Logic-based Knowledge Representation*

Franz Baader

Theoretical Computer Science, RWTH Aachen, 52074 Aachen, Germany
baader@informatik.rwth-aachen.de

Abstract. After a short analysis of the requirements that a knowledge representation language must satisfy, we introduce Description Logics, Modal Logics, and Nonmonotonic Logics as formalisms for representing terminological knowledge, time-dependent or subjective knowledge, and incomplete knowledge respectively. At the end of each section, we briefly comment on the connection to Logic Programming.

1 Introduction

This section is concerned with the question under which conditions one may rightfully claim to have *represented knowledge* about an application domain, and not just *stored data* occurring in this domain.[1] In the early days of Artificial Intelligence and Knowledge Representation, there was a heated discussion on whether logic can at all be used as a formalism for Knowledge Representation (see e.g. [135, 91, 92]). One aspect of the requirements on knowledge representation formalisms that can be derived from the considerations in this section is very well satisfied by logical formalisms. We shall see, however, that some other aspects are not treated satisfactorily, at least not by classical first-order predicate logic. The main purpose of this article is to demonstrate that these deficiencies can be overcome with the help of other logic-based formalisms, namely Description Logics, Modal Logics, and Nonmonotonic Logics. More recently, it has been argued (see e.g. [110, 23]) that Logic Programming can serve as a convenient and universal formalism for Knowledge Representation. However, as indicated by their name, Logic Programming languages are *programming* languages, and thus not necessarily appropriate as *representation* languages.

In a nut-shell (and somewhat exaggerated), the difference between knowledge-based programming (which processes knowledge) and classical programming (which processes data) can be formulated as follows. In classical programming, one designs specialized programs that are tailored to a specific application problem. The knowledge about the problem description and the application domain is implicitly represented in the structure of the program, and must thus be acquired by the programmer. In knowledge-based programming, the knowledge about the

* This is an extended and updated version of an article that has appeared (in German) in the journal KI, 3/96:8-16, 1996.

[1] The division between knowledge and data is, of course, not strict, and this article will not give a definitive answer to the problem of distinguishing between both.

application domain is represented explicitly (in an appropriate representation formalism); ideally, the processing can be done with the help of general (i.e., application-independent) problem solving methods. Thus, the knowledge can be acquired and represented by an application domain expert, who need not be well-acquainted with details of the implementation. As a simple example, let us consider the task of finding out whether two nodes in a directed graph (which may represent the hierarchical organization of a company) are connected (whether one employee is the boss of another one). A first solution, which is very far away from being knowledge-based, might be a program that encodes the structure of the specific graph in its control structure. A second, more reasonable, solution might explicitly represent the graph in an appropriate data structure (which list the nodes that are directly connected), and then code the meaning of "connected" in a program that is independent of the specific graph. This program must be able to compute the connected nodes from the given information about the directly connected nodes. Finally, an even more knowledge-based solution could be one in which the knowledge about the meaning of "connected" is also represented in an explicit way, for example, by the following Horn clauses:

$$\begin{aligned}\text{directly-connected}(x, y) &\rightarrow \text{connected}(x, y),\\ \text{directly-connected}(x, z) \wedge \text{connected}(z, y) &\rightarrow \text{connected}(x, y).\end{aligned}$$

A general problem solving component, which is able to handle such clauses (e.g., the interpreter of a Logic Programming language), could now use these formulae together with the explicit information on the directly connected nodes to infer the connected nodes.

A knowledge representation (KR) formalism should allow for the symbolic representation of all the knowledge relevant in a given application domain. From what we have said so far about the use of knowledge in knowledge-based programming, we can derive two requirements that such a formalism must satisfy. On the one hand, it must be equipped with a *declarative semantics*, that is, the meaning of the entries in a knowledge base (KB) must be defined independently of the programs that operate on the KB (no purely procedural semantics). Otherwise, the knowledge cannot be acquired by a domain expert without detailed knowledge of the implementation of these programs. Usually, such a declarative semantics is given by mapping the symbolic expressions into (an abstraction of) the relevant segment of the "world." In addition, one needs a notion of "truth," which makes it possible to determine which of the symbolic statements hold in the current "world."

On the other hand, one needs an *"intelligent" retrieval mechanism*, which is able to realize the above-mentioned general problem solving component. This mechanism should allow to retrieve knowledge that is only implicitly present in the KB, that is, it should be able to deduce implicitly represented knowledge from the explicit knowledge. In our above example, the information that two nodes are connected is implicit knowledge, which must be deduced from the explicit knowledge (i.e., the information about directly connected nodes and the Horn clauses specifying the meaning of "connected") with the help of an appropriate inference engine. The behaviour of this deductive component should depend only on the semantics of the representation language, and not on the syntactic form of the entries in the KB, i.e., semantically equivalent entries should lead to the same results. If we use first-order predicate logic as semantics for the Horn clauses in our above example, then the usual Prolog interpreters do not fulfill this requirement. According to the semantics, the order of the clauses and of the conjuncts is irrelevant. Any Prolog programmer will know, however, that such a re-ordering may drastically change the behaviour of the program.

Another requirement that is usually imposed on KR formalisms is that of allowing for a *structured representation* of the knowledge. One aspect of a structured representation is that semantically related information (for example, all the knowledge about knowledge representation based on Description Logics) should also syntactically be grouped together. This requirement is, on the one hand, justified by cognitive adequacy.[2] On the other hand, there are purely pragmatic reasons, since a structured representation allows for faster retrieval.

Critics of a logic-based approach to KR often (implicitly) equate logic with first-order predicate logic. If we consider in how far *first-order predicate logic* satisfies the requirements introduced above, it shows in fact some strong deficits. The Tarskian semantics of predicate logic is the prototype of a declarative semantics; however, it does not allow for an adequate treatment of incomplete and contradictory knowledge, or of subjective and time-dependent knowledge. Later on, we shall see that this deficit can be overcome by considering Nonmonotonic Logics and Modal Logics. The usual syntax of first-order predicate logic does not support a structured representation of knowledge. Since all the relevant inference problems are undecidable, it is also not possible to provide for semantically adequate inference procedures. In the following section, we shall describe so-called Description Logics,[3] which are an attempt to overcome both of the last-mentioned problems by using a non-standard syntax and by restricting the expressive power.

According to what was said until now, Logic Programming languages qualify as representation languages only if they are equipped with an appropriate declarative semantics. This is not the case for Prolog, as illustrated by the example. It should be noted that I do not claim that it is not possible to solve a specific representation problem with the help of a Prolog program: since Prolog is computationally complete, this is always possible. However, this is a programming approach in which the knowledge is not encoded independently of the way in which it is processed. Exactly because of Prolog being computationally complete, all the responsibility (e.g., for termination of the inference process) lies in the hands of programmer, and cannot automatically be guaranteed. In a KR system, the intelligent retrieval mechanisms should (ideally) be able to handle all the knowledge that is represented in a syntactically correct way, which means that one must restrict oneself to a sublanguage (such as Datalog) if one wants to

[2] In the human brain, correlated information is also not stored in unrelated parts.

