

Multi-modal nonmonotonic logics of minimal knowledge

Riccardo Rosati

Dipartimento di Informatica e Sistemistica
Università di Roma “La Sapienza”
Via Salaria 113, 00198 Roma, Italy
rosati@dis.uniroma1.it

Abstract

In this paper we introduce multi-modal logics of minimal knowledge. Such a family of logics constitutes the first proposal in the field of epistemic nonmonotonic logic in which the three following aspects are simultaneously addressed: (i) the possibility of formalizing multiple agents through multiple modal operators; (ii) the possibility of using first-order quantification in the modal language; (iii) the possibility of formalizing nonmonotonic reasoning abilities for the agents modeled, based on the principle of minimal knowledge. We illustrate the expressive capabilities of multi-modal logics of minimal knowledge to provide a formal semantics to peer-to-peer data integration systems, which constitute one of the most recent and complex architectures for distributed information systems.

1 Introduction

1.1 Nonmonotonic modal logics

Research in the formalization of commonsense reasoning has pointed out the need of formalizing agents able to reason introspectively about their own knowledge and ignorance [34, 24]. Modal epistemic logics have thus been proposed, in which modalities are interpreted in terms of knowledge or belief. Generally speaking, the conclusions an introspective agent is able to draw depend on both what she knows and what she *does not* know. Hence, any such conclusion may be retracted when new facts are added to the agent’s knowledge. For this reason, many *nonmonotonic* modal formalisms have been proposed in order to characterize the reasoning abilities of an introspective agent: among them, we recall the nonmonotonic modal logics originally proposed by McDermott and Doyle [32, 31, 30], Moore’s autoepistemic logic [35], Lifschitz’s logic of minimal knowledge and negation

as failure MKNF/MBNF [26, 27], Levesque’s logic of only knowing [24], and ground nonmonotonic modal logics [18, 9, 46].

1.2 Modal logics of minimal knowledge

Among the nonmonotonic modal logics proposed in the literature, some are based on the so-called principle of *minimal knowledge*. In particular, Halpern and Moses in [18] defined an epistemic logic, based on the modal system S5, for modeling knowledge and ignorance of processes in a distributed computer system, which is based on a very intuitive semantics: consider only the models of the knowledge base (i.e. the epistemic states of the agent modeled) in which knowledge is minimal (i.e. the ignorance of the agent is maximal). Hence, this formalism is also known as logic of minimal epistemic states, and constitutes the basis of several nonmonotonic modal formalisms proposed in the literature, among which [26, 29, 33, 44]. In particular, Lifschitz [26, 27] has proposed a bimodal logic, known as MKNF,¹ combining the minimal knowledge paradigm with the notion of negation as failure in logic programming.

Notably, it was shown that the logic of minimal epistemic states can be given a fixpoint characterization [46] which actually defines a whole family of logics of minimal knowledge states, the so-called *ground nonmonotonic modal logics* [47, 40], obtained by varying the underlying modal system. Hence, such a family of logics can be considered as obtained through a generalization of the notion of minimal knowledge, according to the different modal system chosen.

MKNF has been used in order to give a declarative semantics to very general classes of logic programs [28, 42, 20], which generalize the stable model semantics of negation as failure in logic programming [13, 14, 15]. Also, MKNF can be viewed as an extension of the theory of epistemic queries to databases [37], which deals with the problem of querying a first-order database about its own knowledge. Due to its ability of expressing many features of nonmonotonic logics [27, 42, 38], MKNF is generally considered as a unifying framework for several nonmonotonic formalisms, including default logic, autoepistemic logic, circumscription, epistemic queries, and logic programming.

1.3 Limitations of current proposals

Let us now point out two limitations of the research in nonmonotonic modal logics:

- The vast majority of the studies in nonmonotonic modal logics in the literature deal with *propositional* modal logics, while there are very

¹Actually, Lifschitz in [26] defined the logic MKNF, while in [27] he presented the logic MBNF, which slightly differs from MKNF.

few proposals concerning nonmonotonic extensions of first-order logic (among them, it is worth citing [21, 22] which study first-order extensions of nonmonotonic modal logics).

- Almost all the modal approaches to nonmonotonic logic use a single modality, i.e., they lack the ability of expressing the knowledge of many agents. In particular, none of the nonmonotonic formalisms based on the principle of minimal knowledge which have been proposed in the literature is *multi-modal*, i.e., is able to express the different epistemic states of a set of agents. The only exceptions are: (i) the work reported in [17], which actually deals with minimal knowledge in a multi-agent setting. However, [17] only deals with a modal propositional language (we refer to Section 3.4 for a more detailed analysis of [17]); (ii) the approach presented in [39], which is able (in a propositional setting) to formalize a multi-agent scenario, but restricts the nonmonotonic abilities (based on the minimal knowledge principle) only to a *single* agent. Conversely, multi-modal languages have been extensively studied in *monotonic* systems for knowledge and belief, in particular $K_n, T_n, S4_n, K45_n, KD45_n, S5_n$ (see [19, 10]).

On the other hand, recent developments in the field of distributed information systems have outlined the need for a multi-modal, first-order, nonmonotonic logic. In particular, several recent studies in the formalization of peer-to-peer distributed systems [2, 5, 12, 7, 6] have clearly pointed out that the intended semantics of information in this kind of applications is naturally captured by an epistemic logic approach based on the principle of minimal knowledge, in which each system is modeled as an autonomous epistemic agent, and the exchange of information in the system is represented by epistemic sentences that express the relationships among the epistemic states of the different agents. Moreover, such studies highlight that, in order to fully represent the peculiar aspects of such application scenarios, we need a logic incorporating first-order quantification that is also able to formalize typical nonmonotonic reasoning features of the epistemic agents [6].

1.4 Our contribution

The aim of this paper is to provide a first proposal of multi-modal nonmonotonic logics. In particular, we define the family of logics \mathcal{S}_n^A , which has the following characteristics:

- each logic in the family \mathcal{S}_n^A is a nonmonotonic logic based on the principle of minimal knowledge, in particular it can be viewed as a generalization of Lifschitz’s logic MKNF;
- each such logic is a multi-modal logic, since it can be viewed as a nonmonotonic extension of the family of multi-modal logics \mathcal{S}_n [19, 10];

- each such logic is a modal first-order logic, i.e., it allows for first-order quantification.

In particular, we point out that the semantic definition of \mathcal{S}_n^A is based on a preference order on possible-world structures, following the studies on a model theory for nonmonotonic modal logics [44, 41, 43, 9]. Also, it can be seen as a generalization of the possible-world semantics of MKNF and MBNF [26, 27].

Then, to show an example of the representational abilities of the logic \mathcal{S}_n^A , we use the logic $\mathbf{K45}_n^A$ to formalize the behaviour of knowledge in distributed, peer-to-peer information systems. For our purposes, this kind of application is of particular interest, since it requires all the three main ingredients of the logics \mathcal{S}_n^A , namely, multiple modal operators, first-order quantification, and nonmonotonic abilities.

