Minimal belief and negation as failure in multi-agent systems

Riccardo Rosati

Dipartimento di Informatica e Sistemistica
Università di Roma "La Sapienza"
Via Salaria 113, 00198 Roma, Italy
E-mail: rosati@dis.uniroma1.it

We propose an epistemic, nonmonotonic approach to the formalization of knowledge in a multi-agent setting. From the technical viewpoint, a family of nonmonotonic logics, based on Lifschitz's modal logic of minimal belief and negation as failure, is proposed, which allows for formalizing an agent which is able to reason about both its own knowledge and other agents' knowledge and ignorance. We define a reasoning method for such a logic and characterize the computational complexity of the major reasoning tasks in this formalism. From the practical perspective, we argue that our logical framework is well-suited for representing situations in which an agent cooperates in a team, and each agent is able to communicate his knowledge to other agents in the team. In such a case, in many situations the agent needs nonmonotonic abilities, in order to reason about such a situation based on his own knowledge and the other agents' knowledge and ignorance. Finally, we show the effectiveness of our framework in the robotic soccer application domain.

Keywords: Multi-agent systems, computational logic, nonmonotonic reasoning, modal epistemic logics

1. Introduction

In this paper we propose an epistemic and nonmonotonic approach to the formalization of knowledge and belief in a multi-agent setting. Our aim is twofold: on the one hand, we want to define a theoretical framework which is epistemologically adequate for reasoning in concrete multi-agent scenarios; on the other hand, we want to analyze the computational properties and to define algorithms for reasoning in such a framework, to the aim of providing the basis for an implementation of such a framework in a real robotic architecture.

Our starting point is Lifschitz's modal logic of minimal belief and negation as failure MBNF [16], which allows for formalizing an agent which is able to reason about both his own beliefs and other agents' beliefs. Lifschitz's logic MBNF is a modal logic with two autoepistemic operators: a "minimal belief" modality B and a "negation as failure" (also called "negation by default") modality *not*; however, for ease of notation, in the following we will use the symbol A (introduced in [17]), which stands for $\neg$*not* and is interpreted in terms of "autoepistemic assumption". It has been proved [16] that MBNF is able to embed many of the best known formalisms for nonmonotonic reasoning, e.g. default logic, autoepistemic logic, circumscription, and extended disjunctive logic programs (under the stable model semantics). Such a logic has therefore been considered as a unifying framework for nonmonotonic reasoning.

In this paper, we define the family of logics MBNF($\mathcal{K}$), obtained as the extension of the multimodal systems for knowledge and belief $\mathsf{K}_n, \mathsf{T}_n, \mathsf{S4}_n, \mathsf{KD45}_n, \mathsf{S5}_n$ [12] with Lifschitz's modalities B and A. We define a reasoning method for such logics and characterize the computational complexity of the major reasoning tasks in these formalisms. In particular, we prove that reasoning in all MBNF($\mathcal{K}$) logics is a PSPACE-complete task, which implies that extending any of the above mentioned multimodal systems with the autoepistemic modalities B and A does not increase the worst-case complexity of reasoning. We also identify the major sources of complexity of reasoning in MBNF($\mathcal{K}$), and establish complexity results for MBNF($\mathcal{K}$) under various syntactic restrictions, which correspond to imposing different bounds on the number of agents modeled in the system and the depth of nesting of the modal operators.

We then argue that our logical framework is well-suited for representing situations in which an agent cooperates in a team, and each agent is able to communicate his beliefs about the world to other agents. In such a case, in many situations the agent needs nonmonotonic abilities for reasoning about such a situation, based on the other agents' knowledge and ignorance. In particular, we exploit the capabilities of MBNF($\mathcal{K}$) in the formalization of nonmonotonic reasoning about many agents, in order to represent nonmonotonic rules for inferring new knowledge about the world, based on the information provided by other agents. We show the usefulness of such kind of rules in multi-agent applications requiring a selective, qualitative fusion of information coming from different agents. To this aim, we have tested the epistemological adequacy of our framework in the robotic soccer application domain (RoboCup).

The paper is structured in the following way. In Section 2, we illustrate syntax and semantics of the MBNF($\mathcal{K}$) framework. In Section 3 we present a general reasoning method for MBNF($\mathcal{K}$), then, in Section 4, we discuss complexity issues and identify the sources of complexity of reasoning in MBNF($\mathcal{K}$). In Section 5, we present the application of our framework to the RoboCup domain. We illustrate related work in Section 6, and conclude in Section 7.

2. The logic MBNF($\mathcal{K}$)

In this section we define the family of logics MBNF($\mathcal{K}$). Informally, such logics can be both seen as the extension of a multimodal logic with the MBNF modalities B and A and as a syntactic restriction of first-order MBNF (since each multimodal logic can be seen as a fragment of first-order logic).

We assume that the reader is familiar with the basics of modal logic. In the following, we denote with $\mathcal{K}$ a multimodal logic among the following formalisms: $\mathsf{K}_n, \mathsf{T}_n, \mathsf{S4}_n, \mathsf{KD45}_n, \mathsf{S5}_n$ [12]. We recall that, e.g., K_n denotes the multimodal extension of normal modal logic K: n different modalities $K_1, \ldots, K_n$ can occur in a K_n-formula. We also recall that T denotes the modal logic interpreted on Kripke structures whose accessibility relation among worlds is reflexive; $\mathsf{S4}$ imposes reflexivity and transitivity on such a relation; $\mathsf{KD45}$ denotes the modal logic interpreted on Kripke structures whose accessibility relation among worlds is serial, transitive and euclidean, while modal logic $\mathsf{S5}$ imposes symmetry, transitivity, and reflexivity on the accessibility relation.

The abstract syntax of MBNF($\mathcal{K}$) is as follows:

$$\psi = p \mid \neg\psi \mid \psi_1 \wedge \psi_2 \mid K_i\psi$$
$$\varphi = \psi \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid B\varphi \mid A\varphi$$

where p is an element from an alphabet of propositional symbols $\mathcal{A}$ and $1 \leq i \leq n$. We also use the logical connectives $\vee, \supset$, which are defined as usual in terms of $\wedge, \neg$. Therefore, ψ denotes a $\mathcal{K}$-formula, i.e. a modal formula of the language $\mathcal{L}$ of the logic $\mathcal{K}$, while φ denotes an MBNF($\mathcal{K}$) formula, i.e. a formula in the language $\mathcal{L}_M$ of the logic MBNF($\mathcal{K}$). Notice that we do not allow occurrences of MBNF operators within the scope of $\mathcal{K}$ modalities.

Moreover, we denote with $\mathcal{L}_M^S$ the set of *subjective* MBNF($\mathcal{K}$) formulas, that is the set of formulas from $\mathcal{L}_M$ of the form φ such that each $\mathcal{K}$-subformula occurring in φ lies within the scope of an MBNF modality (i.e., B or A).

We now define the semantics of MBNF($\mathcal{K}$) formulas.

Let W be a fixed, countably infinite set of elements called worlds, and let $\mathcal{U}_\mathcal{A}$ be the set of propositional valuations over an alphabet $\mathcal{A}$. We call $\mathcal{K}$-*interpretation* a usual interpretation structure for the logic $\mathcal{K}$, i.e. a Kripke structure I of the form $I = \langle w, W, R_1, \ldots, R_n, V\rangle$, where $w \in W$ is called *initial world* of I, $V : W \rightarrow \mathcal{U}_\mathcal{A}$ and $R_i \subseteq W \times W$ for each $i \in \{1, \ldots, n\}$ such that:

- if $\mathcal{K} = \mathsf{T}_n$, then each R_i is a reflexive relation;
- if $\mathcal{K} = \mathsf{S4}_n$, then each R_i is a reflexive and transitive relation;
- if $\mathcal{K} = \mathsf{KD45}_n$, then each R_i is a serial, transitive and euclidean relation;
- if $\mathcal{K} = \mathsf{S5}_n$, then each R_i is a reflexive, symmetric and transitive relation.

$\mathcal{K}$-satisfiability of a formula $\psi \in \mathcal{L}$ in I (which we denote as $I \models \psi$) is defined in the usual way:

1. if ψ is an atom, then $I \models \psi$ iff $V(w)(\psi) = \mathit{TRUE}$;
2. $I \models \neg\psi$ iff $I \not\models \psi$;
3. $I \models \psi_1 \wedge \psi_2$ iff $I \models \psi_1$ and $I \models \psi_2$;
4. $I \models K_i\psi$ iff, for every $w' \in W$ s. t. $(w, w') \in R_i$, $\langle w', W, R_1, \ldots, R_n, V\rangle \models \psi$.

In the following, we call *cluster* a set of $\mathcal{K}$-interpretations M of the above form. An MBNF($\mathcal{K}$) structure is a triple (I, M_b, M_a), where I is a $\mathcal{K}$-interpretation and M_b, M_a are clusters, which are denoted respectively as the B-cluster and the A-cluster of (I, M_b, M_a).

Definition 1. Satisfiability of a formula in an MBNF($\mathcal{K}$) structure is inductively defined as follows:

1. if $\varphi \in \mathcal{L}$, $(I, M_b, M_a) \models \varphi$ iff $I \models \varphi$;
2. $(I, M_b, M_a) \models \neg\varphi$ iff $(I, M_b, M_a) \not\models \varphi$;
3. $(I, M_b, M_a) \models \varphi_1 \wedge \varphi_2$ iff $(I, M_b, M_a) \models \varphi_1$ and $(I, M_b, M_a) \models \varphi_2$;
4. $(I, M_b, M_a) \models B\varphi$ iff, for every $J \in M_b$, $(J, M_b, M_a) \models \varphi$;
5. $(I, M_b, M_a) \models A\varphi$ iff, for every $J \in M_a$, $(J, M_b, M_a) \models \varphi$.

Notice that the above definition evaluates the modality B in M_b (and the modality A in M_a) as in the Kripke model in which each world corresponds to a $\mathcal{K}$-interpretation in M_b and the accessibility relation between worlds is universal.

The nonmonotonic character of MBNF($\mathcal{K}$) is obtained by imposing the following preference semantics over the interpretation structures satisfying a given formula.

Definition 2. An MBNF($\mathcal{K}$) structure of the form (I, M, M) is an MBNF($\mathcal{K}$) *model* (or simply *model*) for $\sigma \in \mathcal{L}_M$ iff $(I, M, M) \models \sigma$ and, for every $\mathcal{K}$-interpretation J and every cluster M', if $(J, M', M) \models \sigma$ then $M' \not\supset M$.

We say that $\varphi \in \mathcal{L}_M$ is MBNF($\mathcal{K}$)-satisfiable if φ has a model (MBNF($\mathcal{K}$)-unsatisfiable otherwise). We say that a formula $\varphi \in \mathcal{L}_M$ is entailed by $\sigma \in \mathcal{L}_M$ (and write $\sigma \models_{\mathrm{MBNF}(\mathcal{K})} \varphi$) iff, for every model (I, M, M) for σ, $(I, M, M) \models \varphi$.

Informally, the above preference semantics provides the modality B with a "minimal belief" meaning, while the modality A is interpreted in terms of an autoepistemic assumption that has to be justified [17]. More precisely, the above semantics realizes a "belief closure", that can be understood in terms of maximization of the formulas not believed by the agent.