[3] Other names used in the literature are "terminological logics," "concept languages," and "KL-ONE-like KR languages."

use a Logic Programming approach for KR. Overviews on the topic of extended Logic Programs with declarative semantics and their application to representing knowledge can be found in [53, 134, 23, 4].

2 Description Logics

The attempt to provide for a structured representation of information was one of the main motivations for introducing early KR formalisms such as Semantic Networks and Frames. Description Logics[4] are logic-based formalisms that are based on the same ideas as Semantic Networks and Frames, but avoid the formal deficiencies that made the use of these precursor formalisms as KR formalisms problematic.

Precursors. *Frames* have been introduced by Minsky [135] as record-like data structures for representing prototypical situations and objects. The key idea was to collect all the information necessary for treating a situation in one place (the frame for this situation). In [135], Minsky combined his introduction of the frame idea with a general rejection of logic as a KR formalism. Hayes [91, 92] criticized that Frames lack a formal semantics, and showed that with an appropriate formalization (of their monotonic, non-procedural aspects), Frames can be seen as a syntactic variant of first-order predicate logic.

Semantic Networks, which we shall consider in somewhat more detail, have been developed by Quillian [155] for representing the semantics of natural language (this explains their name). They allow to represent concepts and objects as nodes in a graph (see Figure 1). Such a graph has two different types of edges: property edges assign properties (like colour) to concepts (e.g., frog) and objects (e.g., Kermit), whereas IS-A-edges introduce hierarchical relationships among concepts and instance relationships between objects and concepts.

Properties are inherited along IS-A-edges. For example, tree frogs, and thus also Kermit, inherit the colour green from frogs. In many systems, inheritance is only by default, that is, grass frogs do not inherit the property green, since this would be in contradiction with the explicit property edge saying that grass frogs are brown. The missing formal *semantics of Semantic Networks* was criticized by Woods [176] and Brachman [31]. The meaning of a given Semantic Network was left to the intuition of the users and the programmers who implemented the programs processing the networks. Consequently, identical networks could lead to very different results, depending on which of the systems was used to process it. As an illustration of this problem, we point out the possible ambiguities in the interpretation of property edges. In Figure 1, the property edge "colour" from "Frog" to "Green" may, on the one hand, mean that green is the only possible colour for frogs (value restriction). On the other hand, it could mean that any frog has at least the colour green, but may have other colours too (it might be

[4] Although there is no overview article that covers all aspects of this research area, [177, 15, 62] can serve as a starting point.

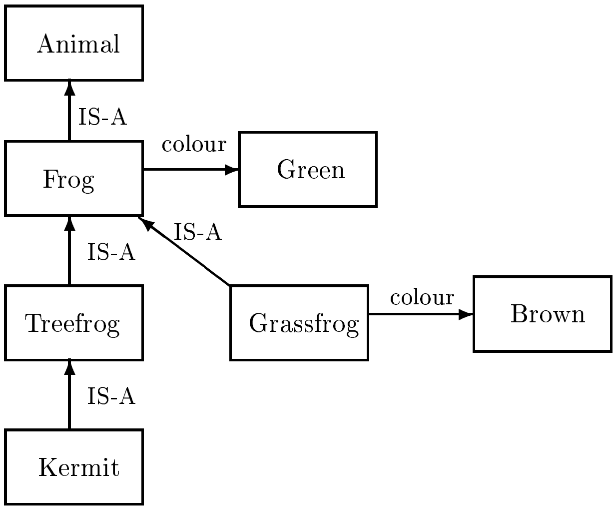

Fig. 1. A semantic network.

green with red stripes). A partial reconstruction of the semantics of Semantic Networks within first-order predicate logic was presented in [167].

Description Logics (DL) make the quantifier that is implicitly present in property edges (universal in the reading as value restrictions, and existential for the other option) explicit (see below).

Systems and applications. Description Logics are descended from so-called "structured inheritance networks" [31, 32], which were first realized in the system KL-ONE [34]. Their main idea is to start with atomic *concepts* (unary predicates) and *roles* (binary predicates), and use a (rather small) set of epistemologically adequate constructors to build complex concepts and roles. This idea has been further developed both from the theoretical and the practical point of view. In particular, there is a great variety of successor systems (e.g., BACK [143, 149, 96], CLASSIC [29, 33], CRACK [36], DLP [148], FACT [98], FLEX [152], K-REP [127, 128], KRIS [14], LOOM [122, 121], SB-ONE [106]), which have been used in different application domains such as natural language processing [154], configuration of technical systems [178, 42, 158, 133], software information systems [50], optimizing queries to databases [41, 25, 24], or planning [107].

Syntax and semantics. Figure 2 introduces syntax and semantics of some of the concept constructors employed in systems or investigated in the literature. Most of the systems do not provide for all of these constructors, and vice versa, they may use additional constructors not introduced here. An extensive list of (most of) the constructors considered until now can be found in [9]. The first column of the figure shows the (Lisp-like) concrete syntax that is used in most

(and $C_1 \ldots C_n$)	$C_1 \sqcap \ldots \sqcap C_n$	$C_1^{\mathcal{I}} \cap \ldots \cap C_n^{\mathcal{I}}$
(or $C_1 \ldots C_n$)	$C_1 \sqcup \ldots \sqcup C_n$	$C_1^{\mathcal{I}} \cup \ldots \cup C_n^{\mathcal{I}}$
(not C)	$\neg C$	$\Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$
(some R C)	$\exists R.C$	$\{d \in \Delta^{\mathcal{I}} \mid \exists e\colon (d,e) \in R^{\mathcal{I}} \wedge e \in C^{\mathcal{I}}\}$
(all R C)	$\forall R.C$	$\{d \in \Delta^{\mathcal{I}} \mid \forall e\colon (d,e) \in R^{\mathcal{I}} \Rightarrow e \in C^{\mathcal{I}}\}$
(atleast n R)	$(\geq n\, R)$	$\{d \in \Delta^{\mathcal{I}} \mid \mathrm{card}(\{e \in \Delta^{\mathcal{I}} \mid (d,e) \in R^{\mathcal{I}}\}) \geq n\}$
(atmost n R)	$(\leq n\, R)$	$\{d \in \Delta^{\mathcal{I}} \mid \mathrm{card}(\{e \in \Delta^{\mathcal{I}} \mid (d,e) \in R^{\mathcal{I}}\}) \leq n\}$
(and $R_1 \ldots R_n$)	$R_1 \sqcap \ldots \sqcap R_n$	$R_1^{\mathcal{I}} \cap \ldots \cap R_n^{\mathcal{I}}$

Fig. 2. Syntax and semantics of concept and role terms.

of the systems, whereas the second column introduces the abstract syntax that is usually employed in theoretical work on Description Logics. Starting with atomic concepts and roles, one can use these constructors to build complex *concept terms* and (in the last row) *role terms*. Using the abstract syntax, the concept "Frog," which has been introduced in the Semantic Network of Figure 1, can be described by the concept term $\mathtt{Animal} \sqcap \forall \mathtt{colour}.\mathtt{Green}$, where `Animal` and `Green` are atomic concepts (concept names), and `colour` is an atomic role (role name). The universal quantifier makes clear that we have interpreted the property edge as a value restriction.