1.5 Structure of the paper

In the next section, we recall standard (monotonic) multi-modal logics. Then, we define syntax and semantics of the nonmonotonic multi-modal logics \mathcal{S}_n^A , and analyze the relationship between the family \mathcal{S}_n^A and nonmonotonic modal logics previously defined. In the subsequent section we illustrate the representation abilities of one of these logics ($\mathbf{K45}_n^A$) in the formalization of knowledge in distributed, peer-to-peer information systems. Finally, we draw some conclusions.

2 Multi-modal logics

In this section we recall multi-modal epistemic logics [19, 10]. We assume that the reader is familiar with the basics of modal logic [3].

The language \mathcal{L}_k is the usual function-free first-order multi-modal language, i.e., it is obtained from a function-free first-order logic language \mathcal{L} by adding a set $\mathbf{K}_1, \dots, \mathbf{K}_n$ of modal operators, involving the formation rule: if ϕ is a (possibly open) formula, then also $\mathbf{K}_i\phi$ is so, for $1 \leq i \leq n$ for a fixed n . Formulas without occurrences of \mathbf{K}_i are called *objective formulas*.

We use ψ_c^x to denote the formula obtained from ψ by substituting each free occurrence of the variable x with the constant c .

To define the semantics, we start from first-order interpretations. In particular, we restrict our attention to first-order interpretations that share a fixed infinite domain Δ . We further assume that for each domain element $d \in \Delta$, we have a unique constant $c_d \in \Gamma$ that denotes exactly d , and, vice versa, that every constant $c_d \in \Gamma$ denotes exactly one domain element $d \in \Delta$. In other words, the constants in Γ act as *standard names* [25]. In the following, we denote by $\mathcal{I}_{\mathcal{L}}$ the set of all first-order interpretations for \mathcal{L} .

Formulas of \mathcal{L}_k are interpreted over \mathcal{S}_n -structures. Given a modal system \mathcal{S} , where $\mathcal{S} \in \{\mathbf{K}, \mathbf{T}, \mathbf{K4}, \mathbf{K45}, \mathbf{KD45}, \mathbf{S4}, \mathbf{S5}\}$, a \mathcal{S}_n -structure is a Kripke structure E of the form $(W, (R_1, \dots, R_n), V)$, where:

- W is a set whose elements are called *possible worlds*;
- V is the world interpretation function, i.e., a function assigning to each $w \in W$ a first-order interpretation $V(w)$;
- each R_i , called the *accessibility relation* for the modality \mathbf{K}_i , is a binary relation over W that satisfies the conditions for the modal system \mathcal{S}_n described below.

Different multi-modal logics are obtained by imposing different conditions that each accessibility relation R_i has to satisfy: in particular,

- when $\mathcal{S}_n = \mathbf{T}_n$, each R_i is reflexive;
- when $\mathcal{S}_n = \mathbf{K4}_n$, each R_i is transitive;
- when $\mathcal{S}_n = \mathbf{K45}_n$, each R_i is transitive and euclidean;
- when $\mathcal{S}_n = \mathbf{KD45}_n$, each R_i is serial, transitive and euclidean;
- when $\mathcal{S}_n = \mathbf{S4}_n$, each R_i is reflexive and transitive;
- when $\mathcal{S}_n = \mathbf{S5}_n$, each R_i is reflexive, transitive and euclidean.

It is well-known that the above four conditions on the accessibility relation (serial, reflexive, transitive, euclidean) of \mathcal{S}_n -structures correspond respectively to the validity of the following axiom schemas:

- *axiom schema D*: $\mathbf{K}_i\phi \supset \neg\mathbf{K}_i\neg\phi$
- *axiom schema T*: $\mathbf{K}_i\phi \supset \phi$
- *axiom schema 4*: $\mathbf{K}_i\phi \supset \mathbf{K}_i\mathbf{K}_i\phi$
- *axiom schema 5*: $\neg\mathbf{K}_i\phi \supset \mathbf{K}_i\neg\mathbf{K}_i\phi$

A \mathcal{S}_n -interpretation is a pair E, w , where $E = (W, (R_1, \dots, R_n), V)$ is a \mathcal{S}_n -structure, and w is a world in W .

The truth of a sentence (i.e., a closed formula) ϕ in an interpretation E, w , written $E, w \models \phi$, is defined inductively as follows:

- $E, w \models P(c_1, \dots, c_n)$ if $V(w) \models P(c_1, \dots, c_n)$
- $E, w \models \phi_1 \wedge \phi_2$ if $E, w \models \phi_1$ and $E, w \models \phi_2$
- $E, w \models \neg\phi$ if $E, w \not\models \phi$

- $E, w \models \exists x. \psi$ if $E, w \models \psi_c^x$ for some constant c
- $E, w \models \mathbf{K}_i \phi$ if $E, w' \models \phi$ for every w' such that $(w, w') \in R_i$

A \mathcal{S}_n -model for ϕ is a \mathcal{S}_n -interpretation E, w such that $E, w \models \phi$.

We say that a sentence ϕ is \mathcal{S}_n -satisfiable if there exists a \mathcal{S}_n -model for ϕ , *unsatisfiable* otherwise. A \mathcal{S}_n -model for a set Σ of sentences is a \mathcal{S}_n -model for every sentence in Σ . A sentence ϕ is \mathcal{S}_n -entailed by a set Σ of sentences, written $\Sigma \models_{\mathcal{S}_n} \phi$, if and only if $E, w \models \phi$ in every \mathcal{S}_n -model E, w of Σ .

3 Multi-modal logics of minimal knowledge and negation as failure

In this section we introduce a nonmonotonic extension of the multi-modal logics recalled in the previous section. Informally, such an extension is obtained by adding a new set of modal operators $\mathbf{A}_1, \dots, \mathbf{A}_n$ to the modal language. Then, following (and generalizing) the semantic construction of the logic MKNF [26], the modal operators $\mathbf{K}_1, \dots, \mathbf{K}_n$ are interpreted as epistemic operators of minimal knowledge, and the modal operators $\mathbf{A}_1, \dots, \mathbf{A}_n$ are interpreted as epistemic operators of *justified assumption* [29], which corresponds to the well-known notion of *negation as failure* [27].

3.1 m -normal Kripke structures

We now restrict the above semantics for \mathcal{S}_n by fixing both the set of worlds W and the world interpretation function V . This restriction is introduced for technical reasons, in order to allow for a well-founded definition of a preference order between structures, which will be introduced in Section 3.3.

However, such a restriction does not affect the semantics of \mathcal{S}_n (with respect to satisfiability of a formula) for each modal system \mathcal{S} studied in this paper.

Given a formula φ of length m , let $k = (n \cdot m)^m$. We define the following set of worlds \mathcal{W}_m :

$$\mathcal{W}_m = \{(I, 1), \dots, (I, k) \mid I \in \mathcal{I}_{\mathcal{L}}\}$$

That is, \mathcal{W}_m contains k distinct elements for each FOL interpretation I .

Moreover, we define the following world interpretation function $V_m : \mathcal{W}_m \rightarrow \mathcal{I}_{\mathcal{L}}$:

$$\text{for each } w = (I, j) \in \mathcal{W}_m, \quad V_m(w) = I$$

Namely, the interpretation that V_m associates with a world (I, j) in \mathcal{W}_m is I .