Example 3. Let $\sigma = B\varphi$, where φ is a $\mathcal{K}$-formula. The only MBNF($\mathcal{K}$) models for σ are of the form (I, M, M), with $M = \{J : J \models \varphi\}$. Hence, $\sigma \models_{\mathrm{MBNF}(\mathcal{K})} B\varphi$, and $\sigma \models_{\mathrm{MBNF}(\mathcal{K})} \neg B\psi$ for each $\psi \in \mathcal{L}$ such that the formula $\varphi \supset \psi$ is not valid in logic $\mathcal{K}$ (i.e., it is satisfied by all $\mathcal{K}$-interpretations). Therefore, the agent modeled by σ has minimal belief, in the sense that she only believes φ and the formulas entailed by φ in the logic $\mathcal{K}$, while she does not believe all other $\mathcal{K}$-formulas. $\square$

As explained by the above example, the meaning of the operator B is provided by the belief closure implied by Definition 2. Such a closure allows the agent for deriving in a nonmonotonic way what she does not believe.

Moreover, the autoepistemic assumption operator A allows for expressing so-called autopistemic or "justified" assumptions [17]. Actually, the modality A corresponds exactly to Moore's autoepistemic operator L [17,21]. As shown by Lifschitz, such a modality (although in its negated form) is also equivalent to the negation-as-failure operator of logic programming under stable model semantics, and has a deep semantic correspondence to justifications of Reiter's default logic [16]. Therefore, the combined usage of B and A allows for easily formalizing many of the best known nonmonotonic reasoning mechanisms. As an example, we now recall the representation of Reiter's default rules in terms of MBNF($\mathcal{K}$) formulas (and refer the reader to [16] for details on embedding other nonmonotonic logics into MBNF). Let d be a propositional default rule of the form

$$\frac{\alpha : \beta}{\gamma}$$

Then, we denote with $\tau(d)$ the MBNF($\mathcal{K}$) formula $\neg B\alpha \vee A\neg\beta \vee B\gamma$. Given a finite default theory (D, W), we denote with $\tau(D, W)$ the MBNF($\mathcal{K}$) formula

$$\tau(D, W) = (\bigwedge_{\psi \in W} \psi) \wedge (\bigwedge_{d \in D} \tau(d))$$

It can be shown that, given a default theory (D, W) and a propositional formula ψ, ψ belongs to each Reiter's default extension of (D, W) if and only if $\tau(D, W) \models_{\mathrm{MBNF}(\mathcal{K})} B\psi$ (see [16]).

We point out that MBNF($\mathcal{K}$) corresponds to the modal propositional fragment of MBNF if we restrict MBNF($\mathcal{K}$) to a single agent, i.e., when the only modalities allowed in MBNF($\mathcal{K}$) formulas are B and A.

Finally, from the knowledge representation viewpoint, the MBNF($\mathcal{K}$) framework can be used for modeling an agent who is able to reason about her own beliefs and other agents' knowledge/beliefs. Precisely:

- the agent's own beliefs can be formalized by the modalities B and A;
- the knowledge/belief of other agents can be expressed by means of the modalities $K_1, \ldots, K_n$: each modality K_i is used to express the knowledge/beliefs of agent i. The way such agents are modeled thus depends on the choice of the modal logic $\mathcal{K}$;
- the agent may have beliefs about other agents' knowledge/beliefs (i.e., $\mathcal{K}$-formulas may occur within the scope of B and A).

Consequently, MBNF($\mathcal{K}$) allows for representing an agent who is able to perform autoepistemic nonmonotonic reasoning about his own beliefs and other agents' knowledge/beliefs. As shown in Section 5, such representational features appear well-suited to actual multi-agent applications.

3. Reasoning in MBNF($\mathcal{K}$)

In this section we study reasoning in MBNF($\mathcal{K}$). We start by introducing some auxiliary definitions and properties.

In the following, we say that an occurrence of a subformula ψ in a formula $\varphi \in \mathcal{L}_M$ is *strict* if it does not lie within the scope of an MBNF modal operator (i.e., B or A). E.g., let $\varphi = B\sigma \wedge A(B\psi \vee \xi)$. The occurrence of $B\sigma$ in φ is strict, while the occurrence of $B\psi$ is not strict. Then, we call a formula of the form $B\psi$ or $A\psi$, with $\psi \in \mathcal{L}_M$, an MBNF *modal atom*, or simply *modal atom*. Moreover, given $\varphi \in \mathcal{L}_M$, we call the set of modal atoms occurring in φ the *modal atoms of* φ (and denote such a set as $MA(\varphi)$).

Definition 4. Let $\varphi \in \mathcal{L}_M$ and let (P, N) be a partition of a set of modal atoms. We denote as $\varphi(P, N)$ the formula obtained from φ by substituting each strict occurrence in φ of a formula in P with *true*, and each strict occurrence in φ of a formula in N with *false*.

Observe that only the occurrences in φ of modal subformulas of the form $B\psi$ or $A\psi$ which are *not* within the scope of another MBNF modality are replaced; notice also that, if $P \cup N$ contains $MA(\varphi)$, then $\varphi(P, N)$ is a $\mathcal{K}$-formula. In this case, the pair (P, N) identifies a guess on all the MBNF modal atoms from φ, i.e. P contains the modal atoms of φ assumed to hold, while N contains the modal atoms of φ assumed not to hold.

Example 5. Suppose $\varphi = B(K_1a \vee Ba) \wedge (\neg A(K_2a \vee \neg d) \vee BAK_2b) \wedge c$. Then,

$$MA(\varphi) = \{B(K_1a \vee Ba), Ba, A(K_2a \vee \neg d), BAK_2b, AK_2b\}$$

A possible partition (P, N) of $MA(\varphi)$ is the following:

$$P = \{B(K_1a \vee Ba)\}$$
$$N = \{Ba, A(K_2a \vee \neg d), BAK_2b, AK_2b\}$$

For such a partition, $\varphi(P, N) = \mathit{true} \wedge (\neg false \vee \mathit{false}) \wedge c$. □

Definition 6. Let $\varphi \in \mathcal{L}_M$ and let (P, N) be a partition of $MA(\varphi)$. We denote as $ob(P, N)$ the $\mathcal{K}$-formula

$$ob(P, N) = \bigwedge_{B\psi \in P} \psi(P, N)$$

Roughly speaking, the $\mathcal{K}$-formula $ob(P, N)$ represents the "objective knowledge" implied by the guess (P, N) on the formulas of the form $B\psi$ belonging to P. From the semantic viewpoint, in each structure (I, M, M') satisfying the guess on the modal atoms given by (P, N) (i.e., such that each modal atom in P is satisfied by (I, M, M') and each modal atom in N is not satisfied by (I, M, M')), the $\mathcal{K}$-formula $ob(P, N)$ constrains the $\mathcal{K}$-interpretations of M, since in each such structure the $\mathcal{K}$-formula $ob(P, N)$ must be satisfied by each $\mathcal{K}$-interpretation $J \in M$, i.e. $J \models ob(P, N)$, while the $\mathcal{K}$-formula $\varphi(P, N)$ constrains the $\mathcal{K}$-interpretation I, since $\varphi(P, N)$ must be satisfied by I.

Example 7. Consider again the partition (P, N) of $MA(\varphi)$ in the previous example:

$$P = \{B(K_1 a \vee Ba)\}$$
$$N = \{Ba, A(K_2 a \vee \neg d), BAK_2 b, AK_2 b\}$$

For such a partition, the only formula of the form $B\psi$ belonging to P is $B(K_1 a \vee Ba)$, and

$$(K_1 a \vee Ba)(P, N) = K_1 a \vee \mathit{false}$$

Therefore, $ob(P, N) = K_1 a \vee \mathit{false}$. □

Definition 8. We say that a pair of sets of $\mathcal{K}$-interpretations (M, M') *induces* the partition (P, N) of $MA(\varphi)$ if, for each modal atom $\xi \in MA(\varphi)$, $\xi \in P$ iff, for each $\mathcal{K}$-interpretation I, $(I, M, M') \models \xi$.

Lemma 9. Let $\varphi \in \mathcal{L}_M$, let I be a $\mathcal{K}$-interpretation, let M, M' be sets of $\mathcal{K}$-interpretations, and let (P, N) be the partition induced by (M, M') on a set of modal atoms S. Then, $(I, M, M') \models \varphi$ iff $(I, M, M') \models \varphi(P, N)$.

Proof. Follows immediately from Definition 4, Definition 8, and Definition 1. □

We now show that, if (I, M, M) is an MBNF($\mathcal{K}$) model for φ which induces the partition (P, N) of $MA(\varphi)$, then the $\mathcal{K}$-formula $ob(P, N)$ completely characterizes the set of $\mathcal{K}$-interpretations M.

Theorem 10. Let $\varphi \in \mathcal{L}_M$, let (I, M, M) be an MBNF($\mathcal{K}$) model for φ, and let (P, N) be the partition of $MA(\varphi)$ induced by (M, M). Then, $M = \{J : J \models ob(P, N)\}$.

Proof. Let $M' = \{J : J \models ob(P, N)\}$. Since (M, M) induces the partition (P, N), by Definition 8 it follows that each $\mathcal{K}$-interpretation in M must satisfy

$ob(P, N)$, hence $M \subseteq M'$. Now suppose $M \subset M'$, and consider the structure (I, M', M). We prove that each modal atom $\xi \in MA(\varphi)$ belongs to P iff $(I, M', M) \models \xi$. The proof is by induction on the depth of formulas in $MA(\varphi)$.

First, consider a modal atom $A\psi$ such that $\psi \in \mathcal{L}$: from the definition of satisfiability of a formula in an MBNF($\mathcal{K}$) structure, it follows immediately that $A\psi \in P$ iff $(I, M', M) \models A\psi$. Then, consider a modal atom $B\psi$ such that $\psi \in \mathcal{L}$: if $B\psi \in P$, then, by definition of $ob(P, N)$, the $\mathcal{K}$-formula $ob(P, N) \supset \psi$ is valid, therefore $(I, M', M) \models B\psi$. If $B\psi \in N$, then there exists a $\mathcal{K}$-interpretation J in M such that $J \not\models \psi$, and since $M' \supset M$, it follows that $(I, M', M) \not\models B\psi$. Hence, each modal atom $\xi \in MA(\varphi)$ of depth 1 belongs to P iff $(I, M', M) \models \xi$.

Suppose now that $\xi \in P$ iff $(I, M', M) \models \xi$ for each modal atom ξ in $MA(\varphi)$ of depth less or equal to i. Consider a modal atom $B\psi$ of $MA(\varphi)$ of depth $i + 1$: by the induction hypothesis, and by Lemma 9, $(I, M', M) \models B\psi$ iff $M' \models B(\psi(P, N))$. Now, if $B\psi \in P$, then, by definition of $ob(P, N)$, the $\mathcal{K}$-formula $ob(P, N) \supset \psi(P, N)$ is valid, and since $M' = \{J : J \models ob(P, N)\}$, it follows that $M' \models B(\psi(P, N))$, which in turn implies $(I, M', M) \models B\psi$; on the other hand, if $B\psi \in N$, then there exists a $\mathcal{K}$-interpretation J in M such that $(J, M, M) \not\models \psi$, hence, by the induction hypothesis and Lemma 9, $J \not\models \psi(P, N)$. Now, since $M' \supset M$, it follows that $M' \not\models B(\psi(P, N))$, hence $(I, M', M) \not\models B\psi$. In the same way it is possible to show that a modal atom of the form $A\psi$ of depth $i + 1$ belongs to P iff $(I, M', M) \models A\psi$.