In order to define the semantics of concept terms, one considers interpretations $\mathcal{I}$, which consist of a non-empty set $\Delta^{\mathcal{I}}$ (the domain of the interpretation) and an interpretation function that assigns to every atomic concept A a set $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ and to every atomic role R a binary relation $R^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. The third column in Figure 2 shows how this interpretation function is extended to concept and role terms. An alternative way of defining the semantics of concept and role terms is to give a translation into formulae of first-order predicate logic: concept terms are translated into formulae with one free variable and role terms into formulae with two free variables. The exact definition of this translation is an easy consequence of the third column of Figure 2. For example, the above concept term for the concept "Frog" yields the following formula with free variable x:

$$\mathtt{Animal}(\mathtt{x}) \wedge \forall \mathtt{y}{:}(\mathtt{colour}(\mathtt{x},\mathtt{y}) \rightarrow \mathtt{Green}(\mathtt{y})).$$

Concept definitions can be used to assign names (abbreviations) to large terms. For example, the definition $\mathtt{Frog} \doteq \mathtt{Animal} \sqcap \forall \mathtt{colour}.\mathtt{Green}$ assigns the name `Frog` to the term $\mathtt{Animal} \sqcap \forall \mathtt{colour}.\mathtt{Green}$, which can then be used as an abbreviation for this term when constructing new concept terms.

A *terminology* (TBox) consists of a finite set of concept and role definitions of the form $A \doteq C$ and $P \doteq R$, where A is a concept name, P is a role name, C is a concept term, and R is a role term. Usually, one imposes the additional requirement that the definitions are unique (i.e., any name may occur at most once as a left-hand side of a definition) and acyclic (i.e., the definition of a name must not, directly or indirectly, refer to this name). As a consequence, definitions can be seen as macros, which can simply be expanded (see [141], Section 3.2.5, and [142]). An interpretation $\mathcal{I}$ is a model of a TBox iff it satisfies all the definitions $A \doteq C$ and $P \doteq R$ contained in the TBox, i.e., if $A^\mathcal{I} = C^\mathcal{I}$ and $P^\mathcal{I} = R^\mathcal{I}$ holds for these definitions.

In addition to this terminological component, the knowledge base of a DL system also has an *assertional component*, in which one can introduce individuals (by giving them names), and assert properties of these individuals. If a, b are names for individuals, C is a concept term, and R a role term, then $C(a)$ and $R(a, b)$ are *assertions*. A finite set of such assertions is called an *ABox*. An interpretation $\mathcal{I}$ (which also assigns elements $a^\mathcal{I} \in \Delta^\mathcal{I}$ to individual names a) is a model of these assertions iff $a^\mathcal{I} \in C^\mathcal{I}$ and $(a^\mathcal{I}, b^\mathcal{I}) \in R^\mathcal{I}$. For example, $\mathtt{Frog}(\mathtt{KERMIT})$ is a concept assertion and $\mathtt{colour}(\mathtt{KERMIT}, \mathtt{Colour07})$ is a role assertion. A *knowledge base* (KB) consist of a TBox $\mathcal{T}$ and an ABox $\mathcal{A}$.

Inference problems. As mentioned in the introduction, one is not just interested in retrieving the knowledge that is explicitly stored in such a KB: one should also like to have access to the knowledge represented implicitly. To make this possible, one must be able to draw inferences from the explicit knowledge. As examples of two typical inference problems in Description Logics, we shall consider the subsumption and the instance problem. *Subsumption* is concerned with the question whether one concept is a subconcept of another one. Formally, we define for given concept terms C and D and a TBox $\mathcal{T}$: C is subsumed by D w.r.t. $\mathcal{T}$ ($C \sqsubseteq_\mathcal{T} D$) iff $C^\mathcal{I} \subseteq D^\mathcal{I}$ holds in all models $\mathcal{I}$ of $\mathcal{T}$. DL systems usually offer the computation of the subsumption hierarchy of the concepts defined in a TBox as a system service (classification). When defining the *instance problem*, one considers a TBox $\mathcal{T}$ and an ABox $\mathcal{A}$. The individual a is an instance of the concept term C w.r.t. $\mathcal{T}$ and $\mathcal{A}$ iff $a^\mathcal{I} \in C^\mathcal{I}$ holds in all interpretations $\mathcal{I}$ that are models of both $\mathcal{T}$ and $\mathcal{A}$. For example, if the TBox contains the above definition of the concept `Frog`, and the ABox contains the above assertions for `KERMIT`, then `COLOUR07` is an instance of `Green` with respect to this TBox and ABox.

Inference algorithms. Many DL systems (for example, Back, Classic, K-Rep, Loom) employ inference procedures that only partially satisfy the requirement that the result should depend only on the semantics of the representation language, and not on the syntactic form of the entries in the KB. These systems use so-called *structural algorithms* (see, e.g., [141], Chapter 4), which are based on a point of view derived from Semantic Networks: the knowledge base is viewed as a directed graph. Structural subsumption algorithms usually proceed in two phases: first, the graphs corresponding to the concepts to be tested for subsumption are normalized, and then one tries to detect similarities in the normalized graphs. They have the advantage that they are usually very efficient (polynomial). An important disadvantage is that they are only sound, but *not complete* for reasonably expressive representation languages. If an incomplete subsumption algorithm answers a subsumption query with "yes," then this answer is correct (soundness). If it answers with "no," then this does not mean anything: because of the incompleteness, the subsumption relationship might nevertheless hold. Thus, the behaviour of the algorithm does not depend only on the semantics, but also on other factors that are not transparent to the user. As a way to overcome this problem without giving up the efficiency of the structural subsumption approach, the use of non-standard semantics has been proposed. In [146], a four-valued semantics characterizing the behaviour of a structural subsumption algorithm is introduced. The system CLASSIC [29, 33] employs an "almost" complete structural subsumption algorithm [30]. Its only incompleteness stems from the treatment of individuals inside concept terms, which can, however, again be characterized with the help of a non-standard semantics.