Definition 1 *An m -normal \mathcal{S}_n -structure is a \mathcal{S}_n -structure whose set of worlds is \mathcal{W}_m and whose world interpretation function is V_m . Moreover, an \mathcal{S}_n -interpretation E, w is m -normal iff the \mathcal{S}_n -structure E is m -normal.*

Notice that, from the above definition, it follows that *all* m -normal \mathcal{S}_n -structures are defined over the *same* set of worlds \mathcal{W}_m and the *same* world interpretation function V_m .

It is known (and easy to see) that, for a formula φ of length less or equal to m , φ is \mathcal{S}_n -satisfiable iff φ is true in an m -normal \mathcal{S}_n -interpretation, for every logic \mathcal{S}_n such that $\mathcal{S}_n \in \{K_n, T_n, K4_n, K45_n, KD45_n, S4_n, S5_n\}$ (see e.g. [30] where this property is shown for the single-agent case). In other words, restricting to the set of worlds \mathcal{W}_m and interpreting \mathcal{W}_m according to V_m does not change satisfiability of formulas of length less or equal to m . Consequently, from now on we restrict our attention to m -normal interpretations only.

3.2 Adding modal operators of negation as failure

First, we introduce the language \mathcal{L}_k^A , which is an extension of \mathcal{L}_k obtained by adding to the first-order modal language a new set of modal operators, $\mathbf{A}_1, \dots, \mathbf{A}_n$.

The semantics of \mathcal{L}_k^A sentences is formally defined as follows. An m -normal \mathcal{S}_n^A -structure E is a tuple $(\mathcal{W}_m, (R_1, \dots, R_n, R_1^a, \dots, R_n^a), V_m)$, where:

- \mathcal{W}_m and V_m are defined as in Section 3.1;
- each R_i and each R_i^a are binary relations over W satisfying the conditions imposed by the system \mathcal{S}_n (described in the previous section).

Therefore, with respect to \mathcal{S}_n -structures, \mathcal{S}_n^A -structures have the additional accessibility relations R_1^a, \dots, R_n^a . Such relations account for the additional modal operators $\mathbf{A}_1, \dots, \mathbf{A}_n$.

An m -normal \mathcal{S}_n^A -interpretation is a pair E, w where $w \in \mathcal{W}_m$.

The notion of truth of a \mathcal{L}_k^A sentence in a world of a \mathcal{S}_n^A -interpretation is analogous to the notion given in Section 2 for \mathcal{L}_k , with the addition of the following rule:

- $E, w \models \mathbf{A}_i \phi$ iff $E, w' \models \phi$ for each w' such that $(w, w') \in R_i^a$

3.3 Nonmonotonic semantics

So far, the family of logics \mathcal{S}_n^A do not appear as a significant extension of the logics \mathcal{S}_n : indeed, according to the above notion of truth, the new modal operators \mathbf{A}_i are treated just like any \mathbf{K}_i operator in \mathcal{S}_n , so there is no apparent reason to distinguish the \mathbf{A}_i 's operators from the \mathbf{K}_i 's.

Actually, for each logic \mathcal{S}_n^A , the different (nonmonotonic) meaning of the two sets of modal operators in \mathcal{S}_n^A with respect to \mathcal{S}_n is due to the following notion of \mathcal{S}_n^A -model for a sentence ϕ , which differs from the (classical) notion of \mathcal{S}_n -model, and is obtained by imposing a preference order over \mathcal{S}_n^A -structures satisfying ϕ .

The next definition defines a relation \leq_K between m -normal \mathcal{S}_n^A -interpretations which agree on the initial world and on the accessibility relations R_i^a 's.

Definition 2 Let $E = (\mathcal{W}_m, (R_1, \dots, R_n, R_1^a, \dots, R_n^a), V_m)$ and $E' = (\mathcal{W}_m, (R'_1, \dots, R'_n, R_1^a, \dots, R_n^a), V_m)$ be m -normal \mathcal{S}_n^A -structures. We say that E is K -contained in E' (denoted by $E \leq_K E'$) if, for each $i \in \{1, \dots, n\}$, $R_i \subseteq R'_i$.

Intuitively, if E is K -contained in E' , then E' has less (or equal) non-introspective knowledge with respect to the modal operators \mathbf{K}_i than E , since adding possible worlds (by adding pairs of worlds to the R_i 's accessibility relations) reduces the non-introspective knowledge represented by the accessibility relations interpreting the \mathbf{K}_i 's operators.

For instance, it can be immediately verified that, if E is K -contained in E' , then, for each first-order sentence ϕ and for each $w \in W$, if $E', w \models \mathbf{K}_i \phi$ then $E, w \models \mathbf{K}_i \phi$, but not necessarily vice-versa.

We now prove that the relation \leq_K between \mathcal{S}_n^A -structures is well-defined, since it constitutes a partial order.

Theorem 3 The relation \leq_K between \mathcal{S}_n^A -structures constitutes a partial order.

Proof. It is immediate to see that, from Definition 1 and Definition 2, reflexivity, antisymmetry and transitivity of \leq_K hold. Consequently, \leq_K is a partial order. \square

Definition 4 Let $\phi \in \mathcal{L}_k^A$ be a formula of size m , let $E = (\mathcal{W}_m, (R_1, \dots, R_n, R_1^a, \dots, R_n^a), V_m)$ be a m -normal \mathcal{S}_n^A -structure, and let $w \in \mathcal{W}_m$. The m -normal \mathcal{S}_n^A -interpretation E, w is a \mathcal{S}_n^A -model for ϕ if the following conditions hold:

- $E, w \models \phi$;
- $R_i = R_i^a$ for each $i \in \{1, \dots, n\}$;
- there exists no m -normal \mathcal{S}_n^A -structure E' such that $E' \neq E$, $E', w \models \phi$, and $E \leq_K E'$.

A \mathcal{S}_n^A -model for a set Σ of sentences is a \mathcal{S}_n^A -model for every sentence in Σ . A sentence ϕ is \mathcal{S}_n^A -entailed by a set Σ of sentences, written $\Sigma \models_{\mathcal{S}_n^A} \phi$, if and only if $E, w \models \phi$ in every \mathcal{S}_n^A -model E, w of Σ .

Let us now try to provide an intuition for the semantics of the logics in the family \mathcal{S}_n^A . The above semantics formalizes the idea of selecting \mathcal{S}_n^A -structures that satisfy two intuitive principles:

1. *knowledge is minimal*, which is realized through the notion of preference between structures;
2. *assumptions are justified by knowledge*, which is realized by the fact that, for each i , the meaning of the operators \mathbf{A}_i and \mathbf{K}_i is the same, since $R_i = R_i^a$.

Such semantic principles of minimal knowledge and justified assumptions are well-known in nonmonotonic reasoning [29, 27, 38]. In particular, we recall that the principle of justified assumption exactly corresponds to the semantics of the modal operator in Moore’s autoepistemic logic [38]. Moreover, as illustrated in [26, 27, 29], the justified assumption operator exactly formalizes the notion of *negation as failure* in logic programming under the stable model semantics.

