Beliefs, Time and Space: A Language for the Yōkai Board Game

Dominique Longin, Emiliano Lorini, Frédéric Maris

HAL Id: hal-03029712
https://hal.science/hal-03029712v2

Submitted on 2 Dec 2020

Beliefs, Time and Space:
A Language for the Yōkai Board Game

Dominique Longin[†], Emiliano Lorini[†], and Frédéric Maris[‡]

[†]CNRS, IRIT/LILaC, Toulouse University, France
[‡]Université Paul Sabatier, IRIT/ADRIA, Toulouse University

Abstract. We present an epistemic language for representing an artificial player's beliefs and actions in the context of the Yōkai board game. Yōkai is a cooperative game which requires a combination of Theory of Mind (ToM), temporal and spatial reasoning to be played effectively by an artificial agent. We show that the language properly accounts for these three dimensions and that its satisfiability problem is NP-complete. This opens up the possibility of exploiting SAT techniques for automating reasoning of an artificial player in the context of the Yōkai board-game. **The full version of this paper can be found in [15].**

1   Introduction

When one wishes to model socio-cognitive agents and, in particular, agents endowed with a Theory of Mind (ToM) who are capable of reasoning about other agents' beliefs, some of the privileged tools are epistemic logic (EL) [12,14] and its extensions by informative and communicative extensions such as public and private announcements [13,20,3]. The latter belongs to the Dynamic Epistemic Logic (DEL) family [9].

The major disadvantage of EL and DEL is that they have most of the time a high complexity thereby making them not very well-suited for practical applications. In particular, it has been shown that extending multi-agent EL by simple notions of state eliminating public announcement or arrow eliminating private announcement does not increase its PSPACE complexity (see, e.g., [19,7]). However, the satisfiability problem of full DEL with public, semi-private and private communicative actions was shown to be NEXPTIME-complete [1]. The situation is even worse in the context of epistemic planning: it is known that epistemic planning in public announcement logic (PAL) is decidable, while it becomes undecidable in full DEL, due to the fact that the epistemic model may grow as a consequence of a private announcement [6].

In [17,16], a variant of epistemic logic with a semantics exploiting belief bases is introduced. It distinguishes explicit belief from implicit belief. The former is modeled as a fact in an agent's belief base, while the latter is modeled as a fact that is deducible from the agent's explicit beliefs. The main advantages of the belief base semantics for epistemic logic compared to the standard possible world semantics based on multi-relational structures (so-called Kripke models)

are (i) its compactness, and (ii) its closeness to the way artificial cognitively-inspired agents are traditionally modeled in AI and in the area of knowledge representation and reasoning (KR) by adopting a database perspective [21]. In [18], it is shown that this variant of epistemic logic provides a valuable abstraction for modeling multi-robot scenarios in which each robot is endowed with a ToM whereby being able to ascribe epistemic states to the other robots and to reason about them.[1]

In this paper, we leverage the belief base semantics for epistemic logic to model interaction in the context of the cooperative board-game Yōkai.[2] We consider its two-player variant in which an artificial agent has to collaborate with a human agent to win it and to obtain the best score as possible. Yōkai is an interesting testbed for artificial agents, as it covers a lot of epistemic and strategic reasoning aspects as well as planning and belief revision aspects. The idea of testing the performance of artificial agents in the context of cooperative board-games in which ToM reasoning plays a role is not new. Some works exist about modeling and implementing artificial players for the card game Hanabi [4,11,10]. Yōkai adds to the ToM dimension, which is central in Hanabi, the temporal and spatial dimension. First of all, in Yōkai a player's performance relies on her/its capacity to remember the cards she/it and the other player have seen in the past. Secondly, the players must move cards in a shared space and there are spatial restrictions on card movements that should be taken into account by the players. More generally, the interesting feature of Yōkai, from the point view of KR, is the combination of epistemic, temporal and spatial reasoning that is required to completely apprehend all the game facets and dimensions.

The main novelty of our approach to modeling artificial board-game players is the use of SAT-based techniques. Specifically, the language we present for representing the artificial player's beliefs about the static and dynamic aspects of the game as well as about the human player's beliefs has the same complexity as SAT and can be polynomially translated into a propositional logic language. This opens up the possibility of exploiting SAT techniques for automating reasoning of the artificial player in the context of the Yōkai board-game.