Since 1988, a new type of algorithms, so-called *tableau-based algorithms*, for reasoning in description logics has been developed. Here, the logical point of view is not only used to define the semantics, but also for the design of algorithms, that is, the inference problems are considered as deduction problems in logics. In principle, these algorithms are methods for generating finite models. They can be seen as specializations of the tableau calculus for first-order predicate logic [26, 170]. The non-standard syntax of DL and the restricted expressiveness of these logics allows to design terminating procedures, i.e., for many description languages one obtains decision procedures for the relevant inference problems (see [15] for an introductory exposition of tableau-based inference methods in DL). The first tableau-based subsumption algorithm was developed in [166] for the language $\mathcal{ALC}$, which allows for the first five constructors of Figure 2. Since then, this approach for designing subsumption algorithms was extended to the instance problem [93, 15] and to various description languages extending $\mathcal{ALC}$ (see, e.g., [95, 94, 20, 21, 8] for languages with number restrictions; [6] for transitive closure of roles and [159, 99, 101] for transitive roles; [12, 89, 87, 22] for constructs that allow to refer to concrete domains such as numbers; and [10, 40, 8] for the treatment of general axioms of the form $C \doteq D$, where C, D may both be complex concept terms).

Undecidability and complexity results. Other important research contributions for DL are concerned with the decidability and the complexity of the subsumption problem in different DL languages. It has turned out that the languages used in early DL systems were too expressive, which led to undecidability of the subsumption problem [165, 147]. More recent undecidability results for extensions of $\mathcal{ALC}$ can be found in [13, 89, 20, 21, 87, 22].

The first complexity results [115, 139] showed that, even for very small languages, there cannot exist subsumption algorithms that are both complete and polynomial. In the meantime, the worst-case complexity of the subsumption problem in a large class of DL languages, the so-called $\mathcal{AL}$-family, has (almost completely) been determined [59, 58, 57]. With the exception of a few polynomially decidable languages, the complexity results range between NP or coNP and PSPACE. Whereas these results are given with respect to an empty TBox (i.e., they consider subsumption of concept terms with respect to all interpretations), Nebel [142] has shown that the expansion of TBox definitions may lead to an exponential blow-up, which may result in a larger complexity (coNP instead of polynomial) for certain languages. In the presence of cyclic TBox definitions, so-called terminological cycles, the subsumption problem becomes PSPACE-complete even for these small languages [140, 5, 7, 112]. The use of general inclusion axioms (in $\mathcal{ALC}$) even causes the subsumption problem to become ExpTime-complete [163]. It has also been shown that for certain languages the instance problem can be harder than the subsumption problem [160, 61].

Optimizations. Considering these complexity results, one may ask whether incomplete, but polynomial algorithms should be preferred over the complete ones, which are necessarily of high worst-case complexity. First experiences [11, 36, 98] with implemented systems using complete algorithms show, however, that on realistic KBs the run time is comparable to that of Classic and Loom (i.e., mature systems using incomplete algorithms). These positive results depend on the use of sophisticated optimization techniques. Whereas [11] concentrated mostly on reducing the number of subsumption tests during classification, more recent work in this direction is concerned with optimizing the subsumption algorithm itself [83, 97, 98, 82, 100, 104].

Connections with other logical formalisms. Before we turn to the connection between DL and Logic Programming, we should like to mention several interesting connections between DL and more traditional areas of logics.

Schild [161] was the first to observe that the language $\mathcal{ALC}$ is a syntactic variant of the propositional multi-modal logic $\mathbf{K}_n$ (see next section), and that the extension of $\mathcal{ALC}$ by transitive closure of roles [6] corresponds to propositional dynamic logic (PDL) [67, 145]. In particular, the algorithms used in modal logics for deciding satisfiability are very similar to the tableau-based algorithms newly developed for DL languages. This connection between DL and modal logics has been used to transfer techniques and results from modal logics and propositional dynamic logic to DL [162, 163, 77–79, 47]. Instead of using tableau-based algorithms, decidability of certain propositional modal logics (and thus of the corresponding DL), can also be shown by establishing the finite model property (see, e.g., [68], Section 1.14) of the logic (i.e., showing that a formula/concept is satisfiable iff it is satisfiable in a finite interpretation) or by employing tree automata (see, e.g, [174]). It should be noted, however, that some of the very expressive DL languages considered in this context (e.g., the language $\mathcal{CIQ}$ introduced in [79]) no longer satisfy the finite model property. For these languages, reasoning with respect to finite models (which is, for example, of interest for database applications) differs from reasoning with respect to all models [43].

Given the translation of DL into first-order predicate logic mentioned above, one might ask whether general first-order theorem provers can be employed for reasoning in DL. In general, this approach will only yield semidecision procedures for DL inference problems such as subsumption. By employing appropriate translation techniques and resolution strategies, general purpose resolution provers can, however, be used to obtain decision procedures for subsumption in the language $\mathcal{ALC}$ [65, 103, 164].

Decidability of the inference problems for $\mathcal{ALC}$ can also be obtained as a consequence of the known decidability result for the two variable fragment of first-order predicate logic. The language L_2 consists of all formulae of first-order predicate logic that can be built with the help of predicate symbols (including equality) and constant symbols (but without function symbols) using only the variables x, y. Decidability of L_2 has been shown in [137]. More precisely, satisfiability of L_2-formulae is a NEXPTIME-complete problem [85]. It is easy to see that, by appropriately re-using variable names, any concept term of the language $\mathcal{ALC}$ can be translated into an L_2-formula with one free variable. A direct translation of the concept term $\forall R.(\exists R.A)$ yields the formula $\forall y{:}(R(x,y) \rightarrow (\exists z{:}(R(y,z) \wedge A(z))))$. Since the subformula $\exists z{:}(R(y,z) \wedge A(z))$ does not contain x, this variable can be re-used: renaming the bound variable z into x yields the equivalent formula $\forall y{:}(R(x,y) \rightarrow (\exists x{:}(R(y,x) \wedge A(x))))$, which uses only two variables (see [28] for details). This connection between $\mathcal{ALC}$ and L_2 shows that any extension of $\mathcal{ALC}$ by constructors that can be expressed with the help of only two variables yields a decidable DL. Number restrictions and composition of roles are examples of constructors that cannot be expressed within L_2. Number restrictions can, however, be expressed in C_2, the extension of L_2 by counting quantifiers, which has recently been shown to be decidable [86, 144].