To better understand the differences between the operators \mathbf{K}_i and \mathbf{A}_i , we show that the two modalities are not equivalent. In particular, suppose that φ is an objective sentence, i.e., a sentence without occurrences of the modal operators. We now prove that adding the formula $\mathbf{K}_i \equiv \mathbf{A}_i$ to a theory T actually changes the set of \mathcal{S}_n^A -models of T . Since $\mathbf{K}_i \equiv \mathbf{A}_i$ corresponds to the conjunction of the two formulas $\mathbf{K}_i\varphi \supset \mathbf{A}_i\varphi$ and $\mathbf{A}_i\varphi \supset \mathbf{K}_i\varphi$, we consider such two formulas:

- first, let ξ be the formula $\mathbf{K}_i\varphi \supset \mathbf{A}_i\varphi$. Given any theory T , it is easy to see that the set of \mathcal{S}_n^A -models of T and the set of \mathcal{S}_n^A -models of $T \cup \{\xi\}$ coincide;
- conversely, we now show that the formula $\mathbf{A}_i\varphi \supset \mathbf{K}_i\varphi$ in general does not preserve the set of \mathcal{S}_n^A -models. For ease of exposition, consider the logic $\mathbf{S5}_n^A$ and the single-agent case: the only $\mathbf{S5}_n^A$ -models of the empty theory are interpretations of the form E, w where E is the structure in which both R_1 and R_1^a are the total relation on \mathcal{W}_m . Conversely, the set of $\mathbf{S5}_n^A$ -models of $\mathbf{A}_i\varphi \supset \mathbf{K}_i\varphi$ also contains all the interpretations of the form E', w where E' is such that both R_1 and R_1^a are the set of all pairs (w_1, w_2) such that w_1 is any world in \mathcal{W}_m and w_2 is any world in \mathcal{W}_m such that the FOL interpretation $V_m(w_2)$ satisfies φ .

From the above proof, it follows that adding the formula $\mathbf{K}_i \equiv \mathbf{A}_i$ to a theory T in general changes the set of \mathcal{S}_n^A -models of T .

Remark. From the technical viewpoint, the above preference semantics for the logics \mathcal{S}_n^A is a non-trivial extension of analogous semantic constructions underlying other nonmonotonic modal logics. The main difference with respect to such previous constructions is that here, due to the presence of multiple modal operators, we cannot impose the condition that the preferred models of a theory always correspond to structures in which each accessibility relation is total (which has a syntactic counterpart in the so-called

stable sets of modal formulas [45]). Consequently, minimality of knowledge in the preferred models is imposed via a different, although simple, condition (formally stated by Definition 2), which can be seen as a generalization of analogous minimality criteria in previous, simpler nonmonotonic modal formalisms like MKNF [26] or ground nonmonotonic modal logics [9].

3.4 The logics \mathcal{S}_n^A vs. nonmonotonic modal logics

We now analyze more in detail the relationship between the family of logics \mathcal{S}_n^A and previous nonmonotonic modal logics. In particular, we want to point out the following deep correspondences between the logics \mathcal{S}_n^A and some well-known nonmonotonic modal logics:

- *Correspondence between MKNF and $S5_n^A$* . First, we analyze the relationship between the logic $S5_n^A$ and Lifschitz’s logic MKNF. More precisely, we start by recalling that the language of MKNF makes use of two modal operators K and *not*. Now, given an MKNF theory Σ , it can be proved that the MKNF-models of Σ coincide with the $S5_n^A$ -models of the theory Σ' obtained from Σ by replacing each occurrence of the modal operator K with the modality \mathbf{K}_1 , and replacing each occurrence of the modal operator *not* with the modality $\neg\mathbf{A}_1$. Therefore, the logic $S5_n^A$ can be viewed as the multi-modal generalization of MKNF, and, more generally, the whole family of logics \mathcal{S}_n^A can be seen as a generalization of the semantic construction underlying the logic MKNF.
- *Correspondence between autoepistemic logic and $S5_n^A$* . As a consequence of the previous correspondence, and since in turn MKNF constitutes a generalization of Moore’s autoepistemic logic [38], it follows that an analogous precise correspondence holds between the logic $S5_n^A$ and Moore’s autoepistemic logic, which allows us to also interpret $S5_n^A$ as a multi-modal generalization of Moore’s autoepistemic logic.
- *Correspondence between ground logic $S5_G$ and $S5_n^A$* . The family of *ground nonmonotonic modal logics* studied in [46, 47, 40, 9] is also deeply related to the logics \mathcal{S}_n^A . More precisely, it can be shown that the ground nonmonotonic modal logic based on the modal system $S5$ and known as $S5_G$ [9] corresponds to the logic $S5_n^A$, in the sense that, given a unimodal theory Σ , the $S5_G$ -models of Σ coincide with the $S5_n^A$ -models of the theory Σ' obtained from Σ by replacing each occurrence of the modal operator with the modality \mathbf{K}_1 .
- *Relationship between \mathcal{S}_n^A and the nonmonotonic multi-modal logics of [17]*. Finally, the work presented in [17] is also related to the logics \mathcal{S}_n^A . In fact, [17] presents a nonmonotonic variant of the multi-modal

propositional logics $K_n, T_n, S4_n, K45_n, KD45_n, S5_n$ based on the minimal knowledge principle. Thus, it is interesting to compare this approach to the logics \mathcal{S}_n^A . First, we observe that the two approaches present the following syntactic differences: (i) the language in [17] only considers the modal propositional language, while our approach considers its extension to the first-order language; (ii) [17] only uses one set of modal operators \mathbf{K}_i , while the logics \mathcal{S}_n^A also use the operators \mathbf{A}_i . Then, if we compare the approach of [17] with the restriction of \mathcal{S}_n^A obtained by disallowing the use of the modalities \mathbf{A}_i and the first-order quantifiers, the crucial difference lies in the definition underlying the nonmonotonic semantics: the nonmonotonic construction in [17] is based on a maximality criterion based on the notion of *possibility*, which changes for the different modal systems, while our maximality criterion, i.e., maximization of the accessibility relations R_i , is independent of the underlying modal system. Of course, such a difference in the semantic construction gives rise to generally different nonmonotonic logics.

4 Modeling knowledge in a P2P system

In this section we show the representational abilities of the multi-modal logics \mathcal{S}_n^A . In particular, we use one of such logics, $K45_n^A$, to provide a formal semantics for *peer-to-peer (P2P) data integration systems*, which constitute one of the most recent and complex architectures in the field of distributed information systems.

For a detailed introduction to P2P data integration systems, we refer the reader to [16], and for more details on the formalization presented in this section, we refer to [6]. In the following, we assume that the reader is familiar with the basics of relational database theory [1].

4.1 P2P data integration systems

We refer to a fixed, infinite, denumerable set Γ of constants. Such constants are shared by all peers, and denote the data items managed by the P2P data integration system (denoted by P2PDIS in the following). Moreover, given a relational alphabet A , we denote with \mathcal{L}_A the set of function-free first-order logic (FOL) formulas whose relation symbols are in A and whose constants are in Γ .

A *P2P data integration system* $\mathcal{P} = \{P_1, \dots, P_n\}$ is constituted by a set of n peers. Each peer $P_i \in \mathcal{P}$ (cf. [16]) is defined as a tuple $P_i = (id, G, S, L, M, \mathcal{L})$, where:

- id is a symbol that identifies the peer P_i within \mathcal{P} , called the identifier of P_i .