We have thus proved that each modal atom $\xi \in MA(\varphi)$ belongs to P iff $(I, M', M) \models \xi$: this in turn implies that $(I, M', M) \models \varphi$ iff $I \models \varphi(P, N)$, and since by hypothesis (I, M, M) satisfies φ and (P, N) is the partition of $MA(\varphi)$ induced by (M, M), by Lemma 9 it follows that $I \models \varphi(P, N)$. Therefore, $(I, M', M) \models \varphi$, which contradicts the hypothesis that (I, M) is an MBNF($\mathcal{K}$) model for φ. Consequently, $M' = M$, which proves the thesis. □

Informally, the above theorem states that each MBNF($\mathcal{K}$) model for φ can be associated with a partition (P, N) of the modal atoms of φ; moreover, the $\mathcal{K}$-formula $ob(P, N)$ exactly characterizes the set of $\mathcal{K}$-interpretations M of an MBNF($\mathcal{K}$) model (I, M, M), in the sense that M is the set of *all* $\mathcal{K}$-interpretations satisfying $ob(P, N)$. This provides a finite way to describe all MBNF($\mathcal{K}$) models for φ.

We now define the notion of a partition of a set of modal atoms induced by a pair of $\mathcal{K}$-formulas.

Definition 11. Let $\sigma \in \mathcal{L}_M$, $\varphi, \psi \in \mathcal{L}$. We denote as $Prt(\sigma, \varphi, \psi)$ the partition of $MA(\sigma)$ induced by (M_1, M_2), where M_1 is the set of $\mathcal{K}$-interpretations satisfying φ and M_2 is the set of $\mathcal{K}$-interpretations satisfying ψ.

In order to simplify notation, we denote as $Prt(\sigma, \varphi)$ the partition $Prt(\sigma, \varphi, \varphi)$. The following theorem provides a constructive way to build the partition $Prt(\sigma, \varphi, \psi)$.

Theorem 12. Let $\sigma \in \mathcal{L}_M$, $\varphi, \psi \in \mathcal{L}$. Let (P, N) be the partition of $MA(\sigma)$ built as follows:

1. start from $P = N = \emptyset$;
2. for each modal atom $B\xi$ in $MA(\sigma)$ such that $\xi(P, N) \in \mathcal{L}$, if the $\mathcal{K}$-formula $\varphi \supset \xi(P, N)$ is valid, then add $B\xi$ to P, otherwise add $B\xi$ to N;
3. for each modal atom $A\xi$ in $MA(\sigma)$ such that $\xi(P, N) \in \mathcal{L}$, if the $\mathcal{K}$-formula $\psi \supset \xi(P, N)$ is valid, then add $A\xi$ to P, otherwise add $A\xi$ to N;
4. iteratively apply the above rules until $P \cup N = MA(\sigma)$.

Then, $(P, N) = Prt(\sigma, \varphi, \psi)$.

Proof. In the following, we say that a formula ξ has MBNF-depth i if each subformula in ξ lies within the scope of at most i modalities B or A, and there exists a subformula in ξ which lies within the scope of exactly i modalities B or A.

The proof is by induction on the structure of the formulas in $MA(\sigma)$. First, from the fact that $Prt(\sigma, \varphi, \psi)$ is the partition induced by (M, M'), with $M = \{I : I \models \varphi\}$, $M' = \{I : I \models \psi\}$, and from the definition of satisfiability in MBNF structures, it follows that, if $\xi \in \mathcal{L}$, then $(M, M') \models B\xi$ if and only if $\varphi \supset \xi$ is a valid $\mathcal{K}$-formula, and $(M, M') \models A\xi$ if and only if $\psi \supset \xi$ is a valid $\mathcal{K}$-formula. Therefore, (P, N) agrees with $Prt(\sigma, \varphi, \psi)$ on all modal atoms of MBNF-depth 1. Suppose now that (P, N) and $Prt(\sigma, \varphi, \psi)$ agree on all modal atoms of MBNF-depth less or equal to i. Consider a modal atom $B\xi$ of $MA(\sigma)$ of MBNF-depth $i+1$. From Lemma 9 and from the definition of satisfiability in MBNF structures, it follows that $(M, M') \models B\xi$ if and only if $\varphi \supset \xi(Prt(\sigma, \varphi, \psi))$ is a valid $\mathcal{K}$-formula, and since by Definition 4 the value of the formula $\xi(Prt(\sigma, \varphi, \psi))$ only depends on the guess of the modal atoms of MBNF-depth less or equal to i in $Prt(\sigma, \varphi, \psi)$, by the induction hypothesis it follows that $\xi(Prt(\sigma, \varphi, \psi)) = \xi(P, N)$, hence $B\xi$ belongs to P if and only if $(M, M') \models B\xi$. Analogously, it can be proven that any modal atom of MBNF-depth $i+1$ of the form $A\xi$ belongs to P if and only if $(M, M') \models A\xi$. Therefore, (P, N) and $Prt(\sigma, \varphi, \psi)$ agree on all modal atoms of MBNF-depth $i+1$. □

Example 13. Let $\mathcal{K} = \mathsf{K}_n$, and let σ be the formula φ of Example 5, i.e.:

$$\sigma = B(K_1 a \vee Ba) \wedge (\neg A(K_2 a \vee \neg d) \vee BAK_2 b) \wedge c$$

We now build the partition $Prt(\sigma, K_1 a, K_1 c)$ of $MA(\sigma)$ according to the above theorem, starting from $P = N = \emptyset$.

1. Since $K_1 a \supset a$ is not K_n-valid, Ba is added to N;
2. since $K_1 c \supset K_2 b$ is not K_n-valid, $AK_2 b$ is added to N;
3. since $K_1 c \supset K_2 a \vee \neg d$ is not K_n-valid, $A(K_2 a \vee \neg d)$ is added to N;
4. now, $(K_1 a \vee Ba)(P, N) = K_1 a \vee \mathit{false}$. Since $K_1 a \supset K_1 a \vee \mathit{false}$ is K_n-valid, $B(K_1 a \vee Ba)$ is added to P;
5. finally, $(AK_2 b)(P, N) = \mathit{false}$. Since $K_1 c \supset \mathit{false}$ is not K_n-valid, $BAK_2 b$ is added to N.

Therefore, $Prt(\sigma, K_1 a, K_1 c) = (P, N)$, where

$$P = \{B(K_1 a \vee Ba)\}$$
$$N = \{Ba, A(K_2 a \vee \neg d), BAK_2 b, AK_2 b\}$$

□

We now define a method for deciding satisfiability of a formula $\varphi \in \mathcal{L}_M$. In particular, we present the algorithm MBNF($\mathcal{K}$)-Sat, reported in Figure 1.

The algorithm extends previous results for the *propositional* fragment of MBNF. In particular, it generalizes the results presented in [21] concerning a finitary characterization of propositional MBNF models and an analogous finite characterization, in terms of partitions of $MA(\varphi)$, of all the models relevant for establishing whether a partition (P, N) of $MA(\varphi)$ identifies a model.

The algorithm checks whether there exists a partition (P, N) of $MA(\varphi)$ satisfying the three conditions (a), (b), (c). Intuitively, the partition cannot be self-contradictory (condition (a)): in particular, the condition $(P, N) = Prt(\varphi, ob(P, N))$ establishes that the objective knowledge implied by the partition (P, N) (that is, the $\mathcal{K}$-formula $ob(P, N)$) identifies a set of $\mathcal{K}$-interpretations $M = \{I : I \models ob(P, N)\}$ such that (M, M) induces the same partition (P, N) on

Algorithm MBNF($\mathcal{K}$)-Sat(φ)
Input: formula $\varphi \in \mathcal{L}_M$;
Output: *true* if φ has a model, *false* otherwise.
begin
if there exists partition (P, N) of $MA(\varphi)$
such that
(a) $(P, N) = Prt(\varphi, ob(P, N))$ **and**
(b) $\varphi(P, N)$ is $\mathcal{K}$-satisfiable **and**
(c) **for each** partition $(P', N') \neq (P, N)$ of $MA(\varphi)$,
(c1) $\varphi(P', N')$ is not $\mathcal{K}$-satisfiable **or**
(c2) $(P', N') \neq Prt(\varphi, ob(P', N'), ob(P, N))$ **or**
(c3) $ob(P, N) \wedge \neg ob(P', N')$ is $\mathcal{K}$-satisfiable
then return *true*
else return *false*
end

Figure 1. Algorithm MBNF($\mathcal{K}$)-Sat.

$MA(\varphi)$. Moreover, the partition must be consistent with φ (condition (b)): such a condition implies that there exists a $\mathcal{K}$-interpretation I such that φ is satisfied in (I, M, M). Moreover, condition (c) corresponds to check whether such a structure (I, M, M) identifies an MBNF($\mathcal{K}$) model for φ according to the preference semantics of MBNF($\mathcal{K}$), i.e. whether there is no pair (J, M') such that $M' \supset M$ and (J, M', M) satisfies φ. Again, the search of such a structure is performed by examining whether there exists a partition of $MA(\varphi)$, different from (P, N), which does not satisfy any of the conditions (c1), (c2), (c3).

We illustrate the algorithm through the following simple example.

Example 14. Suppose

$$\varphi = B(K_1 a \vee Ba) \wedge (\neg A(K_2 a \vee \neg d) \vee BAK_2 b) \wedge c$$

Therefore,

$$MA(\varphi) = \{B(K_1 a \vee Ba), Ba, A(K_2 a \vee \neg d), BAK_2 b, AK_2 b\}$$

1. First, suppose $\mathcal{K} = \mathsf{K}_n$, and consider the partition $(P,N) = (P_1,N_1)$, where

$$P_1 = \{B(K_1a \vee Ba)\}$$
$$N_1 = \{Ba, A(K_2a \vee \neg d), BAK_2b, AK_2b\}$$

Then, $\varphi(P,N) = \mathit{true} \wedge (\neg \mathit{false} \vee \mathit{false}) \wedge c$ (which is K_n-equivalent to c), thus satisfying condition (b) of the algorithm. Moreover, $ob(P,N) = K_1a \vee \mathit{false}$ (which is K_n-equivalent to K_1a). Now, let M be the set of $\mathcal{K}$-interpretations satisfying K_1a: it is easy to see that (M,M) satisfies the modal atoms in P, while it does not satisfy the modal atoms in N, hence $(P,N) = Prt(\varphi, ob(P,N))$, thus satisfying condition (a) of the algorithm. As for condition (c), it is immediate to see that either condition (c1) or condition (c2) holds for each partition of $MA(\varphi)$ different from (P_1,N_1), with the exception of the following one:

$$P' = \{B(K_1a \vee Ba), Ba\}$$
$$N' = \{A(K_2a \vee \neg d), BAK_2b, AK_2b\}$$

Since $ob(P',N') = (K_1a \vee \mathit{true}) \wedge a$, which is K_n-equivalent to a, $ob(P,N) \wedge \neg ob(P',N')$ is K_n-equivalent to $K_1a \wedge \neg a$, hence it is K_n-satisfiable, thus satisfying condition (c3) of the algorithm. Therefore, (P,N) satisfies condition (c) of the algorithm. Consequently, MBNF($\mathcal{K}$)-Sat(φ) returns *true* In fact, the partition (P_1,N_1) identifies the set of MBNF($\mathcal{K}$) models for φ (I,M) such that I is a K_n-interpretation satisfying c and M is the set of K_n-interpretations satisfying K_1a.