Another distinguishing features of the formulae obtained by translating $\mathcal{ALC}$-concept terms into first-order predicate logic is that quantifiers are used only in a very restricted way: the quantified variables are always "guarded" by role expression. For example, in the formula $\forall y{:}(R(x,y) \rightarrow C(y))$, which is obtained as translation of the concept term $\forall R.C$, the quantified variable y is guarded by $R(x,y)$. For this reason, the formulae obtained as translations of $\mathcal{ALC}$-concept terms belong to the so-called *guarded fragment* GF of first-order predicate logic [2], which has the following inductive definition:

– Every atomic formula belongs to GF.
– GF is closed under the Boolean connectives.
– If $\boldsymbol{x}, \boldsymbol{y}$ are tuples of variables, R is a predicate symbol, and ψ is a formula of GF such that every free variable in ψ occurs in $\boldsymbol{x}, \boldsymbol{y}$, then the formulae
  - $\forall \boldsymbol{y}{:}(R(\boldsymbol{x}, \boldsymbol{y}) \rightarrow \psi)$
  - $\exists \boldsymbol{y}{:}(R(\boldsymbol{x}, \boldsymbol{y}) \wedge \psi)$

  also belong to GF.

It should be noted that there is no restriction to unary and binary predicate symbols, and that the free variables of the formula ψ in the third item above are not restricted to the quantified variables $\boldsymbol{y}$. Thus, GF is considerably more expressive than the fragment obtained by translating $\mathcal{ALC}$-concept terms into first-order predicate logic. In [2] it is shown that GF nevertheless has the finite model property, which implies that satisfiability of formulae in GF is decidable. More precisely, the satisfiability problem for GF is complete for double exponential time [84]. Decidability of GF can also be shown with the help of resolution methods [48].

Connection with Logic Programming Since Logic Programming languages are computationally complete and DL languages are usually decidable, one may say that DL languages have less expressive power. If we consider Logic Programming languages as representation languages rather than as programming languages, then one observes that several of the DL constructors cannot be expressed. In fact, disjunction and existential restrictions allow for incompletely specified knowledge. For example, the term

$$(\exists \mathtt{pet}.(\mathtt{Dog} \sqcup \mathtt{Cat}))(\mathtt{BILL})$$

leaves it open which ABox individual is Bill's pet, and whether it is a cat or a dog. Such an "under-specification" is not possible with the Horn clauses of traditional Logic Programming languages. To overcome this deficit, extensions of Logic Programming languages by disjunction and classical negation (in contrast to "negation as failure") have been introduced [76, 151, 119, 23, 35]. However, these extensions treat only some aspects of these constructors: for example, the "classical negation" in these approaches only represents the aspect that a set and its complement are disjoint; the fact that the union of a set with its complement is the whole universe is not taken into account. The integration of Description Logics with a rule calculus that is able to express Horn rules has been investigated in [90]. Other work in this direction can, for example, be found in [117, 116].

3 Modal Logics

This is an area of logics that has been investigated for quite a while, and for which a great variety of methods and results are available. In the following, we give a very short introduction, which emphasizes the connection between Description Logics and Modal Logics. For more detailed introductions and overviews of the area we refer the reader to [44, 102, 88, 68].

The propositional multi-modal logic $\mathbf{K}_n$ extends propositional logic by n pairs of unary operators, which are called box and diamond operators. The $\mathbf{K}$ stands for the basic modal logic on which most modal logics are build (see below), and "multi-modal" means that one considers more than one pair of box and diamond operators. Depending on the intended application, these operators may have different intuitive meanings. For example, if we want to represent time-dependent knowledge, then we can use the diamond operator $\langle \mathit{future} \rangle$ and the box operator $[\mathit{future}]$, where the intended meaning of a formula $\langle \mathit{future} \rangle \phi$ is "Sometime in the future, ϕ holds," whereas $[\mathit{future}]\phi$ is meant to express "Always in the future, ϕ holds." If we want to represent knowledge about the beliefs of intelligent agents, then we can use the operators $\langle \mathit{robi1} \rangle$ and $[\mathit{robi1}]$, where $[\mathit{robi1}]\phi$ should be interpreted as "Robot 1 believes that ϕ holds," and $\langle \mathit{robi1} \rangle \phi$ as "Robot 1 believes that ϕ is possible." The different meanings of the operators are taken into account by using additional axioms or by imposing additional restrictions on the semantic structures (see below).

Syntax and semantics. First, we consider the base logic $\mathbf{K}_n$ in a bit more detail. *Formulae* of this logic are built from atomic propositions p and n different modal parameters m using the Boolean connectives $\wedge, \vee, \neg$[5] and the modal operators $[m]$ and $\langle m\rangle$. For example,

$$[robi1]\langle future\rangle(p \wedge \langle robi2\rangle\neg p)$$

is a formula of $\mathbf{K}_n$, which could be interpreted as saying "Robot 1 believes that, sometime in the future, p will hold, while at the same time Robot 2 will believe that $\neg p$ is possible." Here *robi1*, *robi2*, and *future* are modal parameters, and p is an atomic proposition.

In order to define the semantics of $\mathbf{K}_n$, we consider so-called *Kripke structures* $\mathcal{K} = (\mathcal{W}, \mathcal{R})$, which consist of a set of possible worlds $\mathcal{W}$ and a set $\mathcal{R}$ of transition relations. Each possible world $I \in \mathcal{W}$ corresponds to an interpretation of propositional logic, i.e., it assigns a truth value $p^I \in \{0,1\}$ to every atomic proposition p. The set $\mathcal{R}$ contains for every modal parameter m a transition relation $R_m \subseteq \mathcal{W} \times \mathcal{W}$. Validity of a $\mathbf{K}_n$-formula ϕ in the world I of a Kripke structure $\mathcal{K}$ is defined by induction on the structure of $\mathbf{K}_n$-formulae:

- $\mathcal{K}, I \models p$ iff $p^I = 1$ (for atomic propositions p).
- $\mathcal{K}, I \models \phi \wedge \psi$ iff $\mathcal{K}, I \models \phi$ and $\mathcal{K}, I \models \psi$. (The semantics of the other Boolean operators is also defined in the usual way.)
- $\mathcal{K}, I \models [m]\phi$ iff $\mathcal{K}, J \models \phi$ holds for *all* J such that $(I, J) \in R_m$.
- $\mathcal{K}, I \models \langle m\rangle\phi$ iff $\mathcal{K}, J \models \phi$ holds for *some* J with $(I, J) \in R_m$.

The $\mathbf{K}_n$-formula ϕ is *valid* iff $\mathcal{K}, I \models \phi$ holds for all Kripke structures $\mathcal{K}$ and all worlds I in $\mathcal{K}$.