- G is the *schema* of P_i , which is a finite set of formulas of \mathcal{L}_{A_G} (representing local integrity constraints), where A_G is a relational alphabet (disjoint from the other alphabets in \mathcal{P}) called the *alphabet* of P_i . We assume that the language \mathcal{L}_{A_G} of peer P_i includes the special sentence \perp_i that is false in every interpretation for \mathcal{L}_{A_G} . Intuitively, the peer schema provides an intensional view of the information managed by the peer.
- S is the (*local*) *source schema* of P_i , which is simply a finite relational alphabet (again disjoint from the other alphabets in \mathcal{P}), called the *local alphabet* of P_i . Intuitively, the source schema describes the structure of the data sources of the peer (possibly obtained by wrapping physical sources), i.e., the sources where the real data managed by the peer are stored.
- L is a set of (*local*) *mapping assertions* between G and S . Each local mapping assertion is an expression of the form $cq_S \rightsquigarrow cq_G$, where cq_S and cq_G are two conjunctive queries of the same arity, respectively over the source schema S and over the peer schema G . The local mapping assertions establish the connection between the elements of the source schema and those of the peer schema in P_i . In particular, an assertion of the form $cq_S \rightsquigarrow cq_G$ specifies that all the data satisfying the query cq_S over the sources also satisfy the concept in the peer schema represented by the query cq_G . In the terminology used in data integration, the combination of peer schema, source schema, and local mapping assertions constitutes a GLAV *data integration system* [23] managing a set of sound data sources S defined in terms of a (virtual) global schema G .
- M is a set of *P2P mapping assertions*, which specify the semantic relationships that the peer P_i has with the other peers. Each assertion in M is an expression of the form $cq' \rightsquigarrow cq$, where cq , called the *head* of the assertion, is a conjunctive query over the peer (schema of) P_i , while cq' , called the *tail* of the assertion, is a conjunctive query of the same arity as cq over (the schema of) one of the other peers in \mathcal{P} . A P2P mapping assertion $cq' \rightsquigarrow cq$ from peer P_j to peer P_i expresses the fact that the P_j -concept represented by cq' is mapped to the P_i -concept represented by cq . From an extensional point of view, the assertion specifies that every tuple that can be retrieved from P_j by issuing query cq' satisfies cq in P_i .
- \mathcal{L} is a relational query language specifying the class of queries that the peer P_i can process. We assume that \mathcal{L} is any fragment of FOL that accepts at least conjunctive queries and the sentence \perp_i . We say that the queries in \mathcal{L} are those *accepted by* P_i . Notice that this implies

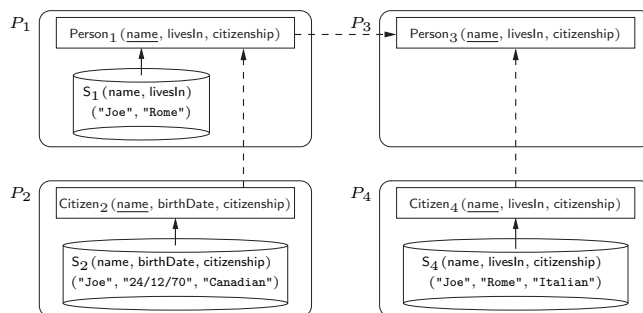


Figure 1: A P2P data integration system

that, for each P2P mapping assertion $cq' \rightsquigarrow cq$ from another peer P_j to peer P_i in M , we have that cq' is accepted by P_j .

An *extension* for a P2PDIS $\mathcal{P} = \{P_1, \dots, P_2\}$ is a set $\mathcal{D} = \{D_1, \dots, D_n\}$, where each D_i is an extension of the predicates in the local source schema of peer P_i .

A P2PDIS, together with an extension, is intended to be queried by external users. A user enquires the whole system by accessing any peer P of \mathcal{P} , and by issuing a *query* q to P . The query q is processed by P if and only if q is expressed over the schema of P and is accepted by P .

Remark. Notice that, in the above abstract model, constraints (i.e., formulas) are not allowed on the source schema, while they can be expressed over the global schema. This is a classical assumption in data integration systems [23]. The reason for such an assumption can be explained by the following observations. In a P2PDIS scenario, sources are autonomous systems which are not under the control of the peer: in other words, the source schema is an input in the construction of the peer. So, in general, a source schema may contain constraints. However, if there are constraints over the source schema, it is assumed that sources provide data to the peer that are consistent with such constraints. Thus, such constraints are actually used by the sources when providing data to the peer. Consequently, constraints on the source schema are only used by the peer in order to derive (at peer design time) constraints on the global schema, but they are not explicitly considered at run-time, i.e., when answering queries. Such an assumption is generally adopted in data integration applications (see [23]) and it is also assumed in other database applications like data exchange [11].

Example 5 Let us consider the P2PDIS in Figure 1, in which we have 4 peers P_1 , P_2 , P_3 , and P_4 (in the following, we assume that each peer P_i is identified by i).

The global schema of peer P_1 is formed by a relation schema $\text{Person}_1(\underline{\text{name}}, \text{livesIn}, \text{citizenship})$, where name is the key (we underline the key of a relation). P_1 contains a local source $S_1(\text{name}, \text{livesIn})$, mapped to the global view by the assertion $\{x, y \mid S_1(x, y)\} \rightsquigarrow \{x, y \mid \exists z. \text{Person}_1(x, y, z)\}$. Moreover, it has a P2P mapping assertion $\{x, z \mid \exists y. \text{Citizen}_2(x, y, z)\} \rightsquigarrow \{x, z \mid \exists y. \text{Person}_1(x, y, z)\}$ relating information in peer P_2 to those in peer P_1 .

P_2 has $\text{Citizen}_2(\underline{\text{name}}, \text{birthDate}, \text{citizenship})$ as global schema, and a local source $S_2(\text{name}, \text{birthDate}, \text{citizenship})$ mapped to the global schema through the local mapping $\{x, y, z \mid S_2(x, y, z)\} \rightsquigarrow \{x, y, z \mid \text{Citizen}_2(x, y, z)\}$. P_2 has no P2P mappings.

P_3 has $\text{Person}_3(\underline{\text{name}}, \text{livesIn}, \text{citizenship})$ as global schema, contains no local sources, and has a P2P mapping $\{x, y, z \mid \text{Person}_1(x, y, z)\} \rightsquigarrow \{x, y, z \mid \text{Person}_3(x, y, z)\}$ with P_1 , and a P2P mapping $\{x, y, z \mid \text{Citizen}_4(x, y, z)\} \rightsquigarrow \{x, y, z \mid \text{Person}_3(x, y, z)\}$ with P_4 .

P_4 has $\text{Citizen}_4(\underline{\text{name}}, \text{livesIn}, \text{citizenship})$ as global schema, and a local source $S_4(\text{name}, \text{livesIn}, \text{citizenship})$ mapped to the global schema through the local mapping $\{x, y, z \mid S_4(x, y, z)\} \rightsquigarrow \{x, y, z \mid \text{Citizen}_4(x, y, z)\}$. P_4 has no P2P mappings.