2. now suppose $\mathcal{K}$ is a reflexive multimodal logic (i.e., either T_n or $\mathsf{S4}_n$ or $\mathsf{S5}_n$). It is easy to see that in this case the partition (P_1,N_1) does not satisfy condition (a) of the algorithm, which is due to the fact that (M,M) satisfies the modal atom Ba (since the formula $K_1a \supset a$ is valid in all reflexive multimodal logics), which is assumed as false in (P_1,N_1). On the other hand, suppose now that $(P,N) = (P_2,N_2)$, where

$$P_2 = \{B(K_1a \vee Ba), Ba, BAK_2b, AK_2b\}$$
$$N_2 = \{A(K_2a \vee \neg d)\}$$

Then, $\varphi(P,N) = \mathit{true} \wedge (\neg \mathit{false} \vee \mathit{true}) \wedge c$ (which is $\mathcal{K}$-equivalent to c), and $ob(P,N) = (K_1a \vee \mathit{true}) \wedge a \wedge \mathit{true}$, which is $\mathcal{K}$-equivalent to a. It is immediate to verify that $(P,N) = Prt(\varphi, ob(P,N))$, thus satisfying condition (a) of the algorithm, and that either condition (c1) or condition (c2) holds for each partition of $MA(\varphi)$ different from (P_2, N_2), therefore condition (c) holds for $(P, N) = (P_2, N_2)$. Consequently, MBNF($\mathcal{K}$)-Sat(φ) returns *true*. In fact, the partition (P_2, N_2) identifies the set of MBNF($\mathcal{K}$) models for φ (I, M) such that I is a $\mathcal{K}$-interpretation satisfying c and M is the set of $\mathcal{K}$-interpretations satisfying a.

□

We now prove soundness and completeness of the algorithm MBNF($\mathcal{K}$)-Sat. To this aim, we need the following preliminary lemma.

Lemma 15. Let $\varphi \in \mathcal{L}_M$, and let (P, N) be the partition of $MA(\varphi)$ induced by (M', M). Let $M'' = \{I : I \models ob(P, N)\}$. Then, (P, N) is the partition induced by (M'', M).

Proof. The proof is by induction on the depth of the modal atoms of $MA(\varphi)$. Let $A\psi \in MA(\varphi)$ such that $\psi \in \mathcal{L}$: then, $(M', M) \models A\psi$ iff, for each $\mathcal{K}$-interpretation $I \in M$, $I \models \psi$, therefore $(M', M) \models A\psi$ iff $(M'', M) \models A\psi$. Now let $B\psi \in MA(\varphi)$ such that $\psi \in \mathcal{L}$: by Definition 6, $(M', M) \models B\psi$ iff the $\mathcal{K}$-formula $ob(P, N) \supset \psi$ is valid, and since $M'' = \{I : I \models ob(P, N)\}$, it follows that $(M', M) \models B\psi$ iff $(M'', M) \models B\psi$.

Now suppose that, for each modal atom ξ of depth i, $(M', M) \models \xi$ iff $(M'', M) \models \xi$, and let (P', N') denote the partition of the modal atoms in $MA(\varphi)$ of depth less or equal to i induced by (M', M). First, consider a modal atom $A\psi$ of depth $i + 1$. Then, by Lemma 9, $(M', M) \models A\psi$ iff $(M', M) \models A(\psi(P', N'))$ and, by the inductive hypothesis and Lemma 9, $(M'', M) \models A\psi$ iff $(M'', M) \models A(\psi(P', N'))$. Then, since ψ has depth i, $\psi(P', N')$ is a $\mathcal{K}$-formula, hence $(M', M) \models A(\psi(P', N'))$ iff, for each $\mathcal{K}$-interpretation $I \in M$, $I \models \psi(P', N')$, which immediately implies that $(M', M) \models A\psi$ iff $(M'', M) \models A\psi$. Now consider a modal atom $B\psi$ of depth $i + 1$. Then, by Lemma 9, $(M', M) \models B\psi$ iff $(M', M) \models B(\psi(P', N'))$ and, by the inductive hypothesis and Lemma 9, $(M'', M) \models B\psi$ iff $(M'', M) \models B(\psi(P', N'))$. By Definition 6, $(M', M) \models B\psi$ iff the $\mathcal{K}$-formula $ob(P, N) \supset \psi(P', N')$ is valid, and since $M'' = \{I : I \models ob(P, N)\}$, it follows that $(M', M) \models B\psi$ iff $(M'', M) \models B\psi$, which proves the thesis. □

Theorem 16. Let $\varphi \in \mathcal{L}_M$. Then, MBNF($\mathcal{K}$)-Sat($\varphi$) returns *true* iff φ is MBNF($\mathcal{K}$)-satisfiable.

Proof. *If part.* Suppose φ is MBNF($\mathcal{K}$)-satisfiable. Then, there exists a $\mathcal{K}$-interpretation I and a cluster M such that (I, M, M) is an MBNF($\mathcal{K}$) model for φ. Let (P, N) be the partition of $MA(\varphi)$ induced by (M, M). By Theorem 10, $M = \{I : I \models ob(P, N)\}$. Therefore, by Definition 11, $(P, N) = Prt(\varphi, ob(P, N))$, hence condition (a) in the algorithm holds. Then, since $(I, M, M) \models \varphi$, by Lemma 9 $I \models \varphi(P, N)$, hence condition (b) in the algorithm holds. Now suppose there exists a partition (P', N') of $MA(\varphi)$ such that $(P', N') \neq (P, N)$ and none of conditions (c1), (c2), and (c3) holds. Then, since $\varphi(P', N')$ is $\mathcal{K}$-satisfiable, there exists a $\mathcal{K}$-interpretation J such that $J \models \varphi(P', N')$, and since $(P', N') = Prt(\varphi, ob(P', N'), ob(P, N))$, from Lemma 9 it follows that there exists a $\mathcal{K}$-interpretation J such that $(J, M', M) \models \varphi$, where $M' = \{I : I \models ob(P', N')\}$. Then, since condition (c3) does not hold, the $\mathcal{K}$-formula $ob(P, N) \supset ob(P', N')$ is valid, which implies that $M' \supseteq M$. Now, if $M' = M$, then (P', N') would be the partition induced by (M, M), thus contradicting the hypothesis $(P', N') \neq (P, N)$. Hence, $M' \supset M$, and since $(J, M', M) \models \varphi$, it follows that (I, M, M) is not an MBNF($\mathcal{K}$) model for φ. Contradiction. Therefore, condition (c) in the algorithm holds, consequently MBNF($\mathcal{K}$)-Sat(φ) returns *true*.

Only-if part. Suppose MBNF($\mathcal{K}$)-Sat(φ) returns *true*. Then, there exists a partition (P, N) of $MA(\varphi)$ such that conditions (a), (b), and (c) hold. Let $M = \{I : I \models ob(P, N)\}$. Since $(P, N) = Prt(\varphi, ob(P, N))$, by Definition 11 (P, N) is the partition induced by (M, M). And since $\varphi(P, N)$ is $\mathcal{K}$-satisfiable, it follows that there exists a $\mathcal{K}$-interpretation I such that $I \models \varphi(P, N)$, hence, by Lemma 9, $(I, M, M) \models \varphi$. Now suppose (I, M, M) is not an MBNF($\mathcal{K}$) model for φ. Then, there exists a cluster M' and a $\mathcal{K}$-interpretation J such that $M' \supset M$ and $(J, M', M) \models \varphi$. Let (P', N') be the partition of $MA(\varphi)$ induced by (M', M). Since $M = \{I : I \models ob(P, N)\}$, it follows that M' contains at least one $\mathcal{K}$-interpretation J which does not satisfy $ob(P, N)$, and since $ob(P, N) = \bigwedge_{B\psi \in P} \psi(P, N)$, J does not satisfy at least one formula of the form $\psi(P, N)$ such that $B\psi \in P$. Therefore, $P' \neq P$, which implies that $(P', N') \neq (P, N)$. Then, since $(J, M', M) \models \varphi$, by Lemma 9 $J \models \varphi(P', N')$, hence $\varphi(P', N')$ is $\mathcal{K}$-satisfiable. Now let $M'' = \{I : I \models ob(P', N')\}$. By Lemma 15, it follows that (P', N') is the partition induced by (M'', M), therefore, by Definition 11, $(P', N') = Prt(\varphi, ob(P', N'), ob(P, N))$. Moreover, since $M' \supset M$, it follows that the $\mathcal{K}$-formula $ob(P, N) \supset ob(P', N')$ is valid, hence the formula $ob(P, N) \wedge \neg ob(P', N')$ is $\mathcal{K}$-unsatisfiable. Consequently, (P', N') does not satisfy condition (c) in the algorithm, thus contradicting the hypothesis. Therefore, (I, M, M) is an MBNF($\mathcal{K}$) model for φ, thus proving the thesis. □

As for entailment in MBNF($\mathcal{K}$), even though the deduction theorem does not hold in this formalism (as in other nonmonotonic logics), it turns out that, for subjective formulas, it is possible to easily reduce entailment to (un)satisfiability in this logic.

In the following, we denote as $\varphi[B/A]$ the formula obtained from $\varphi \in \mathcal{L}_M$ by replacing each occurrence of the modality B with the modality A.

Theorem 17. Let $\sigma \in \mathcal{L}_M$, $\varphi \in \mathcal{L}_M^S$. Then, $\sigma \models_{\mathrm{MBNF}(\mathcal{K})} \varphi$ iff the formula $\sigma \wedge (\neg\varphi[B/A])$ is MBNF($\mathcal{K}$)-unsatisfiable.

Proof. If-part. Suppose $\sigma \not\models_{\mathrm{MBNF}} \varphi$. Then, there exists an MBNF($\mathcal{K}$) model (I, M, M) for σ such that $(I, M, M) \not\models \varphi$. Therefore, $(I, M, M) \models \neg\varphi$, and, by Definition 1, $(I, M, M) \models \neg\varphi[B/A]$, since the B-cluster and the A-cluster coincide in (I, M, M). Since (I, M, M) is an MBNF($\mathcal{K}$) model for σ, it follows that $(I, M, M) \models \sigma \wedge (\neg\varphi[B/A])$; moreover, from Definition 2, for each $\mathcal{K}$-interpretation J and for each cluster M' such that $M' \supset M$, $(J, M', M) \not\models \sigma$, which implies that $(J, M', M) \not\models \sigma \wedge (\neg\varphi[B/A])$. Consequently, (I, M, M) is an MBNF($\mathcal{K}$) model for $\sigma \wedge (\neg\varphi[B/A])$, which implies that $\sigma \wedge (\neg\varphi[B/A])$ is MBNF($\mathcal{K}$)-satisfiable.

Only-if part. Suppose $\sigma \wedge (\neg\varphi[B/A])$ is MBNF($\mathcal{K}$)-satisfiable. Then, there exist a $\mathcal{K}$-interpretation I and a cluster M such that (I, M, M) is an MBNF($\mathcal{K}$) model for $\sigma \wedge (\neg\varphi[B/A])$, hence, by Definition 2, $(I, M, M) \models \sigma \wedge (\neg\varphi[B/A])$ and, for each $\mathcal{K}$-interpretation J and for each cluster M' such that $M' \supset M$, $(J, M', M) \not\models \sigma \wedge (\neg\varphi[B/A])$, that is, either $(J, M', M) \not\models \sigma$ or $(J, M', M) \not\models \neg\varphi[B/A]$. Now, since $\varphi \in \mathcal{L}_M^S$, the evaluation of $\neg\varphi[B/A]$ in (J, M', M) does not depend on the $\mathcal{K}$-interpretation J; moreover, since in $\varphi[B/A]$ there are no occurrences of the operator B, the evaluation of $\neg\varphi[B/A]$ in (J, M', M) does not depend on the B-cluster M'. Therefore, the evaluation of $\neg\varphi[B/A]$ in (J, M', M) only depends on the A-cluster M, and since $(I, M, M) \models \neg\varphi[B/A]$, it follows that, for each $\mathcal{K}$-interpretation J and for each cluster M', $(J, M', M) \models \neg\varphi[B/A]$, and since either $(J, M', M) \not\models \sigma$ or $(J, M', M) \not\models \neg\varphi[B/A]$, it follows that $(J, M', M) \not\models \sigma$. Consequently, by Definition 2, (I, M, M) is an MBNF($\mathcal{K}$) model for σ. Moreover, since $(I, M, M) \models \neg\varphi[B/A]$, from Definition 1 it follows that

$(I, M, M) \models \neg\varphi$, since the B-cluster coincides with the A-cluster in (I, M, M). Therefore, $(I, M, M) \not\models \varphi$, which implies that $\sigma \not\models_{\mathrm{MBNF}(\mathcal{K})} \varphi$. □

Based on the above theorem, we are able to decide entailment in MBNF($\mathcal{K}$) through the algorithm MBNF($\mathcal{K}$)-Sat.