Connection with Description Logics. This definition of the semantics for $\mathbf{K}_n$ looks very similar to the semantics for DL languages. Concept terms C of $\mathcal{ALC}$ can directly be translated into formulae ϕ_C of $\mathbf{K}_n$ by interpreting concept names as atomic propositions and role names as modal parameters. The Boolean connectives of $\mathcal{ALC}$ are simply replaced by the corresponding Boolean connectives of $\mathbf{K}_n$. Universal role restrictions (value restrictions) are replaced by the corresponding box operator, and existential role restrictions are replaced by the corresponding diamond operator. For example, the concept term $\forall R.A \sqcap \exists S.\neg A$ yields the $\mathbf{K}_n$-formula $[R]A \wedge \langle S\rangle\neg A$.

There is an obvious 1–1 correspondence between Kripke structures $\mathcal{K}$ and interpretations $\mathcal{I}$ of $\mathcal{ALC}$: the domain $\Delta^{\mathcal{I}}$ of $\mathcal{I}$ corresponds to the set of possible worlds $\mathcal{W}$, the interpretation $S^{\mathcal{I}}$ of the role name S corresponds to the transition relation R_S, and the interpretation $A^{\mathcal{I}}$ of the concept name A corresponds to the set of worlds in which A has truth value 1. It is easy to show (by induction on the structure of concept terms) that this correspondence also holds for complex

[5] Additional Boolean connectives like implication can, as usual, be introduced as abbreviations.

concept terms C: if $\mathcal{I}$ is an interpretation of $\mathcal{ALC}$ and $\mathcal{K}$ is the corresponding Kripke structure, then $C^{\mathcal{I}} = \{I \mid \mathcal{K}, I \models \phi_C\}$.

In particular, this implies that $C \sqsubseteq_{\emptyset} D$ iff $\phi_C \rightarrow \phi_D$ is valid. This shows that decision procedures for validity in $\mathbf{K}_n$ can be used to decide the subsumption problem in $\mathcal{ALC}$. One should note, however, that this observation does not yield decision procedures for ABox reasoning, or for reasoning in DL languages with number restrictions. More recent work on modal logics with "graded modalities" [66, 64, 172] (which correspond to number restrictions) and "nominals" [72] (which correspond to ABox individuals) did not focus on decidability issues. The extension of results from propositional modal and dynamic logic to logics allowing for number restrictions and individuals was addressed in [77, 46, 79].

Axiomatizations. If one wants to assign the modal operators with a specific meaning (like "knowledge of an intelligent agent" or "in the future"), then using the basic modal logic $\mathbf{K}_n$ is not sufficient since it does not model the specific properties that modal operators with this interpretation should satisfy. In order to describe these properties, one can use an axiomatic approach. Figure 3

All propositional tautologies	**Taut**	axiom schemata
$[m](\phi \rightarrow \psi) \rightarrow ([m]\phi \rightarrow [m]\psi)$	**K**	
$[m]\phi \rightarrow \phi$	**T**	
$[m]\phi \rightarrow [m][m]\phi$	**4**	
$\neg[m]\phi \rightarrow [m]\neg[m]\phi$	**5**	
$\phi \rightarrow \psi$ and ϕ yield ψ	**modus ponens**	inference rules
ϕ yields $[m]\phi$	**necessitation**	

Fig. 3. Axiom schemata and inference rules for modal logics.

introduces some axiom schemata and the corresponding inference rules modus ponens and necessitation. These are axiom schemata rather than axioms since ϕ and ψ may be substituted by arbitrary modal formulae.

The basic modal logic $\mathbf{K}_n$ is characterized by the axiom schemata **Taut** and **K** in the following sense: a formula ϕ of $\mathbf{K}_n$ is valid (i.e., holds in all worlds of all Kripke structures) iff it can be derived from instances of **Taut** and **K** using modus ponens and necessitation (see [88] for a proof). The other three schemata describe possible properties of modal operators that express "knowledge of intelligent agents," i.e., in this interpretation $[m]\phi$ should be read as "agent m knows ϕ." Thus, **T** can intuitively be read as "An intelligent agent does not have incorrect knowledge," or more precisely: "If agent m knows ϕ in a situation, then ϕ holds in this situation." This property distinguishes knowledge from belief: an agent may very well believe in incorrect facts, but it cannot know them. The axiom schema **4** describes positive introspection, i.e., "An intelligent agent knows what it knows," whereas axiom schema **5** describes negative introspection, i.e., "An intelligent agent knows what it does not know." While **T** and **4** are generally accepted as reasonable axioms for knowledge, **5** is disputable: negative introspection implies that the agent can asses its own competence in the sense that it knows when its own knowledge is not sufficient to solve a certain task. Consequently, there are two different modal logics that model knowledge of intelligent agents. The logic **S4** dispenses with negative introspection, i.e., it uses only the schemata **Taut**, **K**, **T**, **4**, whereas **S5** additionally allows for **5**.

Properties of transition relations. As an alternative to this axiomatic approach for defining **S4** and **S5**, one can also characterize these logics in a semantic way by restricting the admissible transition relations in Kripke structures [88]. The logic **S4** corresponds to the restriction to *reflexive and transitive* transition relations, i.e., the formula ϕ holds in all worlds of all Kripke structures with reflexive and transitive transition relations iff it can be derived from instances of **Taut**, **K**, **T** and **4** using modus ponens and necessitation. Analogously, **S5** corresponds to the restriction to transition relations that are *equivalence relations*.

These correspondences can be used to design tableau algorithms for **S4** and **S5** by integrating the special properties of the transition relations into the rules of the tableau calculus. It should be noted that a naive integration of a tableau rule for transitivity would lead to a nonterminating procedure. This problem can, however, be overcome by testing for cyclic computations. This yields a PSPACE-algorithm for **S4**. For a single **S5**-modality, the satisfiability problem is "only" NP-complete, whereas the problem becomes PSPACE-complete if more than one such modality is available (see [88] for detailed proofs of these and other complexity results).

The properties of modal operators that model time-dependent knowledge have also been investigated in detail (see, e.g., [63, 171] for overview articles on this topic).

Integration with Description Logics. In order to represent time-dependent and subjective knowledge in Description Logics, one can integrate modal operators into DL languages. Because of the above mentioned close connection between DL and modal logics, such an integrations appeared to be rather simple. It has turned out, however, that it is more complex than expected, both from the semantic and the algorithmic point of view [113, 19, 18, 162, 175]. It should be noted that some of these combined languages [18, 175] are first-order modal logics rather than propositional modal logics.

Connection with Logic Programming. We close this section by mentioning some work that is concerned with the connection between modal logics and logic programming. Gelfond [73] extends Disjunctive Logic Programs to Epistemic Logic Programs by introducing modal operators **K** and **M**: for a literal L, the expression $\mathbf{K}L$ should be read as "L is known" and $\mathbf{M}L$ should be read as "L may be assumed." Giordano and Martelli [81] use modal logic to obtain a uniform representation of different approaches for structuring Logic Programs with the help of blocks and modules. In [49, 138, 1, 70], modal or temporal Logic Programming languages are introduced.