Finally, Figure 1 shows also an extension of the P2P data integration system, which includes $S_1(\text{"Joe"}, \text{"Rome"})$, $S_2(\text{"Joe"}, \text{"24/12/70"}, \text{"Canadian"})$, and $S_4(\text{"Joe"}, \text{"Rome"}, \text{"Italian"})$. ■

4.2 Formalization of P2P systems in K45_n^A

In order to logically formalize a P2PDIS, several aspects of the intended meaning of information in such a system must be taken into account. Due to lack of space,² here we only focus on *inconsistency tolerance*, which is the characteristic that enforces the need of a nonmonotonic logic for the above purpose. Informally, inconsistency tolerance corresponds to the ability of providing a semantics to the system even in the presence of contradicting information (e.g., data contradicting integrity constraints on the peer schemas).

More specifically, we want the P2PDIS to be inconsistency-tolerant in the following sense:

1. When a peer is *locally inconsistent*, i.e., data at the sources in P_i contradict, via the local mapping, the peer schema, making the whole peer inconsistent, the P2PDIS should be equivalent to the one obtained by eliminating the peer P_i from the system. In other words, an inconsistent peer should be “isolated” from the other peers: in this way, a local inconsistency does not affect the overall consistency (and

²For a detailed description of the intended semantics of information in a P2PDIS we refer to [16, 6].

meaning) of the system. The choice of isolating locally inconsistent peers is motivated by the modularity of P2PDISs pursued by our approach, in which each peer is considered as a black box. Of course, the study of inconsistency might be also interesting in an alternative setting not focused on modularity. However, this is outside the scope of the present paper.

2. In the presence of *P2P inconsistency*, i.e., when in a peer P_i the data coming from another peer P_j (through a P2P mapping) contradict the local data of P_i (or the data coming to P_i from another peer P_k), the peer P_i should not reach an inconsistent state: rather, it should discard a *minimal* amount of the data retrieved from the other peers in order to preserve consistency.

Due to the characteristics mentioned above, $K45_n^A$ is well-suited to formalize P2PDISs. Let $\mathcal{P} = \{P_1, \dots, P_n\}$ be a P2PDIS in which each peer P_i has identifier i . We use the modal operators \mathbf{K}_i and \mathbf{A}_i to model the peer i . More precisely, for each peer $P_i = (i, G, S, L, M, \mathcal{L})$ we define the theory $\mathcal{T}_K(P_i)$ in $K45_n^A$ as the union of the following sentences:

- Global schema G of P_i : for each sentence ϕ in G , we have

$$\mathbf{K}_i \phi$$

Observe that ϕ is a first-order sentence expressed in the alphabet of P_i , which is disjoint from the alphabets of all the other peers in \mathcal{P} .

- Local mapping assertions L between G and the local source schema S : for each mapping assertion $\{\mathbf{x} \mid \exists \mathbf{y}. \text{body}_{cq_S}(\mathbf{x}, \mathbf{y})\} \rightsquigarrow \{\mathbf{x} \mid \exists \mathbf{z}. \text{body}_{cq_G}(\mathbf{x}, \mathbf{z})\}$ in L , we have

$$\mathbf{K}_i(\forall \mathbf{x}. \exists \mathbf{y}. \text{body}_{cq_S}(\mathbf{x}, \mathbf{y}) \supset \exists \mathbf{z}. \text{body}_{cq_G}(\mathbf{x}, \mathbf{z}))$$

- P2P mapping assertions M : for each P2P mapping assertion $\{\mathbf{x} \mid \exists \mathbf{y}. \text{body}_{cq_j}(\mathbf{x}, \mathbf{y})\} \rightsquigarrow \{\mathbf{x} \mid \exists \mathbf{z}. \text{body}_{cq_i}(\mathbf{x}, \mathbf{z})\}$ between the peer j and the peer i in M , we have

$$\forall \mathbf{x}. \neg \mathbf{A}_j \perp_j \wedge \mathbf{K}_j(\exists \mathbf{y}. \text{body}_{cq_j}(\mathbf{x}, \mathbf{y})) \wedge \neg \mathbf{A}_i(\neg \exists \mathbf{z}. \text{body}_{cq_i}(\mathbf{x}, \mathbf{z})) \supset \mathbf{K}_i(\exists \mathbf{z}. \text{body}_{cq_i}(\mathbf{x}, \mathbf{z})) \quad (1)$$

Informally, the above sentence specifies the following rule: for each tuple of values \mathbf{t} , if peer j is *locally consistent* and knows the sentence $\exists \mathbf{y}. \text{body}_{cq_j}(\mathbf{t}, \mathbf{y})$, and the sentence $\exists \mathbf{z}. \text{body}_{cq_i}(\mathbf{t}, \mathbf{z})$ is *consistent with what peer i knows*, then peer i knows the sentence $\exists \mathbf{z}. \text{body}_{cq_i}(\mathbf{t}, \mathbf{z})$. In other words, information flows from peer j to peer i through a P2P

mapping assertion only if j is locally consistent and if adding such information to peer i does not give rise to a P2P inconsistency in peer i . More precisely, the meaning of the above sentence in $\mathbf{K45}_n^A$ is that exactly a *maximal* amount of information (i.e., a maximal set of tuples) consistent with peer i flows from peer j to peer i through the P2P mapping assertion. Moreover, under such a formalization the P2PDIS is tolerant to local inconsistency, in the sense that the peers that are locally inconsistent are “isolated” from the rest of the system (i.e., they cannot propagate information).

We denote by $\mathcal{T}_K(\mathcal{P})$ the theory corresponding to the P2PDIS \mathcal{P} , i.e., $\mathcal{T}_K(\mathcal{P}) = \bigcup_{i=1, \dots, n} \mathcal{T}_K(P_i)$.

The extension $\mathcal{D} = \{D_1, \dots, D_n\}$ of a P2PDIS \mathcal{P} is modeled as a sentence constituted by the conjunction of all facts corresponding to the tuples stored in the sources, i.e., $DB(\mathcal{D}) = \bigwedge_{i=1}^n DB(D_i)$ where $DB(D_i) = \mathbf{K}_i(\bigwedge_{t \in r^{D_i}} r(t))$.

A client of the P2PDIS interacts with one of the peers, say peer P_i , posing a *query* to it. A query q is an open formula $q(\mathbf{x})$ with free variables \mathbf{x} expressed in the language accepted by the peer P_i (we recall that such a language is a subset of first-order logic). The semantics of a query $q \in \mathcal{L}$ posed to a peer $P_i = (i, G, S, L, M, \mathcal{L})$ of \mathcal{P} with respect to an extension \mathcal{D} is defined as the set of tuples $\{\mathbf{t} \mid \mathcal{T}_K(\mathcal{P}) \cup DB(\mathcal{D}) \models_{\mathbf{K45}_n^A} \mathbf{K}_i q(\mathbf{t})\}$, where $q(\mathbf{t})$ denotes the sentence obtained from the open formula $q(\mathbf{x})$ by replacing all occurrences of the free variables in \mathbf{x} with the corresponding constants in \mathbf{t} .