As for the implementation of a reasoning procedure for MBNF($\mathcal{K}$), we remark that the algorithm MBNF($\mathcal{K}$)-Sat is essentially a method for reducing satisfiability in MBNF($\mathcal{K}$) to a number of satisfiability problems in the logic $\mathcal{K}$. This allows for easily implementing a procedure for MBNF($\mathcal{K}$)-satisfiability on top of an already implemented satisfiability solver for multimodal logics, like KSAT [8] or FaCT [14].

4. Computational characterization

We now provide a computational characterization of reasoning in MBNF($\mathcal{K}$), and analyze some syntactic restrictions of MBNF($\mathcal{K}$) which reduce the worst-case complexity of reasoning.

In particular, we have identified four different syntactic restrictions that affect the complexity of reasoning in MBNF($\mathcal{K}$):

1. the number of modeled agents;
2. the choice of the multimodal system $\mathcal{K}$;
3. the depth of nesting of the $\mathcal{K}$ modalities $K_1, \ldots, K_n$;
4. the depth of nesting of the MBNF operator B.

In the following, we say that a modal formula φ has *B-depth i* ($i \geq 0$) if there is a subformula of φ which lies within the scope of i nested occurrences of the modality B, and there is no subformula of φ which lies within the scope of $i+1$ occurrences of the modality B. Moreover, we say that φ has *K-depth* i ($i \geq 0$) if there is a subformula of φ which lies within the scope of i nested occurrences of the modalities $K_1, \ldots, K_n$, and there is no subformula of φ which lies within the scope of $i+1$ occurrences of the modalities $K_1, \ldots, K_n$.

The rest of this section is organized as follows: we first recall the complexity of reasoning in the multimodal logics $\mathcal{K} \in \{\mathsf{K}_n, \mathsf{T}_n, \mathsf{S4}_n, \mathsf{KD45}_n, \mathsf{S5}_n\}$ under various syntactic restrictions [12,9]; then, we first establish the complexity of MBNF($\mathcal{K}$)-satisfiability with respect to the different choices of the multimodal logic $\mathcal{K}$, and then establish the complexity of MBNF($\mathcal{K}$)-satisfiability for the various multimodal logics $\mathcal{K}$ when we impose a bound on the depth of nesting of the modality B; finally, we summarize the complexity results obtained.

4.1. Complexity results for multimodal logics

We first briefly recall the complexity classes in the *polynomial hierarchy*, and refer to [15,19] for further details about the complexity classes mentioned in the paper. The class NP contains all problems that can be solved by a nondeterministic Turing machine in polynomial time. The class coNP comprises all problems that are the complement of a problem in NP. P^A (NP^A) is the class of problems that are solved in polynomial time by deterministic (nondeterministic) Turing machines using an oracle for A (i.e. that solves in constant time any problem in A). The classes Σ_k^p, Π_k^p and Δ_k^p of the polynomial hierarchy are defined by $\Sigma_0^p = \Pi_0^p = \Delta_0^p = \mathrm{P}$, and for $k \geq 0$, $\Sigma_{k+1}^p = \mathrm{NP}^{\Sigma_k^p}$, $\Pi_{k+1}^p = \mathrm{co}\Sigma_{k+1}^p$ and $\Delta_{k+1}^p = \mathrm{P}^{\Sigma_k^p}$. In particular, the complexity class Σ_2^p is the class of problems that are solved in polynomial time by a nondeterministic Turing machine that uses an NP-oracle, and Π_2^p is the class of problems that are complement of a problem in Σ_2^p, while Σ_3^p is the class of problems that are solved in polynomial time by a nondeterministic Turing machine that uses an Σ_2^p-oracle, and Π_3^p is the class of problems that are complement of a problem in Σ_3^p. It is generally assumed that the polynomial hierarchy does not collapse: hence, a problem in the class Σ_2^p or Π_2^p is considered computationally easier than a Σ_3^p-hard or Π_3^p-hard problem. Finally, PSPACE is the class of problems that can be solved by a Turing machine that uses a polynomially bounded amount of memory. It is known that PSPACE contains all problems in the polynomial hierarchy: moreover, a PSPACE-hard problem is generally considered computationally harder than a Σ_k^p-hard or Π_k^p-hard problem, for any given k.

The computational results for multimodal logics obtained in [12,9] are summarized in the table reported in Figure 2. The table must be read as follows: for each row of the table, the complexity of checking $\mathcal{K}$-satisfiability under the syntactic restriction (on the depth of nesting of the modal operators) reported in the first column (i.e., considering only the subset of $\mathcal{K}$-formulas satisfying the restriction) is complete with respect to the complexity class appearing in the other columns, where each column corresponds to a different choice of the multimodal

	$\mathcal{K} \in \{\mathsf{KD45}_1, \mathsf{S5}_1\}$	$\mathcal{K} \in \{\mathsf{KD45}_n, \mathsf{S5}_n\}$ $(n \geq 2)$	$\mathcal{K} \in \{\mathsf{K}_n, \mathsf{T}_n\}$ $(n \geq 1)$	$\mathcal{K} = \mathsf{S4}_n$ $(n \geq 1)$
K-depth ≤ 1	NP	NP	NP	NP
K-depth $\leq k$ $(k \geq 2)$	NP	NP	NP	PSPACE
no restrictions	NP	PSPACE	PSPACE	PSPACE

Figure 2. Complexity of $\mathcal{K}$-satisfiability [12,9]

logic $\mathcal{K}$ and of the number of agents modeled.

4.2. Complexity of MBNF($\mathcal{K}$)*-satisfiability*

In the following, we say that a language $\mathcal{L}' \subseteq \mathcal{L}$ is *closed under boolean composition* if, for each $\varphi_1, \varphi_2 \in \mathcal{L}'$, $\varphi_1 \wedge \varphi_2 \in \mathcal{L}'$ and $\neg\varphi_1 \in \mathcal{L}'$. Moreover, we denote as $\mathcal{L}'_M$ the subset of $\mathcal{L}_M$ built upon $\mathcal{L}'$, i.e., the modal extension of $\mathcal{L}'$ with the modalities B and A obtained according to the following abstract syntax:

$$\varphi = \psi \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid B\varphi \mid A\varphi$$

where $\psi \in \mathcal{L}'$.

Theorem 18. Let $\mathcal{K} \in \{\mathsf{K}_n, \mathsf{T}_n, \mathsf{S4}_n, \mathsf{KD45}_n, \mathsf{S5}_n\}$, let $\mathcal{L}' \subseteq \mathcal{L}$, and let $\mathcal{L}'$ be closed under boolean composition. If $\mathcal{K}$-satisfiability for formulas from $\mathcal{L}'$ is a PSPACE-complete problem, then MBNF($\mathcal{K}$)-satisfiability for formulas from $\mathcal{L}'_M$ is PSPACE-complete.

Proof. PSPACE-hardness follows from the fact that $\mathcal{L}'_M \supset \mathcal{L}'$ and from the fact that, if $\varphi \in \mathcal{L}'$, φ is MBNF($\mathcal{K}$)-satisfiable if and only if φ is $\mathcal{K}$-satisfiable.

To prove membership in PSPACE, we analyze the complexity of the algorithm MBNF($\mathcal{K}$)-Sat reported in Figure 1. In particular, observe that:

- given (P, N), the formula $ob(P, N)$ belongs to $\mathcal{L}'$, because $\varphi \in \mathcal{L}'_M$ and $\mathcal{L}'$ is closed under boolean composition. Moreover, $ob(P, N)$ can be computed in polynomial time with respect to the size of P, hence, by Theorem 12, since $MA(\varphi)$ has size linear with respect to the size of φ, construction of the partition $Prt(\varphi, ob(P, N))$ can be performed by solving a linear number (with respect to the size of φ) of $\mathcal{K}$-satisfiability problems for formulas from $\mathcal{L}'$. Therefore, condition (a) can be checked through a linear number (in the size of the input) of calls to a PSPACE-oracle;

- Since, given φ and (P, N), $\varphi(P, N)$ can be computed in polynomial time with respect to the size of the input, and since $\varphi(P, N) \in \mathcal{L}'$, it follows that condition (b) can be computed in PSPACE;
- given a partition (P', N'), each of the conditions (c1), (c2) and (c3) (analogous to conditions (a) and (b)) can be checked in polynomial time, with respect to the size of φ, using a PSPACE-oracle. Therefore, since the guess of the partition (P', N') of $MA(\varphi)$ requires a nondeterministic choice, falsity of condition (c) can be decided in $\mathrm{NP}^{\mathrm{PSPACE}}$, and since $\mathrm{NP}^{\mathrm{PSPACE}}$=PSPACE [15], verifying whether condition (c) holds can be decided in PSPACE.

Since the guess of the partition (P, N) of $MA(\varphi)$ requires a nondeterministic choice, it follows that the algorithm MBNF($\mathcal{K}$)-Sat, if considered as a nondeterministic procedure, decides satisfiability of φ in nondeterministic polynomial time (with respect to the size of φ), using a PSPACE-oracle. And since $\mathrm{NP}^{\mathrm{PSPACE}}$=PSPACE, from Theorem 16 we obtain an upper bound of PSPACE for the satisfiability problem in MBNF($\mathcal{K}$). □

The above theorem and the results reported in the table of Figure 2 imply that, for each language $\mathcal{L}'$ and choice of the modal system $\mathcal{K}$ corresponding to an entry of such a table for which $\mathcal{K}$-satisfiability is PSPACE-complete, MBNF($\mathcal{K}$)-satisfiability in the corresponding language $\mathcal{L}'_M$ is PSPACE-complete as well.

In particular, in the general case, i.e., when $\mathcal{L}' = \mathcal{L}$ and hence $\mathcal{L}'_M = \mathcal{L}_M$, MBNF($\mathcal{K}$)-satisfiability is PSPACE-complete for each choice of $\mathcal{K}$ in $\{\mathsf{K}_n, \mathsf{T}_n, \mathsf{S4}_n, \mathsf{KD45}_n, \mathsf{S5}_n\}$. Moreover, from the above theorem and Theorem 17 it immediately follows that deciding entailment in MBNF($\mathcal{K}$) (when the formula entailed is in $\mathcal{L}^S_M$) is also a PSPACE-complete problem for each choice of $\mathcal{K}$ in $\{\mathsf{K}_n, \mathsf{T}_n, \mathsf{S4}_n, \mathsf{KD45}_n, \mathsf{S5}_n\}$.

Theorem 19. Let $\mathcal{K} \in \{\mathsf{K}_n, \mathsf{T}_n, \mathsf{S4}_n, \mathsf{KD45}_n, \mathsf{S5}_n\}$, let $\mathcal{L}' \subseteq \mathcal{L}$, and let $\mathcal{L}'$ be closed under boolean composition. If $\mathcal{K}$-satisfiability for formulas from $\mathcal{L}'$ is NP-complete, then MBNF($\mathcal{K}$)-satisfiability for formulas from $\mathcal{L}'_M$ is Σ^p_3-complete.

Proof. Σ_3^p-hardness follows from the fact that, since $\mathcal{L}'$ is closed under boolean composition, $\mathcal{L}'_M$ is a superset of the extension of the propositional language with the modalities B and A. For such a language, it has been proven in [21] that satisfiability in MBNF is Σ_3^p-complete. Therefore, MBNF($\mathcal{K}$)-satisfiability in $\mathcal{L}'_M$ is Σ_3^p-hard.