4 Nonmonotonic Logics

This research area has also created a huge number of approaches and results, which we cannot describe in detail here. The following is a brief introduction into the existing approaches and the problems treated by these approaches. Overviews of this research area can, for example, be found in [71, 38]. In addition, there are several monographs on the topic [37, 120, 126, 3]. An annotated collection of influential papers in the area can be found in [80].

Motivation. Knowledge representation languages based on classical logics (e.g., first-order predicate logic) are monotonic in the following sense: if a statement ϕ can be derived from a knowledge base, then ϕ can also be derived from any larger knowledge base. This property has the advantage that inferences once drawn need not be revised when additional information comes in. However, this leads to the disadvantage that adding information contradicting one of the drawn inferences leads to an inconsistency, and thus makes the knowledge base useless. In many applications, the knowledge about the world (the application domain) is represented in an incomplete way. Nevertheless, one wants to draw *plausible* conclusions from the available knowledge, that is, inferences that are not justified by reasoning in classical logic, but are plausible considering the available knowledge. In this situation, newly acquired information may show that some of these plausible conclusions were wrong. This should not lead to inconsistency of the knowledge base, but rather to a withdrawal of some of the plausible conclusions.

Let us start with three simple examples that illustrate the situations in which nonmonotonic inference methods are desirable. *Default rules* apply to most individuals (resp. typical or normal individuals), but not to all. As an example, we may consider property edges in Semantic Networks. In the network of Figure 1 we had the default rule "Frogs are *normally* green." This rule should be applied whenever there is no information contradicting it. As long as we only know that Kermit is a frog, we deduce that Kermit is green. The rule is not applied to grass frogs, of which it is known that they are not green (since they are brown).

The *Closed World Assumption (CWA)* [156] assumes by default that the available information is complete. If an assertion cannot be derived (using classical inference methods) from the knowledge base, then CWA deduces its negation. This assumption is, for example, employed in relational databases and in Logic Programming languages with "Negation as Failure" [45, 168]. As an example, we may consider train connections in a timetable: if a connection is not contained in the timetable, we conclude that it does not exist. If we learn later on that there is such a connection (which was not contained in the timetable), then we must withdraw this conclusion.

The *Frame problem* [130, 39] comes from the domain of modelling actions. In this context, it is important to describe which properties are changed and which remain unchanged by an application of the action. Usually, the application of an action changes very few aspects of the world. For example, by sending a manuscript of this article to a publisher, I have changed its location, but (hopefully) not its content. The so-called "frame axioms" describe which properties remain unchanged by the application of an action. Since there usually is a very large number of these axioms, one should try to avoid having to state them explicitly. A possible solution to this problem is to employ a nonmonotonic inference rule, which says that (normally) all aspects of the world that are not explicitly changed by the action remain invariant under its application.

Four approaches to nonmonotonic reasoning. In the literature, a great variety of different approaches to nonmonotonic reasoning has been introduced, of which none seems to be "the best" approach. A very positive development is the fact that recently several results clarifying the connection between different approaches have been obtained (see, e.g., [105, 108, 109, 118, 111]). In the following, we briefly introduce the four most important types of approaches.

Consistency-based approaches, of which Reiter's Default Logic [157] is a typical example, consider nonmonotonic rules of the form "A normally implies B." Such a rule can be applied (i.e., B is inserted into the knowledge base), if A holds, and inserting B does not destroy consistency of the knowledge base. As an example, we may again use the default rule "Frogs are normally green." Let us first assume that the knowledge base contains the information: "Grass frogs are brown. An individual cannot be both brown and green. Grass frogs are frogs. Kermit is a frog. Scooter is a grass frog." In this case, the default rule can be applied to Kermit, but not to Scooter. The major problem that a consistency-based approach must solve is the question of how to resolve conflicts between different rules: applying default rules in different order may lead to different results. To illustrate this problem, let us assume that the strict information "Grass frogs are brown" is replaced by the defeasible information "Grass frogs are normally brown," whereas all the other information remains the same. Now, both default rules are applicable to Scooter. As soon as one of them is applied, the other one is no longer applicable. Thus, one must decide which of the possible results should be preferred, or, if the conflict cannot be resolved due to priorities among rules, how much information can still be deduced from such an unresolved situation. For example, the conclusions concerning Kermit should not be influenced by the conflict for Scooter.

A *modal nonmonotonic logic* was, for example, proposed by McDermott and Doyle [132, 131]. This logic allows for a diamond operator, which is written as $\mathbf{M}$, where $\mathbf{M}\phi$ should intuitively be read as "ϕ is consistent." The default rule "Frogs are normally green" can then be written as the implication $\mathtt{Frog} \wedge \mathbf{M}\mathtt{Green} \rightarrow \mathtt{Green}$. In order to treat this implication correctly, the logic needs an additional inference rule, which is able to derive formulae of the form $\mathbf{M}\phi$ according to their intended semantics. In principle, this inference rule allows us to deduce $\mathbf{M}\phi$, if $\neg\phi$ cannot be deduced. This is not as simple as it may sound since we are faced with the following cyclic dependency: which formulae of the form $\neg\phi$ are deducible already depends on the inference rule to be defined. Doyle and McDermott solve this problem by introducing an appropriated fixed-point semantics. Another representative of the class of modal nonmonotonic logics is the well-known autoepistemic logic [136].

In classical predicate logic, the notion of logical consequence is defined with respect to *all models* of a set of formulae. *Preferential semantics* [169] takes as logical consequences all the formulae that hold in *all preferred models*. The preferred models are usually defined as the minimal models with respect to a given *preference relation* on interpretations. To illustrate this idea, we consider a simple example from propositional logic. If we assume that there are only two propositional variables p, q, then we have four different interpretations: $I_{\neg p,\neg q}$, in which p and q are false, $I_{p,\neg q}$, in which p is true and q is false, $I_{\neg p,q}$, in which p is false and q is true, and $I_{p,q}$, in which p and q are true. Assume that the preference relation $<$ is given by

$$I_{\neg p,\neg q} < I_{p,q},\quad I_{\neg p,\neg q} < I_{\neg p,q},\quad I_{p,q} < I_{p,\neg q},\quad \text{and}\quad I_{\neg p,q} < I_{p,\neg q},$$

where the smaller interpretation is preferred over the larger one. The empty set of formulae $\emptyset$ has all interpretations as model, which means that $I_{\neg p,\neg q}$ is the only minimal model. In particular, $\neg p$ is a consequence of $\emptyset$ with respect to this preferential semantics. The set $\{q\}$ excludes the two interpretations $I_{p,\neg q}$ and $I_{\neg p,\neg q}$. The remaining models, $I_{p,q}, I_{\neg p,q}$ are incomparable w.r.t. $<$, and are thus both minimal models. Since $\neg p$ does not hold in both, it can no longer be deduced. This shows that preferential semantics yields a nonmonotonic formalism. An important example for preferential semantics for the case of predicate logic is *circumscription* [129]. Here, the goal is to minimize the extension of a given predicate P, i.e., an interpretation $\mathcal{I}$ is preferred over an interpretation $\mathcal{J}$, if $P^{\mathcal{I}} \subset P^{\mathcal{J}}$ holds. Default rules like "Frogs are normally green" can then be expressed with the help of an "abnormality predicate," which is minimized:

$$\mathtt{Frog(x)} \wedge \neg\mathtt{Ab(x)} \rightarrow \mathtt{Green(x)}.$$

Exceptions to this rule can now be introduced by implications like $\mathtt{Brown(x)} \rightarrow \mathtt{Ab(x)}$. The fact that `Ab` is minimized makes sure that the default rule is applied to an individual unless it is a known exception.