The paper is organized as follows. In Section 2, we introduce the specification language for modeling the artificial player's actions and beliefs about the game properties and about the human player's beliefs. It is a timed language for explicit and implicit belief with a semantics exploiting belief bases. The main novelty compared to the epistemic language presented in [17] is the temporal component: the artificial player modeled in the language has knowledge about the current time of the game and beliefs about current and past beliefs of the human player.

---

[1] See also [5,8] for a DEL-based approach to modeling and implementing ToM on social robots.

[2] https://www.ultraboardgames.com/yokai/game-rules.php

2 A timed language for explicit and implicit belief

This section presents a two-agent timed variant of the language and the semantics of the logic of explicit and implicit belief presented in [17]. The two agents are the artificial agent (or machine) $\mathfrak{m}$ and the human user $\mathfrak{h}$. Agents $\mathfrak{m}$ and $\mathfrak{h}$ are treated asymmetrically. Our language allows us to represent (i) $\mathfrak{h}$'s explicit beliefs at different points in a game sequence, and (ii) $\mathfrak{m}$'s actual explicit and implicit beliefs, namely, $\mathfrak{m}$'s explicit and implicit beliefs at the current time point of the game sequence. Following [17], explicit beliefs are defined to be beliefs in an agent's belief base, while implicit beliefs are those beliefs that are derivable from the agent's explicit beliefs.

We first present the static language in which agent $\mathfrak{m}$'s beliefs do not change. Then, we present a dynamic extension in which agent $\mathfrak{m}$'s belief base can be expanded by new information.

2.1 Static language

Assume a countably infinite set of atomic propositions ATM. We define the language in two steps.

We first define the language $\mathcal{L}_0(ATM)$ by the following grammar in BNF:

$$\alpha ::= p^t \mid \triangle_{\mathfrak{h}}^t \alpha \mid now^{\geq t} \mid \neg\alpha \mid \alpha_1 \wedge \alpha_2 \mid \triangle_{\mathfrak{m}}\alpha,$$

where p ranges over ATM and t ranges over $\mathbb{N}$. $\mathcal{L}_0(ATM)$ is the language for representing agent $\mathfrak{h}$'s timed explicit beliefs and agent $\mathfrak{m}$'s actual explicit beliefs. Specifically, the formula $\triangle_{\mathfrak{h}}^t \alpha$ is read "agent $\mathfrak{h}$ explicitly believes that α at time t", whilst $\triangle_{\mathfrak{m}}$ is read "agent $\mathfrak{m}$ actually has the explicit belief that α". Atomic propositions are assumed to be timed: p^t is read "atomic proposition p is true at time t". Finally, formula $now^{\geq t}$ provides information about the current time point. It is read "the actual time of the game play is at least t".

Then, we define $\mathcal{L}_0^T(ATM)$ to be the subset $\mathcal{L}_0(ATM)$ including only timed formulas, that is:

$$\begin{aligned}
\mathcal{L}_0^T(ATM) = &\{p^t : p \in ATM \text{ and } t \in \mathbb{N}\} \cup \\
&\{\triangle_{\mathfrak{h}}^t \alpha : \alpha \in \mathcal{L}_0(ATM) \text{ and } t \in \mathbb{N}\} \cup \\
&\{now^{\geq t} : t \in \mathbb{N}\}.
\end{aligned}$$

Elements of $\mathcal{L}_0^T(ATM)$ are denoted by $x, y, \ldots$

The language $\mathcal{L}(ATM)$ extends the language $\mathcal{L}_0(ATM)$ by a modal operator of implicit belief for agent $\mathfrak{m}$ and is defined by the following grammar:

$$\varphi ::= \alpha \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \square_{\mathfrak{m}}\alpha,$$

where α ranges over $\mathcal{L}_0(ATM)$. For notational convenience we write $\mathcal{L}_0$ instead of $\mathcal{L}_0(ATM)$, $\mathcal{L}_0^T$ instead of $\mathcal{L}_0^T(ATM)$ and $\mathcal{L}$ instead of $\mathcal{L}(ATM)$, when the context is unambiguous. The formula $\square_{\mathfrak{m}}\alpha$ is read "agent $\mathfrak{m}$ actually has the implicit belief that α". The other Boolean constructions $\top, \bot, \vee, \rightarrow$ and $\leftrightarrow$ are

defined in the standard way. Notice that only formulas from the sublanguage $\mathcal{L}_0$ can be in the scope of the implicit belief operator $\Box_{\mathfrak{m}}$. Therefore, nesting of this operator is not allowed (e.g., $\Box_{\mathfrak{m}}\neg\Box_{\mathfrak{m}}p^t$ is not a well-formed formula). As we will show at the end of the section, this syntactic restriction on our language is useful to make the complexity of its satisfiability problem the same as the complexity of SAT.