To prove membership in Σ_3^p, we analyze the complexity of the algorithm MBNF($\mathcal{K}$)-Sat reported in Figure 1. In particular, observe that:

- since, by Theorem 12, construction of the partition $Prt(\varphi, ob(P, N))$ can be performed by solving a linear number (with respect to the size of φ) of $\mathcal{K}$-satisfiability problems for formulas from $\mathcal{L}'$. Therefore, condition (a) can be checked through a linear number (in the size of the input) of calls to an NP-oracle;
- Since, given φ and (P, N), $\varphi(P, N)$ can be computed in polynomial time with respect to the size of the input, and since $\varphi(P, N) \in \mathcal{L}'$, it follows that condition (b) can be computed in NP;
- given a partition (P', N'), each of the conditions (c1), (c2) and (c3) (analogous to conditions (a) and (b)) can be checked in polynomial time, with respect to the size of φ, using an NP-oracle. In particular, falsity of condition (c3) can be decided in coNP. Therefore, since the guess of the partition (P', N') of $MA(\varphi)$ requires a nondeterministic choice, falsity of condition (c) can be decided in $\mathrm{NP}^{\mathrm{coNP}} = \Sigma_2^p$.

Since the guess of the partition (P, N) of $MA(\varphi)$ requires a nondeterministic choice, it follows that the algorithm MBNF($\mathcal{K}$)-Sat, if considered as a nondeterministic procedure, decides satisfiability of φ in nondeterministic polynomial time (with respect to the size of φ), using a Σ_2^p-oracle. Therefore, from Theorem 16 we obtain an upper bound of Σ_3^p for the satisfiability problem in MBNF($\mathcal{K}$). □

The above theorem and the results reported in the table of Figure 2 imply that, for each language $\mathcal{L}'$ and choice of the modal system $\mathcal{K}$ corresponding to an entry of such a table for which $\mathcal{K}$-satisfiability is NP-complete, MBNF($\mathcal{K}$)-satisfiability in the corresponding language $\mathcal{L}'_M$ is Σ_3^p-complete.

4.3. Complexity for B-depth 1

We now analyze reasoning in MBNF($\mathcal{K}$) when we restrict to formulas of B-depth 1. In particular, we define a specialized algorithm for deciding

Algorithm B-depth-1-Sat(φ)
Input: formula $\varphi \in \mathcal{L}^1_M$;
Output: *true* if φ is MBNF($\mathcal{K}$)-satisfiable, *false* otherwise.
begin
if there exists partition (P, N) of $MA(\varphi)$
such that
(a) $(P, N) = Prt(\varphi, ob(P, N))$ **and**
(b) $\varphi(P, N)$ is $\mathcal{K}$-satisfiable **and**
(c) **for each** partition (P'', N'') of $P - P_A$ **such that** $N'' \neq \emptyset$,
(c1) $\varphi(P_A \cup P'', N \cup N'')$ is not $\mathcal{K}$-satisfiable **or**
(c2) **there exists** $B\psi \in N''$
such that $(\bigwedge_{B\xi \in P''} \xi) \wedge \neg\psi$ is not $\mathcal{K}$-satisfiable
then return *true*
else return *false*
end

Figure 3. Algorithm B-depth-1-Sat.

MBNF($\mathcal{K}$)-satisfiability of formulas of B-depth 1, and characterize the complexity of MBNF($\mathcal{K}$)-satisfiability for such formulas.

In the following, we denote with $\mathcal{L}^1_M$ the set of MBNF($\mathcal{K}$) formulas of B-depth 1, that is the set of formulas from $\mathcal{L}_M$ of the form φ such that each subformula occurring in φ lies within the scope of at most one modality B.

In Figure 3 we report the algorithm B-depth-1-Sat for computing MBNF($\mathcal{K}$)-satisfiability of a formula of B-depth 1. In the algorithm, P_A denotes the subset of modal atoms from P prefixed by the modality A, i.e. $P_A = \{A\psi : A\psi \in P\}$.

Informally, the algorithms MBNF($\mathcal{K}$)-Sat and B-depth-1-Sat only differ in the way in which it is verified whether the MBNF($\mathcal{K}$) structure associated with a partition (P, N) satisfies the preference semantics provided by Definition 2, which is implemented through condition (c) in both algorithms. In the algorithm MBNF($\mathcal{K}$)-Sat, a partition is checked against all other partitions of $MA(\varphi)$, while in the algorithm B-depth-1-Sat, due to the fact that each formula within the scope of a modality B is a $\mathcal{K}$-formula, it is sufficient to verify, in a simpler way, the partition (P, N) against the partitions of $MA(\varphi)$ that agree with (P, N) in the evaluation of all modal atoms in N and in the evaluation of all modal atoms of the form $A\psi$ belonging to P. As shown in the following, such a difference reflects the different computational properties of the MBNF($\mathcal{K}$)-satisfiability problem in the two cases.

Theorem 20. Let $\varphi \in \mathcal{L}_M^1$. Then, B-depth-1-Sat(φ) returns *true* iff φ is MBNF($\mathcal{K}$)-satisfiable.

Proof. We prove that, given $\varphi \in \mathcal{L}_M^1$, B-depth-1-Sat(φ) returns *true* if and only if MBNF($\mathcal{K}$)-Sat(φ) returns *true*, that is, given a partition (P, N) of $MA(\varphi)$, condition (c) in the algorithm B-depth-1-Sat holds for (P, N) if and only if condition (c) in the algorithm MBNF($\mathcal{K}$)-Sat holds for (P, N). Let (P'', N'') be a partition of $P - P_A$ such that $N' \neq \emptyset$ and let $P' = P_A \cup P''$ and $N' = N \cup N''$. It is immediate to see that if (P'', N'') does not satisfy condition (c1) in the algorithm B-depth-1-Sat, then (P', N') does not satisfy condition (c1) in the algorithm MBNF($\mathcal{K}$)-Sat; moreover, if (P'', N'') does not satisfy condition (c2) in the algorithm B-depth-1-Sat, then, by Definition 11, (P', N') does not satisfy condition (c2) in the algorithm MBNF($\mathcal{K}$)-Sat. Therefore, if condition (c) in the algorithm B-depth-1-Sat does not hold for (P, N), then condition (c) in the algorithm MBNF($\mathcal{K}$)-Sat does not hold for (P, N).

Conversely, suppose condition (c) in the algorithm MBNF($\mathcal{K}$)-Sat does not hold for (P, N). Therefore, there exists a partition $(P', N') \neq (P, N)$ of $MA(\varphi)$ that does not satisfy any of the conditions (c1), (c2) and (c3) in the algorithm MBNF($\mathcal{K}$)-Sat. First, since condition (c2) does not hold, it follows that $(P', N') = Prt(\varphi, ob(P', N'), ob(P, N))$, which implies that for each modal atom of the form $A\psi$ of $MA(\varphi)$, $A\psi \in P$ iff $A\psi \in P'$; moreover, since condition (c3) does not hold, $ob(P, N) \wedge \neg ob(P', N')$ is $\mathcal{K}$-unsatisfiable, which implies that, if $B\psi \in P'$, then $B\psi \in P$. Finally, since $(P, N) \neq (P', N')$, it follows that there exists at least one modal atom of the form $B\psi$ such that $B\psi \in P$ and $B\psi \in N'$. Now let $P'' = P' - P_A$, $N'' = N' - N$: by the above considerations, it follows that (P', N') is a partition of $P - P_A$, and $N'' \neq \emptyset$. Furthermore, since $(P_A \cup P'', N \cup N'') = (P', N')$ and condition (c1) in the algorithm MBNF($\mathcal{K}$)-Sat does not hold, it follows that condition (c1) in the algorithm B-depth-1-Sat does not hold, and since $(P', N') = Prt(\varphi, ob(P', N'), ob(P, N))$, it follows that condition (c2) in the algorithm B-depth-1-Sat does not hold. Consequently, condition (c) in the algorithm B-depth-1-Sat does not hold for (P, N). □

We now prove that bounding the depth of nesting of the modality B to 1 has the effect of lowering the complexity of MBNF($\mathcal{K}$)-satisfiability in all the cases examined in which $\mathcal{K}$-satisfiability is NP-complete.

In the following, we denote with $\mathcal{L}'^1_M$ the restriction of the language $\mathcal{L}'_M$ to formulas of B-depth 1, i.e., $\mathcal{L}'^1_M = \mathcal{L}'_M \cap \mathcal{L}^1_M$.

Theorem 21. Let $\mathcal{K} \in \{\mathsf{K}_n, \mathsf{T}_n, \mathsf{S4}_n, \mathsf{KD45}_n, \mathsf{S5}_n\}$, let $\mathcal{L}' \subseteq \mathcal{L}$, and let $\mathcal{L}'$ be closed under boolean composition. If $\mathcal{K}$-satisfiability for formulas from $\mathcal{L}'$ is an NP-complete problem, then MBNF($\mathcal{K}$)-satisfiability for formulas from $\mathcal{L}'^1_M$ is Σ^p_2-complete.

Proof. Σ^p_2-hardness follows from the fact that, since $\mathcal{L}'$ is closed under boolean composition, $\mathcal{L}'^1_M$ is a superset of the extension of the so-called *flat* fragment of propositional MBNF, i.e., the extension of the propositional language with the modalities B and A and such that each propositional symbol lies within the scope of exactly one modality. For such a language, it has been proven in [21] that satisfiability in MBNF is Σ^p_2-complete. Therefore, MBNF($\mathcal{K}$)-satisfiability in $\mathcal{L}'^1_M$ is Σ^p_2-hard.

To prove membership in Σ^p_2, we analyze the complexity of the algorithm B-depth-1-Sat reported in Figure 1. In particular, observe that, given a partition (P, N), falsity of condition (c) can be decided in nondeterministic polynomial time. In fact, given a partition (P'', N''), verifying that both conditions (c1) and (c2) do not hold corresponds to decide that a linear number of $\mathcal{K}$-formulas from $\mathcal{L}'$ are $\mathcal{K}$-satisfiable, which can be decided in polynomial time after a single nondeterministic choice. Therefore, finding a partition (P'', N'') such that that both conditions (c1) and (c2) do not hold can be computed in nondeterministic polynomial time. Then, since the guess of the partition (P, N) of $MA(\varphi)$ requires a nondeterministic choice, and, as shown in the proof of Theorem 19, conditions (a) and (b) can be decided by solving a linear number of $\mathcal{K}$-satisfiability problems for formulas from $\mathcal{L}'$, it follows that the algorithm B-depth-1-Sat, if considered as a nondeterministic procedure, decides satisfiability of φ in nondeterministic polynomial time (with respect to the size of φ), using an NP-oracle. Therefore, from Theorem 20 we obtain an upper bound of Σ^p_2 for the satisfiability problem in MBNF($\mathcal{K}$). □

Informally, the different computational behavior of the algorithms MBNF($\mathcal{K}$)-

	$\mathcal{K} \in \{\mathsf{KD45}_1, \mathsf{S5}_1\}$	$\mathcal{K} \in \{\mathsf{KD45}_n, \mathsf{S5}_n\}$ $(n \geq 2)$	$\mathcal{K} \in \{\mathsf{K}_n, \mathsf{T}_n\}$ $(n \geq 1)$	$\mathcal{K} = \mathsf{S4}_n$ $(n \geq 1)$
B-depth $\leq 1 \wedge$ K-depth ≤ 1	Σ_2^p	Σ_2^p	Σ_2^p	Σ_2^p
B-depth $\leq 1 \wedge$ K-depth $\leq k$ $(k \geq 2)$	Σ_2^p	Σ_2^p	Σ_2^p	PSPACE
B-depth ≤ 1	Σ_2^p	PSPACE	PSPACE	PSPACE
K-depth ≤ 1	Σ_3^p	Σ_3^p	Σ_3^p	Σ_3^p
K-depth $\leq k$ $(k \geq 2)$	Σ_3^p	Σ_3^p	Σ_3^p	PSPACE
no restrictions	Σ_3^p	PSPACE	PSPACE	PSPACE

Figure 4. Complexity of MBNF($\mathcal{K}$)-satisfiability

Sat and B-depth-1-Sat is due to the presence of condition (c3) in the algorithm MBNF($\mathcal{K}$)-Sat: deciding falsity of such a condition requires to solve a coNP-complete problem, which implies that deciding falsity of condition (c) is Σ_2^p-hard. Conversely, without condition (c3) falsity of condition (c) can be decided in NP, due to the fact that, in both algorithms, falsity of conditions (c1) and (c2) can be decided in NP.

