall games are a special case of memoryless games, and properties of objective and subjective abilities of agents are incomparable.

Acknowledgements. Wojciech Jamroga acknowledges the support of the FNR (National Research Fund) Luxembourg under project S-GAMES – C08/IS/03. The authors would like to express their thanks to Jan Broersen for the argument that triggered this study, Valentin Goranko for discussions on (non-)validity of fixpoint properties in various semantics of ATL, and the anonymous reviewers for their remarks.

References

- [1] T. Ågnes, V. Goranko, and W. Jamroga. Alternating-time temporal logics with irrevocable strategies. In D. Samet, editor, *Proceedings of TARK XI*, pages 15–24, 2007.
- [2] T. Agotnes and D. Walther. A logic of strategic ability under bounded memory. *Journal of Logic, Language and Information*, 18(1):55–77, 2009.
- [3] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. *Journal of the ACM*, 49:672–713, 2002.
- [4] H. Ghaderi, H. Levesque, and Y. Lespérance. A logical theory of coordination and joint ability. In *Proceedings of AAI'07*, pages 421–426, 2007.
- [5] V. Goranko and G. van Drimmelen. Complete axiomatization and decidability of alternating-time temporal logic. *Theoretical Computer Science*, 353(1):93–117, 2006.
- [6] A. Harding, M. Ryan, and P. Schobbens. Approximating ATL^* in ATL. In *VMCAI '02*, pages 289–301. Springer-Verlag, 2002.
- [7] W. Jamroga. Some remarks on alternating temporal epistemic logic. In *Proceedings of FAMAS 2003*, pages 133–140, 2003.
- [8] W. Jamroga and J. Dix. Model checking abilities of agents: A closer look. *Theory of Computing Systems*, 42(3):366–410, 2008.
- [9] W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63(2–3):185–219, 2004.
- [10] F. Laroussinie, N. Markey, and G. Oreiby. On the expressiveness and complexity of ATL. *LMCS*, 4:7, 2008.
- [11] S. Schewe. ATL^* satisfiability is 2ExpTime-complete. In *Proceedings of ICALP 2008*, volume 5126 of *Lecture Notes in Computer Science*, pages 373–385. Springer-Verlag, 2008.
- [12] P. Y. Schobbens. Alternating-time logic with imperfect recall. *ENTCS*, 85(2):82–93, 2004.
- [13] W. van der Hoek, A. Lomuscio, and M. Wooldridge. On the complexity of practical ATL model checking. In *Proceedings of AAMAS'06*, pages 201–208, 2006.
- [14] 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.
- [15] D. Walther, C. Lutz, F. Wolter, and M. Wooldridge. ATL satisfiability is indeed EXPTIME-complete. *Journal of Logic and Computation*, 16(6):765–787, 2006.