The interpretation of the language $\mathcal{L}$ exploits the notion of belief base. While the notions of possible world and epistemic alternative are primitive in the standard Kripke semantics for epistemic logic [12], they are defined from the primitive concept of belief base in our semantics.

Definition 1 (State). *A state is a tuple $S = (B, V)$ where (i) $B \subseteq \mathcal{L}_0$ is agent $\mathfrak{m}$'s belief base (or, agent $\mathfrak{m}$'s subjective view of the actual situation), (ii) $V \subseteq \mathcal{L}_0^T$ is the actual situation, and such that, for every $t, t' \in \mathbb{N}$,*

$$now^{\geq 0} \in V, \tag{1}$$

$$\text{if } now^{\geq t} \in V \text{ and } t' \leq t \text{ then } now^{\geq t'} \in V, \tag{2}$$

$$now^{\geq t} \in V \text{ iff } now^{\geq t} \in B, \text{ and} \tag{3}$$

$$now^{\geq t} \notin V \text{ iff } \neg now^{\geq t} \in B. \tag{4}$$

The set of all states is denoted by $\mathbf{S}$.

Conditions (1) and (2) in the previous definition guarantees time consistency, namely, that the current time should be at least 0 and that if the current time is at least t and $t' \leq t$, then it should be at least t'. Conditions (3) and (4) capture agent $\mathfrak{m}$'s time-knowledge, namely, the assumption that $\mathfrak{m}$ has complete information about the current time. Note that the actual situation V includes timed formulas in $\mathcal{L}_0^T$ and not simply atomic propositions from ATM.

The sublanguage $\mathcal{L}_0(ATM)$ is interpreted w.r.t. states as follows:

Definition 2 (Satisfaction). *Let $S = (B, V) \in \mathbf{S}$. Then:*

$$S \models x \Longleftrightarrow x \in V,$$
$$S \models \neg\alpha \Longleftrightarrow S \not\models \alpha,$$
$$S \models \alpha_1 \wedge \alpha_2 \Longleftrightarrow S \models \alpha_1 \text{ and } S \models \alpha_2,$$
$$S \models \triangle_{\mathfrak{m}}\alpha \Longleftrightarrow \alpha \in B.$$

Observe in particular the set-theoretic interpretation of the explicit belief operator for agent $\mathfrak{m}$: agent $\mathfrak{m}$ actually has the explicit belief that α if and only if α is included in her actual belief base. This highlights the asymmetry between agent $\mathfrak{m}$ and agent $\mathfrak{h}$ in our semantics. We adopt agent $\mathfrak{m}$'s *internal* perspective, that is, the point of view of its belief base.[3] On the contrary, agent $\mathfrak{h}$'s explicit beliefs are modeled from an *external* point of view and semantically interpreted in the same way as the other timed formulas in $\mathcal{L}_0^T(ATM)$.

---

[3] See [2] for an in-depth logical analysis of the internal perspective on modeling knowledge and belief.

A multi-agent belief model (MAB) is defined to be a state supplemented with a set of states, called *context*. The latter includes all states that are compatible with agent $\mathfrak{m}$'s background knowledge.

Definition 3 (Model). *A model is a pair (S, Cxt), where $S \in \mathbf{S}$ and $Cxt \subseteq \mathbf{S}$. The class of all models is denoted by $\mathbf{M}$.*

Note that we do not impose that $S \in Cxt$. When $Cxt = \mathbf{S}$ then (S, Cxt) is said to be *complete*, since $\mathbf{S}$ is conceivable as the complete (or universal) context which contains all possible states.