The above theorem and the results reported in the table of Figure 2 imply that, for each language $\mathcal{L}'$ and choice of the modal system $\mathcal{K}$ corresponding to an entry of such a table for which $\mathcal{K}$-satisfiability is NP-complete, MBNF($\mathcal{K}$)-satisfiability in the corresponding language $\mathcal{L}'^1_M$ with B-depth 1 is Σ_2^p-complete.

4.4. Summary of results

The computational results obtained in the two previous subsections are summarized in the table reported in Figure 4. The table must be read as follows:

for each row of the table, the complexity of checking satisfiability in MBNF($\mathcal{K}$) under the syntactic restriction reported in the first column (i.e., considering only the subset of MBNF($\mathcal{K}$)-formulas satisfying the restriction) is complete with respect to the complexity class appearing in the other columns, where each column corresponds to a different choice of the multimodal logic $\mathcal{K}$, both in the modal system and in the number of agents modeled.

As shown in the table, four major sources of complexity can be identified:

1. The first one is the number of agents allowed. As shown in the table, bounding the number of agents affects the worst-case complexity of the satisfiability task only for the modal systems KD45 and S5, and only when the bound is equal to 1. Indeed, for $\mathcal{K} = \mathsf{KD45}_1$ or $\mathcal{K} = \mathsf{S5}_1$, MBNF($\mathcal{K}$)-satisfiability is Σ_3^p-complete in the general case, thus it is computationally easier than for the other choices of $\mathcal{K}$. We remark that choosing $\mathcal{K} = \mathsf{KD45}_1$ (or $\mathcal{K} = \mathsf{S5}_1$) does not imply that MBNF($\mathcal{K}$) is only able to formalize a single agent: rather, it is possible in such a logic to represent the knowledge of an agent (through the modalities B and A) which is able to reason about another agent's beliefs (through the modality K_1).

2. The second restriction which affects complexity is the choice of the underlying multimodal system $\mathcal{K}$. As illustrated by the table, for $\mathcal{K} = \mathsf{S4}$ the complexity of MBNF($\mathcal{K}$)-satisfiability is harder than for the other choices of $\mathcal{K}$, when the K-depth is bounded to a value greater than 1.

3. The third complexity source is the K-depth: bounding the K-depth to any integer value generally lowers the worst-case complexity of the MBNF($\mathcal{K}$)-satisfiability problem, with the exception of the choice $\mathcal{K} = \mathsf{S4}_n$. In this case, it is necessary to bound the K-depth to 1 in order to affect the complexity.

4. Finally, bounding the B-depth to 1 has also an impact on the worst-case complexity of MBNF($\mathcal{K}$)-satisfiability, which can be seen by comparing the first three rows with the last three rows of the table. Notice that, conversely, it is possible to prove that bounding the nesting of the modality A does not affect the worst-case complexity of MBNF($\mathcal{K}$)-satisfiability in all the cases considered.

5. Application to RoboCup

In this section we apply our logical framework to the domain of robotic soccer, by experimenting the expressive abilities of MBNF($\mathcal{K}$) in the domain of RoboCup [1], the world championship for robotic soccer teams. To this purpose, we have analyzed the robotic architecture of the italian competitor in the RoboCup middle-size league, the ART team [18].[1] Such an architecture is shared by other teams participating in the RoboCup competition.

In the middle size league, each team is composed of four autonomous mobile robots, playing in a soccer field of about 8x4 meters. In the ART team, each robot has an internal representation (both at a symbolic and at a numerical level) of the environment. Each robot has some high-level planning ability, which uses the symbolic representation of the environment, and a set of low-level behaviors (i.e. control programs for executing simple actions) that are executed based on the numerical representation of the environment.

During the game, each robot receives data from his own sensors (video-camera, sonars, etc.) and communicates periodically part of such data to all the teammates. Then, each robot merges at the numerical level his own knowledge (i.e. information coming from his sensors) with the information coming from the other teammates, which may contradict the robot's own knowledge, and then updates his symbolic representation accordingly.

Such an approach has several limitations, which are mostly due to the fact that the only way for a robot to take into account the other agents' knowledge is through a quantitative fusion of information, while in many situations a qualitative approach appears in principle better suited to this application domain. In particular, we have modeled the information obtained by an agent's sensor as a *belief*, since such information may not correspond to the actual situation in the field, and have analyzed the possibility of reasoning, through nonmonotonic rules, about the other agents' beliefs (and ignorance).

To this purpose, we have exploited two representational features of MBNF($\mathcal{K}$):

1. the formal distinction done in MBNF($\mathcal{K}$) between the real world and the agents' beliefs about the world, which allows for easily representing rules for

[1] The ART team is constituted by a consortium among six italian universities. It obtained the second place in the middle-size league during the 1999 RoboCup edition.

deriving knowledge about the real world based on the data coming from the agents' sensors;

2. the possibility of representing nonmonotonic default rules in MBNF($\mathcal{K}$) by a combined usage of the two epistemic operators A and B.

Moreover, in order to properly model the information coming from the other agents as beliefs, we have chosen $\mathcal{K} = \mathsf{KD45}_n$ as the underlying modal system [12]. Therefore, we have studied MBNF($\mathsf{KD45_n}$) as a logical framework for reasoning in the RoboCup domain.

Specifically, in the following examples we represent the knowledge of a robot player who communicates with two teammates (for ease of exposition, we do not consider the fourth robot, who plays as the goalkeeper). Such a knowledge is formalized by means of a set of MBNF($\mathsf{KD45_n}$) formulas, in which:

- properties concerning the robot are not subscripted;
- properties concerning the robot's teammates (player 1 and player 2) are subscripted with indices 1 and 2;
- the modalities B and A are used to represent the robot's beliefs and assumptions;
- the modalities K_1 and K_2 are used to represent the beliefs of players 1 and 2, respectively. E.g., the formula $BK_1\varphi$ means that the robot believes that player 1 believes the property φ: such a belief derives from the information communicated to the robot by player 1;
- objective formulas represent properties that hold in the real world.

We have exploited such an approach in order to solve two crucial issues for the robotic soccer application: (i) "intelligent" fusion of sensing information coming from the robot's own sensors and from the other agents; (ii) deciding which role the robot has to play in the current situation, according to the team strategy [23].

The following simple examples have the purpose to show the use of MBNF($\mathcal{K}$) with respect to the above two issues. Hence, for the sake of simplicity, such examples do not make use of all the expressive abilities of MBNF($\mathcal{K}$): in particular, the examples only use literals (instead of general propositional formulas) within the scope of the modal operators K_1, K_2, and all $\mathcal{K}$-subformulas have K-depth 1.

5.1. Merging beliefs about the current situation

In the MBNF($\mathsf{KD45_n}$) framework we are able to represent different forms of qualitative fusion rules, according to the different nature of the sensing data to be merged. Below we report three different examples of such rules.

1. The first kind of rule is applied to those properties for which the robot can completely rely on his own beliefs, regardless of the other agents' beliefs. This corresponds to the case when the robot's sensing data relative to a certain property are the most reliable. As an example, if the robot believes (through his own sensor data) that he is in control of the ball, then he can conclude that he is actually in control of the ball, regardless of the information coming from the other players, since his sensors generally detect such a property more precisely than the other robots' sensors. We can formalize such a fusion rule through the following MBNF($\mathsf{KD45_n}$) formula:

$$B\,\textit{ball-control} \supset \textit{ball-control}$$

2. Then, there are some properties for which the robot can rely on his own sensors, unless the data coming from another agent contradicts such a piece of information. This case corresponds to the situation in which the different sensing data coming from different robots, relative to a certain property, have the same reliability. As an example, if the robot sees an opponent on player 1, then he can conclude there is actually an opponent on player 1, unless player 2 communicates that he sees no opponent on player 1.

$$B\,\textit{opponent-on}_1 \wedge \neg AK_2 \neg \textit{opponent-on}_1 \supset \textit{opponenton}_1$$

3. Finally, there are properties for which the robot's sensing data are the less reliable, therefore the robot should always believe his teammates. As an example, if player 1 or player 2 sees an opponent player on the robot, then the robot concludes there is actually an opponent on him, regardless of his own beliefs about such a property.

$$BK_1\,\textit{opponent-on} \vee BK_2\,\textit{opponent-on} \supset \textit{opponent-on}$$

5.2. Inferring the robot's role in the team

Another very useful application of epistemic reasoning in current robotic soccer implementations (and more generally, in multi-agent applications [23]) concerns the problem of deducing the role to assume within the team in a given situation. In the ART team, each robot runs a different control program (behavior) according to the role played in the team (goalkeeper, defender, forward). Each player is initially provided with a role, however such a role may change dynamically according to the situation in the field. Therefore, it is very important that a single robot is able to understand the "right" role he has to play in each situation according to the team strategy. Nonmonotonic reasoning about the other agents' knowledge appears well-suited to this purpose, which we now illustrate by a very simple example.

Example 22. Suppose the team plays with the following strategy: during a defensive action, two players must defend and one must stay in a forward position. Therefore, a cautious way for a robot to reasoning about such a situation is to assume a defensive role, unless he knows that both other players have already assumed a defensive role. Such a nonmonotonic rule can be expressed through the following MBNF($\mathsf{KD45_n}$) formula:

$$B\,\textit{defensive-sit} \wedge (\neg AK_1\,\textit{defender}_1 \vee \neg AK_2\,\textit{defender}_2) \supset B\,\textit{defender} \qquad (1)$$

For instance, suppose that, during a defensive action, the agent receives no information from agent 2, while agent 1 communicates to the agent that he is playing as a defender (i.e., $BK_1\,\textit{defender}_1$ holds). In this case, from (1) the agent concludes that he must assume a defensive role. In fact, all MBNF($\mathsf{KD45_n}$) models of the formula

$$B\,\textit{defensive-sit} \wedge BK_1\,\textit{defender}_1(1) \qquad (2)$$

are of the form (I, M, M) where M is the set of $\mathsf{KD45}_n$-interpretations satisfying the formula $\textit{defensive-sit} \wedge K_1\,\textit{defender}_1 \wedge \textit{defender}$, since it can be seen that each other structure of the form (J, M', M') satisfying the above formula is such that $(J, M', M') \not\models B\,\textit{defender}$, therefore from formula (1) it follows that (J, M', M') should satisfy $AK_2\,\textit{defender}_2$, but such an assumption is not "justified" by the formula (2), i.e., under the assumption that $AK_2\,\textit{defender}_2$ holds, it does not follow that $BK_2\,\textit{defender}_2$ is a consequence of (2). Conversely, if also agent 2 tells the agent that he is playing as a defender, then the agent does not conclude from (1) that he must assume a defensive role, since both $BK_1\,\textit{defender}_1$ and $BK_2\,\textit{defender}_2$ hold, which implies that $(\neg AK_1\,\textit{defender}_1 \vee \neg AK_2\,\textit{defender}_2)$ does not hold.

