Unawareness in Multi-Agent Systems with Partial Valuations

Line van den Berg, Manuel Atencia, Jérôme Euzenat
{line.van-den-berg,manuel.atencia,jerome.euzenat}@inria.fr
Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, F-38000 Grenoble France

ABSTRACT

Public signature awareness is satisfied if agents are aware of the vocabulary, propositions, used by other agents to think and talk about the world. However, assuming that agents are fully aware of each other's signatures prevents them to adapt their vocabularies to newly gained information, from the environment or learned through agent communication. Therefore this is not realistic for open multi-agent systems. We propose a novel way to model awareness with partial valuations that drops public signature awareness and can model agent signature unawareness, and we give a first view on definining the dynamics of raising and forgetting awareness on this framework.

KEYWORDS

Awareness; dynamic epistemic logic; partial valuations; multi-agent systems

1 INTRODUCTION

Agents use propositions to represent the information they have about the world. They may use different propositions and may not be aware of the propositions used by other agents, i.e. their *signature*, yet they may still need to communicate. In multi-agent modal logics and in particular Dynamic Epistemic Logic (DEL), all agents share the same signature. However, this is not desirable nor practical for open multi-agent systems because it prevents agents from acquiring new vocabulary or adapting their current signatures when learning new information from the environment or through agent communication.

This problem lies at the core of DEL: dynamic upgrades shrink or re-arrange the models so that the carried information becomes knowledge or belief in the resulting model. But this requires agents to already be aware of the possible future evolutions of their knowledge and beliefs and are not able to adapt their signatures.

We propose a novel way to model agent awareness with partial valuations that (i) allows agents to be unaware of other agents' signature and that (ii) enables knowledge representations to dynamically evolve. This enables us to drop public signature awareness and raise awareness of agents when they acquire new vocabulary.

2 RELATED WORK

Partial valuations have already been introduced for (Dynamic) Epistemic Logic [3–5, 7], but not connected to (dynamic) agent awareness. However, we are not the first to capture unawareness and awareness of agents. In [1], epistemic logic is extended with an operator $A\phi$ to denote "awareness of ϕ" and a complete dynamic logic with upgrades for increasing and decreasing agent awareness was developed in [6, 8–10]. In this approach, each proposition is evaluated at each world and only awareness is defined as a partial function. That is, all the propositions that agents may become aware of in the future are already specified in the initial setting. As a consequence, increasing agent awareness also uncovers the underlying truth values. Awareness is then used to distinguish between 'implicit' and 'explicit' knowledge [8].

In this paper, we propose a different viewpoint and consider becoming aware of a proposition and becoming aware of its truth value as two different acts. This enables models to evolve openly in their entirety.

3 UNAWARENESS

With partial semantics, lack of truth and falsity are not the same. This enables agents to be *uncertain* about a statement p, i.e. not knowing whether it is true or false (in the figure below on the left), like in the case with standard semantics, but also to be *unaware* of it, i.e. not considering it (in the figure below on the right). An agent is unaware of p if p is not evaluated at the worlds the agent considers plausible, where plausibility from a world w to a world v for agent a is defined as follows: 1) there is an arrow from w to v for a (wR_av), and 2) there is a (reflexive) arrow from v to v for a (vR_av).

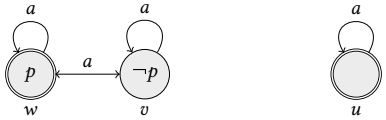

To allow agents to have different knowledge representations about the world, and to be unaware of each others signatures, there is only a 'weak reflexivity' requirement: wR_aw and wR_av implies vR_av. Reflexivity and the lack of reflexivity allow us to control what agents are aware of and therefore can have knowledge (or beliefs) about. For example, consider two agents a and b that represent the world with the propositions p and q, respectively, that they each know but that the other agent is unaware of - and therefore cannot know or believe anything about. The states of the agents are described as follows, where from w_a^* (w_b^*) agent b (a) does not have a (reflexive) arrow to w_a^* (w_b^*) but instead only to another world v_a (v_b) where p (q) is undefined:

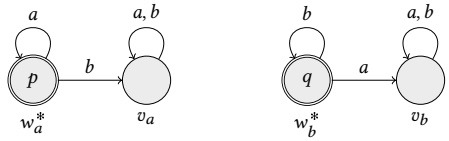

We model the knowledge and beliefs of agents from an agent-perspective, where each agent can use a different signature, or vocabulary. Thus, instead of one actual world as with standard semantics for DEL, agents have different ways to represent the actual world: these are *reflections* of the actual world, representing the actual world as the agent sees it.

We require that the reflections are consistent. More specifically, that for each agent, there is a reflection that is consistent with a

reflection of each other agent. In the example above, the reflections are w_a^* and w_b^* for agent a and b, respectively, and they are indeed consistent: p and q do not contradict each other.

This enables models to be truly open: even the reflections of the actual world are not constrained to interpret the same propositions.

3.1 Properties of awareness

We require that awareness cannot be lost over the relations R_a, but is preserved. Similar properties for awareness were already motivated in [1, 2]. In [1], awareness is assumed to only increase over time and in [2] awareness is considered constant for all the worlds the agent has access to.

In our semantics, preserving agent awareness over the relations R_a comes two-fold:

- whenever an agent a has a (reflexive) relation from w to w, she also has a (reflexive) relation from v to v for any v such that wR_av(*weak reflexivity*);
- and the propositions that are evaluated (defined) at w, remain evaluated at any v such that wR_av.

The latter property is specified as follows:

- the evaluated propositions cannot increase over R_a (*specification*);
- and any two worlds that can be reached from a world w by the same agent via R_a, share the same evaluated propositions (*consideration consistency*).

Together, the requirements of awareness enforce that agents are consistent in their considerations: if an agent a considers a proposition p or its negation plausible at a world w, she considers p or its negation plausible at every world she can reach via R_a from w.

Definition 3.1 (Properties of awareness). Let W be a set of states, a be an agent with a relation $R_a \subseteq W \times W$, and V a valuation function that assigns to each state a partial function $\cdot^{V_w} : \mathcal{P} \to \{0, 1\}$. Then the properties of awareness are formalized as:

- *Weak reflexivity:* $\forall w, v \in W$: $wR_aw \wedge wR_av \Rightarrow vR_av$
- *Specification:* $\forall w, v \in W$: $wR_av \Rightarrow Dom(\cdot^{V_v}) \subseteq Dom(\cdot^{V_w})$
- *Consideration consistency:* $\forall w, v, u \in W$: $wR_av \wedge wR_au \Rightarrow Dom(\cdot^{V_v}) = Dom(\cdot^{V_u})$

where the set of *evaluated propositions* at world w, the domain ($Dom(\cdot^{V_w})$), is defined as $Dom(\cdot^{V_w}) = \{p \in \mathcal{P} \mid p^{V_w} \in \{0, 1\}\}$.

3.2 Semantics

The semantics that we use are different from the semantics of Partial (Dynamic) Epistemic Logic in [4] in two ways:

- knowledge and belief are defined as truth in all accessible and all most plausible worlds, respectively, in which *reflexivity is satisfied*;
- and formulas ϕ are only true (or false) whenever all propositions occurring in ϕ are *defined*.

The first condition shapes our epistemic ($\sim_a$) and doxastic ($\rightarrow_a$) relations via R_a: $w \sim_a v$ iff vR_av and either wR_av or vR_aw, and $w \rightarrow_a v$ iff $v \in Max_{R_a}\{u \mid wR_au \wedge uR_au\}$. Requiring reflexivity enables us to control that agents can only know or believe a proposition if they are aware of it.

The second condition strenghtens this: it ensures that agents can only know (or believe) a formula if they have full awareness of the propositions that occur in it. For example, unlike the work in [4], this means that an agent a can only know (or believe) a disjunction, i.e. $K_a(p \vee q)$, if she is aware of both disjuncts p and q.