Definition 4 (Epistemic alternatives). *We define $\mathcal{R}$ to be the binary relation on the set $\mathbf{S}$ such that, for all $S = (B, V), S' = (B', V') \in \mathbf{S}$:*

$$S\mathcal{R}S' \text{ if and only if } \forall \alpha \in B : S' \models \alpha.$$

$S\mathcal{R}S'$ means that S' is an epistemic alternative for the artificial agent $\mathfrak{m}$ at S. So $\mathfrak{m}$'s set of epistemic alternatives at S, noted $\mathcal{R}(S) = \{S' \in \mathbf{S} : S\mathcal{R}S'\}$, includes exactly those states that satisfy $\mathfrak{m}$'s explicit beliefs.

Definition 5 extends Definition 2 to the full language $\mathcal{L}$. Its formulas are interpreted with respect to models. We omit Boolean cases that are defined in the usual way.

Definition 5 (Satisfaction, cont.). *Let $(S, Cxt) \in \mathbf{M}$. Then:*

$$(S, Cxt) \models \alpha \Longleftrightarrow S \models \alpha;$$
$$(S, Cxt) \models \Box_{\mathfrak{m}}\varphi \Longleftrightarrow \forall S' \in Cxt, \text{if } S\mathcal{R}S' \text{ then } (S', Cxt) \models \varphi.$$

A formula $\varphi \in \mathcal{L}$ is valid in the class $\mathbf{M}$, noted $\models_{\mathbf{M}} \varphi$, if and only if $(S, Cxt) \models \varphi$ for every $(S, Cxt) \in \mathbf{M}$; it is satisfiable in $\mathbf{M}$ if and only if $\neg\varphi$ is not valid in $\mathbf{M}$. As the following theorem indicates, the satisfiability problem for $\mathcal{L}(ATM)$ has the same complexity as SAT.

Theorem 1. *Checking satisfiability of $\mathcal{L}(ATM)$ formulas in the class $\mathbf{M}$ is an NP-complete problem.*

SKETCH OF PROOF. Hardness is clear since $\mathcal{L}(ATM)$ extends the propositional logic language. As for membership, we can find a polysize satisfiability preserving translation from $\mathcal{L}(ATM)$ to propositional logic. The translation is divided in three steps. First, we transform the input formula in $\mathcal{L}(ATM)$ into negated normal form (NNF). Secondly, we translate the formula in NNF into a restricted mono-modal language with no nesting of the modal operator. Thirdly, we translate the latter into a propositional logic language in a way similar to the standard translation of modal logic into FOL. We take care of translating a finite theory including axioms corresponding to the four constraints of Definition 1. The axioms have the following form: $now^{\geq 0}$, $now^{\geq t} \rightarrow now^{\geq t'}$ for $t' \leq t$, $now^{\geq t} \leftrightarrow \triangle_{\mathfrak{m}} now^{\geq t}$ and $\neg now^{\geq t} \leftrightarrow \triangle_{\mathfrak{m}} \neg now^{\geq t}$. The theory is finite since we only need to consider instances of the axioms whose symbols occur in the input formula. For example, if $t' \leq t$ and both $now^{\geq t}$ and $now^{\geq t'}$ occur in the input formula, then $now^{\geq t} \rightarrow now^{\geq t'}$ should be included in the theory, otherwise not. ∎

2.2 Dynamic extension

Let us now move from a static to a dynamic perspective by presenting a extension of the language $\mathcal{L}(ATM)$ with belief expansion operators. Specifically, we introduce the following language $\mathcal{L}^+(ATM)$:

$$\varphi ::= \alpha \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \Box_{\mathfrak{m}}\alpha \mid [+_{\mathfrak{m}}^t\alpha]\varphi,$$

where α ranges over $\mathcal{L}_0$ and t ranges over $\mathbb{N}$. The formula $[+_{\mathfrak{m}}^t\alpha]\varphi$ is read "φ holds after agent $\mathfrak{m}$ has privately learnt that α and that the current time is at least t" or simply "φ holds after agent $\mathfrak{m}$ has privately learnt that α at time at least t".