Therefore, when all robots adopt the strategy formalized by formula (1), they might all assume a defensive role, if they fail to communicate their conclusions to the teammates: however, this situation may be accepted and preferred to the case in which none or only one player assumes a defensive role. □

Another strategy that can be easily formalized in terms of an MBNF($\mathsf{KD45_n}$) formula is the following: if the robot knows that both other teammates have assumed a defensive role, then he can play in a forward position. Such a rule can be formalized as follows:

$$B\,\mathit{defensive\text{-}sit} \wedge BK_1\,\mathit{defender}_1 \wedge BK_2\,\mathit{defender}_2 \supset B\mathit{forward}$$

We conclude with an example of a strategy for an offensive situation, and reasoning in such a scenario.

Example 23. Suppose that, during offensive actions, we want two players to go forward while one player stays in a defensive position. Furthermore, we distinguish between the role of going forward while keeping control of the ball and going forward without carrying the ball (*forward-no-ball*), since such roles are implemented by two different behaviors. A player who is not in control of the ball must assume the "forward without ball" role, unless he knows that another player has already assumed such a role. This rule can be expressed through the following MBNF($\mathsf{KD45_n}$) formula:

$$\begin{array}{l} B\,\mathit{offensive\text{-}sit} \wedge B\neg \mathit{ball\text{-}control} \wedge \\ (\neg AK_1\mathit{forward\text{-}no\text{-}ball}_1 \vee \neg AK_2\mathit{forward\text{-}no\text{-}ball}_2) \supset B\mathit{forward\text{-}no\text{-}ball} \end{array} \quad (3)$$

Suppose now that the agent is not in control of the ball during an offensive action, i.e., both $B\mathit{offensive\text{-}sit}$ and $B\neg\mathit{ball\text{-}control}$ hold, and suppose that the agent receives no information from the other agents. In this case, the agent concludes that he must assume the "forward without ball" role. In fact, in a way analogous to the above described defensive situation, it can be seen that all MBNF($\mathsf{KD45_n}$) models of the formula

$$B\mathit{offensive\text{-}sit} \wedge B\neg\mathit{ball\text{-}control} \wedge (3) \quad (4)$$

are of the form (I, M, M) where M is the set of $\mathsf{KD45}_n$-interpretations satisfying the formula $\mathit{offensive\text{-}sit} \wedge \mathit{ball\text{-}control} \wedge \mathit{forward\text{-}no\text{-}ball}$, since each other structure satisfying the above formula should satisfy either $AK_1\mathit{forward\text{-}no\text{-}ball}_1$ or $AK_2\mathit{forward\text{-}no\text{-}ball}_2$, but such assumptions are not "justified" by formula (4). □

5.3. Other forms of epistemic reasoning

In addition, more sophisticated forms of (auto)epistemic reasoning can be realized within the MBNF($\mathsf{KD45_n}$) framework. For instance, suppose it is reasonable to assume that, during an offensive action, there is always at least one player and at most two players which are blocked by an opponent player. Suppose the robot sees one opponent on player 2 and no opponent on player 1; moreover, suppose player 2 has not communicated to the robot that there is an opponent on him. Then, the robot can conclude that there is an opponent player blocking him (if player 2 saw no opponent on the robot, he would have concluded that there is an opponent blocking him). Such a rule can be formalized by the following MBNF($\mathsf{KD45_n}$) formula:

$$B\neg \mathit{opponent\text{-}on}_1 \wedge B\,\mathit{opponent\text{-}on}_2 \wedge \neg A K_2 \mathit{opponenton}_2 \supset \mathit{opponent\text{-}on}$$

6. Related work

Several works in the recent literature have proposed modal approaches to the formalization of knowledge in a multi-agent setting. Among them, the most similar proposal to the one presented in this paper is [11,10], in which Levesque's logic of only knowing is extended to the case of many agents. Besides major technical aspects (the logic of only knowing is a monotonic logic while MBNF($\mathcal{K}$) is nonmonotonic), an important difference with our approach lies in the fact that the logic of only knowing does not allow for easily representing nonmonotonic rules and default rules with prerequisites, which are important tools for modeling the knowledge of agents, as shown in the previous section.

Moreover, recent proposals have defined modal logics with the aim of merging the beliefs of many agents [3]: however, such approaches are based on probabilistic frameworks, thus making it impossible a comparison with the logic MBNF($\mathcal{K}$).

Also, a great amount of work has been devoted to the definition of modal logics for reasoning about actions [7,2], in particular Dynamic Logics (see e.g., [22,4]): however, differently from our framework, such approaches use modalities for modeling the actions of agents rather than their epistemic abilities.

Finally, let us briefly comment on the relationship between the MBNF($\mathcal{K}$) framework and logic programming. Indeed, as illustrated in Section 2, the logic MBNF can be seen as a generalization of logic programming with negation as failure under the stable model/answer set semantics: actually, the main reason for the definition of MBNF was the logical reconstruction (and generalization) of logic programs with negation as failure. Therefore, the fragment of MBNF here presented, namely the MBNF($\mathcal{K}$) framework, can also be thought of as a generalization of a special class of logic programs, whose goals correspond to multimodal formulas (and negation-as-failure of such formulas). Consequently, we can think of exploiting the technology of logic programming under stable model semantics in order to reason about MBNF($\mathcal{K}$) theories, or, alternatively, to use current implementations of logic programming for reasoning in syntactically restricted MBNF($\mathcal{K}$) theories. However, such an issue is outside the scope of the present paper, whose main purpose is to illustrate the computational properties and the epistemological adequacy of the MBNF($\mathcal{K}$) framework.

7. Conclusions

In this paper we have proposed the logic MBNF($\mathcal{K}$) as a logical framework for representing the knowledge of an agent who reasons about its own beliefs and the beliefs of other agents. The main contributions of the paper are the following:

1. we have defined a theoretical framework for reasoning in multi-agent scenarios;
2. we have characterized the computational properties and defined reasoning algorithms for such a framework;
3. we have experimented the epistemological adequacy of our framework in a concrete multi-agent application (RoboCup);
4. we have started analyzing the computational issues related to the actual implementation of our framework in a real multi-robotic architecture. In particular, we have identified the main sources of the complexity of reasoning in MBNF($\mathcal{K}$), and, consequently, some syntactic restrictions which make the reasoning task computationally easier.

The above results show that the MBNF($\mathcal{K}$) framework appears well-suited for representing actual multi-agent settings. However, a lot of important questions still need to be addressed. In particular:

- on the one hand, the high worst-case complexity of reasoning in MBNF($\mathcal{K}$) arises the question of whether is it possible to arrive at a concrete, useful implementation of this framework on a real robotic architecture. On the other hand, we recall the very good experimental results of the most recent implementations of satisfiability solvers for multimodal logics, which already have to deal with PSPACE-hard problems. Therefore, it would be very interesting to construct a solver for MBNF($\mathcal{K}$)-satisfiability on top of an already existing system for multimodal logics, and test its performance e.g. in the RoboCup domain;

- the issue of belief revision in MBNF($\mathcal{K}$) appears also very important for the application of this framework to multi-agent scenarios. E.g., in the RoboCup application described in the previous section, it could be interesting to apply belief revision when new information comes from the sensors (or from other agents), instead of executing a cycle in which, at each iteration, first a new knowledge base is created from the current sensor data and the information communicated by the other agents, then all conclusions are recomputed with respect to such a knowledge base;
- an important aspect of reasoning about knowledge that is missing in the MBNF($\mathcal{K}$) framework is the so-called *common knowledge* [5,13], which allows for sophisticated forms of epistemic inference. Therefore, it would be interesting to add common knowledge to MBNF($\mathcal{K}$) and study the epistemological and computational properties of such an augmented framework.

References

[1] M. Asada, H. Kitano, I. Noda, and M. Veloso. Robocup: Today and tomorrow – what we have learned. *Artificial Intelligence*, 110(2):193–214, 1999.

[2] M. Baldoni, L. Giordano, A. Martelli, and V. Patti. Modeling agents in a logic action language. In *Proc. of Workshop on Practical Reasoning Agents at FAPR2000*, 2000.

[3] L. Boldrin and A. Saffiotti. A modal logic for fusing partial belief of multiple reasoners. *J. of Logic and Computation*, 9(1), 1999.

[4] G. De Giacomo and M. Lenzerini. PDL-based framework for reasoning about actions. In *Proc. of the 4th Conf. of the Ital. Assoc. for Artificial Intelligence (AI*IA'95)*, volume 992 of *Lecture Notes in Artificial Intelligence*, pages 103–114. Springer-Verlag, 1995.

[5] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge.* The MIT Press, 1995.

[6] M. L. Ginsberg, editor. *Readings in Nonmonotonic Reasoning.* Morgan Kaufmann, Los Altos, 1987.

[7] L. Giordano, A. Martelli, and C. Schwind. Dealing with concurrent actions in modal action logic. In *Proc. of the 13th Eur. Conf. on Artificial Intelligence (ECAI'98)*, pages 537–541, 1998.

[8] F. Giunchiglia and R. Sebastiani. Building decision procedures for modal logics from propositional decision procedure – the case study of modal K. In *Proc. of the 13th Int. Conf. on Automated Deduction (CADE'96)*, pages 583–597, 1996.

[9] J. Y. Halpern. The effect of bounding the number of primitive propositions and the depth of nesting on the complexity of modal logic. *Artificial Intelligence*, 75(2):361–372, 1995.

[10] J. Y. Halpern. A theory of knowledge and ignorance for many agents. *J. of Logic and Computation*, 7(1):79–108, 1997.

[11] J. Y. Halpern and G. Lakemeyer. Multi-agent only knowing. In *Proc. of the 6th Conf. on Theoretical Aspects of Reasoning about Knowledge (TARK'96)*. Morgan Kaufmann, 1996.

[12] J. Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54:319–379, 1992.

[13] J. Y. Halpern and R. Shore. Reasoning about common knowledge with infinitely many agents. In *Proc. of the 14th IEEE Symp. on Logic in Computer Science (LICS'99)*, pages 384–393, 1999.

[14] I. Horrocks. Using an expressive description logic: FaCT or fiction? In *Proc. of the 6th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'98)*, pages 636–647, 1998.

[15] D. S. Johnson. A catalog of complexity classes. In *Handbook of Theoretical Computer Science*, volume A, chapter 2. Elsevier Science Publishers (North-Holland), Amsterdam, 1990.

[16] V. Lifschitz. Minimal belief and negation as failure. *Artificial Intelligence*, 70:53–72, 1994.

[17] F. Lin and Y. Shoham. Epistemic semantics for fixed-point non-monotonic logics. *Artificial Intelligence*, 57:271–289, 1992.

[18] D. Nardi, G. Adorni, A. Bonarini, R. Cassinis, G. Clemente, A. Chella, E. Pagello, and M. Piaggio. ART: Azzurra Robot Team. Available at `http://www.ida.liu.se/ext/RoboCup-99`, 1999.

[19] C. H. Papadimitriou. *Computational Complexity.* Addison Wesley Publ. Co., Reading, Massachussetts, 1994.

[20] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980. Republished in [6].

[21] R. Rosati. Reasoning about minimal belief and negation as failure. *J. of Artificial Intelligence Research*, 11:277–300, 1999.

[22] S. J. Rosenschein. Plan synthesis: A logical perspective. In *Proc. of the 7th Int. Joint Conf. on Artificial Intelligence (IJCAI'81)*, pages 331–337, 1981.

[23] P. Stone and M. Veloso. Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *Artificial Intelligence*, 110(2):241–273, 1999.