3.3 Raising awareness

Traditionally, dynamic upgrades for DEL reduce or re-organize the possible worlds and, with this, increase the knowledge and beliefs of agents. With a formal notion of awareness, we can additionally extend (or decrease) the valuation function to raise (or forget) agent awareness, both locally or globally. This allows agents to naturally extend their vocabularies, and hence knowledge and beliefs, with newly gained information.

Formally, to raise awareness of p ($+p$), all the worlds (globally), or all accessible worlds for an agent (locally), in which p was initially not defined are duplicated, accessibility to and from duplicated worlds being preserved, and p is made true in one world and false in the other, while preserving the relations. This means that unaware agents (p is not defined in their accessible worlds) are transformed to uncertain agents (considering p true or p false) after raising awareness.

3.4 Forgetting

A dual, inverse operator for forgetting awareness can similarly be defined. Naturally, to forget awareness of a proposition p ($-p$) all valuations of p are deleted from the model (globally), or from all accessible worlds of an agent (locally), while preserving accessibility relations. After awareness of p is raised and subsequently forgotten, i.e. $\mathcal{M}^{+p;-p}$, this way of forgetting forces us back to the original model $\mathcal{M}$, up to bisimilarity. However, after a more complex upgrade sequence like $+p; !p; !(p \rightarrow q); -p$, where a proposition ($p$) is used as evidence for another proposition (q) before it is forgotten, we have a choice: to arrive back at the original state (and therefore forgetting the truth value learned of q), or to keep the conclusions and view forgetting as a generalization operator (abstracting from the evidence p).

4 DISCUSSION AND CONCLUSION

We have provided a first view on a new semantics for modeling agent unawareness using partial valuations. This semantics allows communicating agents to be unaware of the signatures of other agents and to raise their awareness when new information is acquired.

Besides its theoretical interest, this can be used to show that public signature awareness is reached in the limit of the raising awareness upgrade. The intuition behind this is that as long as agents share all the propositions in their signature, the other agents will raise their awareness accordingly.

Future research is required to formally explore the necessary conditions for successful communication without public signature awareness and to explore the practical implications of this semantics.

ACKNOWLEDGMENTS

The authors thank the anonymous reviewers for their valuable comments and helpful suggestions. This work has been partially supported by MIAI @ Grenoble Alpes (ANR-19-P3IA-0003).

REFERENCES

[1] Ronald Fagin and Joseph Y Halpern. 1987. Belief, awareness, and limited reasoning. *Artificial intelligence* 34, 1 (1987), 39–76.

[2] Joseph Y Halpern. 2001. Alternative semantics for unawareness. *Games and Economic Behavior* 37, 2 (2001), 321–339.

[3] Jens Ulrik Hansen. 2014. Modeling truly dynamic epistemic scenarios in a partial version of DEL. *The Logica Yearbook 2013* (2014), 63–75.

[4] Jan Jaspars and Elias Thijsse. 1996. Fundamentals of partial modal logic. *Studies in Logic Language and Information* (1996).

[5] Elias Thijsse. 1994. Partial logic and knowledge representation. (1994).

[6] Johan Van Benthem and Fernando R Velázquez-Quesada. 2010. The dynamics of awareness. *Synthese* 177, 1 (2010), 5–27.

[7] Wiebe Van der Hoek, Jan Jaspars, and Elias Thijsse. 1996. Honesty in partial logic. *Studia Logica* 56, 3 (1996), 323–360.

[8] Hans Van Ditmarsch and Tim French. 2009. Awareness and forgetting of facts and agents. In *2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology*, Vol. 3. IEEE, 478–483.

[9] Hans Van Ditmarsch and Tim French. 2011. Becoming aware of propositional variables. In *Indian Conference on Logic and Its Applications*. Springer, 204–218.

[10] Hans Van Ditmarsch, Andreas Herzig, Jérôme Lang, and Pierre Marquis. 2009. Introspective forgetting. *Synthese* 169, 2 (2009), 405–423.