Our extension has the following semantics relative to a model:

Definition 6 (Satisfaction relation, cont.). *Let* $S = (B, V) \in \mathbf{S}$ *and* $(S, Cxt) \in \mathbf{M}$. *Then,*

$$(S, Cxt) \models [+_{\mathfrak{m}}^t\alpha]\varphi \iff (S^{+_{\mathfrak{m}}^t\alpha}, Cxt) \models \varphi$$

with

$$S^{+_{\mathfrak{m}}^t\alpha} = (B^{+_{\mathfrak{m}}^t\alpha}, V^{+_{\mathfrak{m}}^t\alpha}),$$
$$V^{+_{\mathfrak{m}}^t\alpha} = V \cup \{now^{\geq t'} : t' \leq t\},$$
$$B^{+_{\mathfrak{m}}^t\alpha} = B \cup \{\alpha\} \cup \{now^{\geq t'} : t' \leq t\}.$$

Intuitively speaking, agent $\mathfrak{m}$'s private learning that α at time at least t simply consists in (i) adding the information α to $\mathfrak{m}$'s belief base, and (ii) moving the objective time and $\mathfrak{m}$'s subjective view of time to index t.

As the following proposition indicates, the dynamic semantics given in Definition 6 is well-defined, as it guarantees that the structure resulting from a belief expansion operation belongs to the class $\mathbf{M}$, if the initial structure also belongs to $\mathbf{M}$.

Proposition 1. *Let* $(S, Cxt) \in \mathbf{M}$. *Then,* $(S^{+_{\mathfrak{m}}^t\alpha}, Cxt) \in \mathbf{M}$.

Interestingly, adding belief expansion operators to the language $\mathcal{L}$ does not increase the complexity of the corresponding satisfiability problem.

Theorem 2. *Checking satisfiability of* $\mathcal{L}^+(ATM)$ *formulas in the class* $\mathbf{M}$ *is an NP-complete problem.*

SKETCH OF PROOF. The theorem is a consequence of Theorem 1 and the fact that we can a find a polysize reduction of the satisfiability problem for $\mathcal{L}^+(ATM)$ to the satisfiability problem for $\mathcal{L}(ATM)$. The reduction makes use of reduction axioms which allow us to eliminate dynamic operators from the input formula and to obtain a logically equivalent formula in $\mathcal{L}(ATM)$. ∎

Let $EVT = \{+_{\mathfrak{m}}^t\alpha : \alpha \in \mathcal{L}_0, t \in \mathbb{N}\}$ be the set of belief expansion events. It is reasonable to assume that such events have executability preconditions that

are specified by the following function $\mathcal{P} : EVT \longrightarrow \mathcal{L}(ATM)$. So, we can define the following operator of successful occurrence of an event in EVT:

$$\langle\!\langle +_{\mathfrak{m}}^{t} \alpha \rangle\!\rangle \varphi \overset{\text{def}}{=} \mathcal{P}(+_{\mathfrak{m}}^{t} \alpha) \wedge [+_{\mathfrak{m}}^{t} \alpha] \varphi.$$

The formula $\langle\!\langle +_{\mathfrak{m}}^{t} \alpha \rangle\!\rangle \varphi$ has to be read "agent $\mathfrak{m}$ can privately learn that α at time t and φ holds after the occurrence of this learning event".

In the next section, we will provide a formalization of the Yōkai board-game with the aid of the language $\mathcal{L}^{+}(ATM)$. We will represent agent $\mathfrak{m}$'s actions in the game as events in EVT affecting $\mathfrak{m}$'s beliefs. For every action of $\mathfrak{m}$, we will specify the corresponding executability precondition.

3   Conclusion

We have introduced a simple epistemic language for representing an artificial player's knowledge and actions in the context of the cooperative board-game Yōkai. In the extended version of this paper [15]: 1) we explain the rules of Yōkai and clarify the representation and reasoning requirements that are necessary for the artificial player to be able to play the game in a clever way; 2) we have shown that this game requires a combination of Theory of Mind (ToM), temporal and spatial reasoning to be played effectively by the artificial agent; 3) we illustrate how this epistemic language is able to represent strategic reasoning rules. Our approach relies on SAT given the existence of a polysize satisfiability preserving translation of the epistemic language into propositional logic.