Example 6 We are now able to provide the formalization of the P2PDIS of Example 5. The theories $\mathcal{T}_K(P_1), \dots, \mathcal{T}_K(P_4)$ modeling the four peers are reported in Figure 2. ■

It can be shown (see [6] for details) that the above formalization in $\mathbf{K45}_n^A$ provides a formal semantics to P2PDISs that, besides other things, exactly captures the two notions of inconsistency tolerance above defined. Indeed, from the above formalization it follows that:

- when inconsistency arises between local data and non-local data in a peer, i.e., when data coming from the peer sources through the local mapping contradicts the data retrieved by a peer through a P2P mapping, then the peer always prefers the local data. Formally, in this case there is one $\mathbf{K45}_n^A$ -model for the P2PDIS, which represents the situation in which non-local data is discarded;
- when inconsistency arises between two different pieces of non-local data, i.e., when a piece of data retrieved by a peer through a P2P mapping contradicts another piece of data retrieved through the P2P

<hr/> $\mathbf{K}_1(\forall x, y, y', z, z'. \text{Person}_1(x, y, z) \wedge \text{Person}_1(x, y', z') \supset y = y' \wedge z = z')$ $\mathbf{K}_1(\forall x, y. \mathbf{S}_1(x, y) \supset \exists z. \text{Person}_1(x, y, z))$ $\forall x, z. \neg \mathbf{A}_2 \perp_2 \wedge \mathbf{K}_2(\exists y. \text{Citizen}_2(x, y, z)) \wedge \neg \mathbf{A}_1 \neg(\exists y. \text{Person}_1(x, y, z)) \supset$ $\mathbf{K}_1(\exists y. \text{Person}_1(x, y, z))$ <hr/>
theory $\mathcal{T}_K(P_1)$
<hr/> $\mathbf{K}_2(\forall x, y, y', z, z'. \text{Citizen}_2(x, y, z) \wedge \text{Citizen}_2(x, y', z') \supset y = y' \wedge z = z')$ $\mathbf{K}_2(\forall x, y, z. \mathbf{S}_2(x, y, z) \supset \text{Citizen}_2(x, y, z))$ <hr/>
theory $\mathcal{T}_K(P_2)$
<hr/> $\mathbf{K}_3(\forall x, y, y', z, z'. \text{Person}_3(x, y, z) \wedge \text{Person}_3(x, y', z') \supset y = y' \wedge z = z')$ $\forall x, y. \neg \mathbf{A}_1 \perp_1 \wedge \mathbf{K}_1(\exists z. \text{Person}_1(x, z, y)) \wedge \neg \mathbf{A}_3 \neg(\exists z. \text{Person}_3(x, z, y)) \supset$ $\mathbf{K}_3(\exists z. \text{Person}_3(x, z, y))$ $\forall x, y, z. \neg \mathbf{A}_4 \perp_4 \wedge \mathbf{K}_4(\text{Citizen}_4(x, y, z)) \wedge \neg \mathbf{A}_3 \neg \text{Person}_3(x, y, z) \supset$ $\mathbf{K}_3 \text{Person}_3(x, y, z)$ <hr/>
theory $\mathcal{T}_K(P_3)$
<hr/> $\mathbf{K}_4(\forall x, y, y', z, z'. \text{Citizen}_4(x, y, z) \wedge \text{Citizen}_4(x, y', z') \supset y = y' \wedge z = z')$ $\mathbf{K}_4(\forall x, y, z. \mathbf{S}_4(x, y, z) \supset \text{Citizen}_4(x, y, z))$ <hr/>
theory $\mathcal{T}_K(P_4)$

Figure 2: Theories $\mathcal{T}_K(P_1), \dots, \mathcal{T}_K(P_4)$ modeling the P2P system of Figure 1 in $\mathbf{K45}_n^A$

mappings, then no preference is made between these two pieces of information, in the sense that in this case there are two $K45_n^A$ -models for the P2PDIS, each of which represents the situation in which one of the two pieces of data is discarded.

Notably, based on the above formalization in $K45_n^A$, it is possible to define sound and complete algorithms for computing answers to queries in a P2PDIS setting. We refer the interested reader to [6].

Finally, we observe that, in principle, the above peer-to-peer setting could also be formalized through alternative, non-modal, nonmonotonic formalisms, like default logic [36] or preferred subtheories [4]. However, we believe that the epistemic multi-modal logic allows for a more natural and general logical reconstruction of the peer-to-peer scenario. In particular, we are currently working at exploiting the introspective abilities of the logic in order to formalize further characteristics of peers in a distributed information system, like the so-called *authorization views* [48], i.e., mechanisms for controlling information access.

5 Conclusions

In this paper we have proposed a first attempt to define a *multi-modal, first-order, nonmonotonic* family of logics. In particular, the logics \mathcal{S}_n^A presented in this paper generalize recent approaches in epistemic logic and nonmonotonic modal logics in many respects.

We have also illustrated the need for multi-modal nonmonotonic logics in the field of distributed systems. Interestingly, the possibility of modeling knowledge in distributed systems was also the initial motivation behind one of the first nonmonotonic modal logics, i.e., Halpern and Moses' logic of minimal knowledge [18].

An interesting extension of the present work is towards reasoning in the logics \mathcal{S}_n^A . The first results in this direction appear in [6], in which an algorithm is presented for reasoning in the restricted fragment of the logic $K45_n$ which is able to logically model information in P2P systems. However, the important issue of defining a general proof system for $K45_n^A$ and for other logics in the \mathcal{S}_n^A family is still open. So far, the results in this direction only concern: (i) either single-agent, propositional fragments of \mathcal{S}_n^A , in particular ground nonmonotonic modal logics, for which a tableaux method has been defined in [8]; or (ii) the single-agent propositional fragment of $S5_n^A$, i.e., the logic MKNF, for which reasoning has been studied in [38].

Finally, it would be very interesting to investigate whether the logics \mathcal{S}_n^A can be characterized by fix-point semantics, in a way analogous to other families of nonmonotonic modal logics [46, 30, 41].

Acknowledgments

The author wishes to thank the anonymous reviewers for their precious suggestions. This research has been partially supported by the projects INFOMIX (IST-2001-33570), and TONES (IST-007603) funded by the EU, and by the project HYPER, funded by IBM through a Shared University Research Award Grant.

References

- [1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison Wesley Publ. Co., 1995.
- [2] P. A. Bernstein, F. Giunchiglia, A. Kementsietsidis, J. Mylopoulos, L. Serafini, and I. Zaihrayeu. Data management for peer-to-peer computing: A vision. In *Proc. of the 5th Int. Workshop on the Web and Databases (WebDB 2002)*, 2002.
- [3] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2001.
- [4] G. Brewka. Preferred subtheories: an extended logical framework for default reasoning. In *Proc. of the 11th Int. Joint Conf. on Artificial Intelligence (IJCAI'89)*, pages 1043–1048, 1989.
- [5] D. Calvanese, E. Damaggio, G. De Giacomo, M. Lenzerini, and R. Rosati. Semantic data integration in P2P systems. In *Proc. of the Int. Workshop on Databases, Information Systems and Peer-to-Peer Computing (DBISP2P 2003)*, 2003.
- [6] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati. Inconsistency tolerance in P2P data integration: an epistemic logic approach. In *Proc. of the 10th Int. Sym. on Database Programming Languages (DBPL 2005)*, pages 90–105, 2005.
- [7] D. Calvanese, G. De Giacomo, M. Lenzerini, and R. Rosati. Logical foundations of peer-to-peer data integration. In *Proc. of the 23rd ACM SIGACT SIGMOD SIGART Symp. on Principles of Database Systems (PODS 2004)*, pages 241–251, 2004.
- [8] F. M. Donini, F. Massacci, D. Nardi, and R. Rosati. A uniform tableaux method for nonmonotonic modal logics. In J. J. Alferes, L. M. Pereira, and E. Orłowska, editors, *Proc. of the 5th Eur. Workshop on Logics in Artificial Intelligence (JELIA'96)*, volume 1126 of *Lecture Notes in Artificial Intelligence*, pages 87–103. Springer, 1996.