Future work will be organized in two steps. First, we intend to specify a belief revision module for the artificial player which spells out how it should change its beliefs after the human player's moves. Secondly, we plan to implement an artificial player based on our formalization and to use existing SAT solvers for automating the reasoning and planning for the artificial player during the game play. We plan to formalize and implement a variety of game strategies for the artificial player, in line with the methodology sketched in [15, Section 4.3], and to the test their performances experimentally.

References

1. Aucher, G., Schwarzentruber, F.: On the complexity of dynamic epistemic logic. In: Proceedings of the 14th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2013) (2013)
2. Aucher, G.: Private announcement and belief expansion: an internal perspective. Journal of Logic and Computation **22**(3), 451–479 (2012)
3. Baltag, A., Moss, L., Solecki, S.: The logic of public announcements, common knowledge and private suspicions. In: Gilboa, I. (ed.) Proceedings of the Seventh Conference on Theoretical Aspects of Rationality and Knowledge (TARK'98). pp. 43–56. Morgan Kaufmann, San Francisco, CA (1998)

4. Bard, N., Foerster, J.N., Chandar, S., Burch, N., Lanctot, M., Song, H.F., Parisotto, E., Dumoulin, V., Moitra, S., Hughes, E., Dunning, I., Mourad, S., Larochelle, H., Bellemare, M.G., Bowling, M.: The hanabi challenge: A new frontier for AI research. Artificial Intelligence **280** (2020)

5. Bolander, T.: Seeing is believing: Formalising false-belief tasks in dynamic epistemic logic. In: Herzig, A., Lorini, E. (eds.) Proceedings of the European conference on Social Intelligence (ECSI-2014). pp. 87–107 (2014)

6. Bolander, T., Andersen, M.B.: Epistemic planning for single- and multi-agent systems. Journal of Applied Non-Classical Logics **21**(1), 656–680 (2011)

7. Bolander, T., van Ditmarsch, H., Herzig, A., Lorini, E., Pardo, P., Schwarzentruber, F.: Announcements to attentive agents. Journal of Logic, Language and Information **25**(1), 1–35 (2015)

8. Dissing, L., Bolander, T.: Implementing theory of mind on a robot using dynamic epistemic logic. In: Bessiere, C. (ed.) Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, (IJCAI 2020). pp. 1615–1621 (2020)

9. van Ditmarsch, H.P., van der Hoek, W., Kooi, B.: Dynamic Epistemic Logic. Kluwer Academic Publishers (2007)

10. Eger, M., Martens, C.: Practical specification of belief manipulation in games. In: Magerko, B., Rowe, J.P. (eds.) Proceedings of the Thirteenth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE-17). pp. 30–36. AAAI Press (2017)

11. Eger, M., Martens, C., Cordoba, M.A.: An intentional AI for hanabi. In: IEEE Conference on Computational Intelligence and Games, CIG 2017. pp. 68–75. IEEE (2017)

12. Fagin, R., Halpern, J., Moses, Y., Vardi, M.: Reasoning about Knowledge. MIT Press, Cambridge (1995)

13. Gerbrandy, J., Groeneveld, W.: Reasoning about information change. Journal of Logic, Language, and Information **6**, 147–196 (1997)

14. Halpern, J.Y., Moses, Y.: A guide to completeness and complexity for modal logics of knowledge and belief. Artificial Intelligence **54**(2), 319–379 (1992)

15. Longin, D., Lorini, E., Maris, F.: Beliefs, Time and Space: A Language for the Yōkai Board Game. Research report, IRIT - Institut de recherche en informatique de Toulouse (Oct 2020)

16. Lorini, E.: In praise of belief bases: Doing epistemic logic without possible worlds. In: Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18). pp. 1915–1922. AAAI Press (2018)

17. Lorini, E.: Rethinking epistemic logic with belief bases. Artificial Intelligence **282** (2020)

18. Lorini, E., Romero, F.: Decision procedures for epistemic logic exploiting belief bases. In: Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019). pp. 944–952. IFAAMAS (2019)

19. Lutz, C.: Complexity and succinctness of public announcement logic. In: Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006). pp. 137–143. ACM (2006)

20. Plaza, J.A.: Logics of public communications. In: Emrich, M., Pfeifer, M., Hadzikadic, M., Ras, Z. (eds.) Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems. 201-216 (1989)

21. Shoham, Y.: Logical theories of intention and the database perspective. Journal of Philosophical Logic **38**(6), 633–648 (2009)