- [9] F. M. Donini, D. Nardi, and R. Rosati. Ground nonmonotonic modal logics. *J. of Logic and Computation*, 7(4):523–548, Aug. 1997.
- [10] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge*. The MIT Press, 1995.
- [11] R. Fagin, P. G. Kolaitis, R. J. Miller, and L. Popa. Data exchange: Semantics and query answering. In *Proc. of the 9th Int. Conf. on Database Theory (ICDT 2003)*, pages 207–224, 2003.
- [12] E. Franconi, G. Kuper, A. Lopatenko, and L. Serafini. A robust logical and computational characterisation of peer-to-peer database systems. In *Proc. of the VLDB International Workshop On Databases, Information Systems and Peer-to-Peer Computing (DBISP2P 2003)*, 2003.
- [13] M. Gelfond and V. Lifschitz. The stable model semantics for logic programming. In *Proceedings of the Fifth Logic Programming Symposium*, pages 1070–1080. The MIT Press, 1988.
- [14] M. Gelfond and V. Lifschitz. Logic programs with classical negation. In *Proceedings of the Seventh International Conference on Logic Programming*, pages 579–597. The MIT Press, 1990.
- [15] M. Gelfond and V. Lifschitz. Classical negation in logic programs and disjunctive databases. *New Generation Computing*, 9:365–385, 1991.
- [16] A. Halevy, Z. Ives, D. Suciu, and I. Tatarinov. Schema mediation in peer data management systems. In *Proc. of the 19th IEEE Int. Conf. on Data Engineering (ICDE 2003)*, pages 505–516, 2003.
- [17] J. Y. Halpern. A theory of knowledge and ignorance for many agents. *J. of Logic and Computation*, 7(1):79–108, 1997.
- [18] J. Y. Halpern and Y. Moses. Towards a theory of knowledge and ignorance: Preliminary report. In K. Apt, editor, *Logic and models of concurrent systems*. Springer, 1985.
- [19] 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.
- [20] K. Inoue and C. Sakama. On positive occurrences of negation as failure. In *Proc. of the 4th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR'94)*, pages 293–304. Morgan Kaufmann, 1994.
- [21] M. Kaminski and G. Rey. First-order non-monotonic modal logics. *Fundamenta Informaticae*, 42(3–4):303–333, 2000.

- [22] M. Kaminski and G. Rey. Revisiting quantification in autoepistemic logic. *ACM Trans. on Computational Logic*, 3(4):542–561, 2002.
- [23] M. Lenzerini. Data integration: A theoretical perspective. In *Proc. of the 21st ACM SIGACT SIGMOD SIGART Symp. on Principles of Database Systems (PODS 2002)*, pages 233–246, 2002.
- [24] H. J. Levesque. All I know: a study in autoepistemic logic. *Artificial Intelligence*, 42:263–310, 1990.
- [25] H. J. Levesque and G. Lakemeyer. *The Logic of Knowledge Bases*. The MIT Press, 2001.
- [26] V. Lifschitz. Nonmonotonic databases and epistemic queries. In *Proc. of the 12th Int. Joint Conf. on Artificial Intelligence (IJCAI'91)*, pages 381–386, 1991.
- [27] V. Lifschitz. Minimal belief and negation as failure. *Artificial Intelligence*, 70:53–72, 1994.
- [28] V. Lifschitz and T. Woo. Answer sets in general nonmonotonic reasoning (preliminary report). In *Proc. of the 3rd Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR'92)*, pages 603–614. Morgan Kaufmann, 1992.
- [29] F. Lin and Y. Shoham. A logic of knowledge and justified assumptions. *Artificial Intelligence*, 57(2–3):271–289, 1992.
- [30] W. Marek and M. Truszczyński. *Nonmonotonic Logics: Context-Dependent Reasoning*. Springer, 1993.
- [31] D. McDermott. Non-monotonic logic II: Non-monotonic modal theories. *J. of the ACM*, 29:33–57, 1982.
- [32] D. McDermott and J. Doyle. Non-monotonic logic I. *Artificial Intelligence*, 13:41–72, 1980.
- [33] J. J. Meyer and W. van der Hoek. *Epistemic Logic for AI and Computer Science*. Cambridge University Press, 1995.
- [34] R. C. Moore. Semantical considerations on nonmonotonic logic. *Artificial Intelligence*, 25:75–94, 1985.
- [35] R. C. Moore. Semantical considerations on nonmonotonic logic. *Acta Informatica*, 25:75–94, 1985.
- [36] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

- [37] R. Reiter. What should a database know? *J. of Logic Programming*, 14:127–153, 1990.
- [38] R. Rosati. Reasoning about minimal belief and negation as failure. *J. of Artificial Intelligence Research*, 11:277–300, 1999.
- [39] R. Rosati. Minimal belief and negation as failure in multi-agent systems. *Ann. of Mathematics and Artificial Intelligence*, 37:5–32, 2003.
- [40] G. Schwarz. Bounding introspection in nonmonotonic logics. In *Proc. of the 3rd Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR'92)*, pages 581–590, 1992.
- [41] G. Schwarz. Minimal model semantics for nonmonotonic modal logics. In *Proc. of the 7th IEEE Symp. on Logic in Computer Science (LICS'92)*, pages 34–43. IEEE Computer Society Press, 1992.
- [42] G. Schwarz and V. Lifschitz. Extended logic programs as autoepistemic theories. In *Proc. of the 2nd Int. Workshop on Logic Programming and Non-monotonic Reasoning (LPNMR'93)*, pages 101–114. The MIT Press, 1993.
- [43] G. Schwarz and M. Truszczyński. Minimal knowledge problem: a new approach. *Artificial Intelligence*, 67:113–141, 1994.
- [44] Y. Shoham. Nonmonotonic logics: Meaning and utility. In *Proc. of the 10th Int. Joint Conf. on Artificial Intelligence (IJCAI'87)*, pages 388–392, 1987.
- [45] R. Stalnaker. A note on non-monotonic modal logic. *Artificial Intelligence*, 64(2):183–196, 1993.
- [46] M. Tiomkin and M. Kaminski. Nonmonotonic default modal logics. In *Proc. of the 3rd Conf. on Theoretical Aspects of Reasoning about Knowledge (TARK'90)*, pages 73–84, 1990.
- [47] M. Truszczyński. Modal nonmonotonic logic with restricted application of the negation as failure to prove rule. *Fundamenta Informaticae*, 14:355–366, 1991.
- [48] Z. Zhang and A. Mendelzon. Authorization views and conditional query containment. In *Proc. of the 10th Int. Conf. on Database Theory (ICDT 2005)*, volume 3363 of *Lecture Notes in Computer Science*, pages 259–273. Springer, 2005.