Properties that a "reasonable" *nonmonotonic inference relation* $\mid\!\sim$ should satisfy were, on the one hand, introduced for the purpose of comparing and evaluating different approaches to nonmonotonic reasoning [69, 124]. On the other hand, these properties can also be interpreted as inference rules (like modus ponens), which can be used to generate new nonmonotonic consequences [114]. Figure 4 gives several examples of such reasonable properties. It has also turned out that there is a close connection between preferential semantics and inferences relations satisfying certain properties [123, 111].

$\phi \mathrel{|\!\sim} \phi$		**Reflexivity**
If $\phi \mathrel{|\!\sim} \psi$ and ϕ equivalent to ϕ'	then $\phi' \mathrel{|\!\sim} \psi$	**Left equivalence**
If $\phi \mathrel{|\!\sim} \psi$ and ψ implies ψ'	then $\phi \mathrel{|\!\sim} \psi'$	**Right weakening**
If $\phi \wedge \phi' \mathrel{|\!\sim} \psi$ and $\phi \mathrel{|\!\sim} \phi'$	then $\phi \mathrel{|\!\sim} \psi$	**Cut**
If $\phi \mathrel{|\!\sim} \phi'$ and $\phi \mathrel{|\!\sim} \psi$	then $\phi \wedge \phi' \mathrel{|\!\sim} \psi$	**Cautious monotony**

Fig. 4. Properties of nonmonotonic inference relations.

Connection with Logic Programming A similar approach for evaluating and comparing the semantics for Logic Programs was used in [51, 52]. As mentioned above, the "Closed World Assumption" in Logic Programs, and the corresponding treatment of negation as "Negation as Failure," leads to a nonmonotonic behaviour of Logic Programs. Thus, it is not surprising that there is a close connection between approaches for defining declarative semantics for (extended) logic programs (e.g., [74, 75, 27, 173]) and formalisms for nonmonotonic reasoning. In principle, these semantics depend on a preference relation between models. Their development was strongly influenced by the semantics for autoepistemic logic and for default logic. A first overview of these connections is given in [150]. More recent work in this direction can be found in the proceedings of the conferences "Non-Monotonic Extensions of Logic Programming" [55, 56] and "Logic Programming and Nonmonotonic Reasoning" [125, 54].

Connection with Description Logics The integration of default rules into Description Logics was investigated in [153, 16, 17]. In [60], an epistemic operator **K** is added to the DL $\mathcal{ALC}$. This operator is similar to the modal operators employed in modal nonmonotonic logic, and it can, for example, be used to impose a "local" closed world assumption.

References

1. M. Abadi and Z. Manna. Temporal logic programming. *Journal of Symbolic Computation*, 8:277–295, 1989.
2. H. Andréka, J. van Benthem, and I. Németi. Modal languages and bounded fragments of predicate logic. Research Report ML-96-03, ILLC, Amsterdam, 1996.
3. G. Antoniou. *Nonmonotonic Reasoning.* MIT Press, Cambridge, Mass., 1997.
4. K. R. Apt and R. N. Bol. Logic programming and negation: A survey. *Journal of Logic Programming*, 19–20:9–71, 1994.
5. F. Baader. Terminological cycles in KL-ONE-based knowledge representation languages. In T. Dietterich and W. Swartout, editors, *Proceedings of the 8th National Conference on Artificial Intelligence*, pages 621–626, Boston, Mass., 1990. AAAI Press / MIT Press.
6. F. Baader. Augmenting concept languages by transitive closure of roles: An alternative to terminological cycles. In J. Mylopoulos and R. Reiter, editors,

Proceedings of the 12th International Joint Conference on Artificial Intelligence, pages 446–451, Sydney, Australia, 1991. Morgan Kaufmann, San Francisco.

7. F. Baader. Using automata theory for characterizing the semantics of terminological cycles. *Annals of Mathematics and Artificial Intelligence*, 18(2–4):175–219, 1996.
8. F. Baader, M. Buchheit, and B. Hollunder. Cardinality restrictions on concepts. *Artificial Intelligence*, 88(1–2):195–213, 1996.
9. F. Baader, H.-J. Bürckert, J. Heinsohn, J. Müller, B. Hollunder, B. Nebel, W. Nutt, and H.-J. Profitlich. Terminological knowledge representation: A proposal for a terminological logic. DFKI Technical Memo TM-90-04, Deutsches Forschungszentrum für Künstliche Intelligenz, Kaiserslautern, 1990.
10. F. Baader, H.-J. Bürckert, B. Hollunder, W. Nutt, and J. Siekmann. Concept logics. In J. W. Lloyd, editor, *Proceedings of the Symposium on Computational Logic*, pages 177–201, Brussels, 1990. Springer–Verlag.
11. F. Baader, E. Franconi, B. Hollunder, B. Nebel, and H.-J. Profitlich. An empirical analysis of optimization techniques for terminological systems. *Journal of Applied Intelligence*, 4:109–132, 1994.
12. F. Baader and P. Hanschke. A scheme for integrating concrete domains into concept languages. In J. Mylopoulos and R. Reiter, editors, *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, pages 452–457, Sydney, Australia, 1991. Morgan Kaufmann, San Francisco.
13. F. Baader and P. Hanschke. Extensions of concept languages for a mechanical engineering application. In *Proceedings of the 16th German AI-Conference, GWAI-92*, volume 671 of *Lecture Notes in Computer Science*, pages 132–143, Bonn (Germany), 1993. Springer–Verlag.
14. F. Baader and B. Hollunder. $\mathcal{KRIS}$: $\mathcal{K}$nowledge $\mathcal{R}$epresentation and $\mathcal{I}$nference $\mathcal{S}$ystem. *SIGART Bulletin*, 2(3):8–14, 1991.
15. F. Baader and B. Hollunder. A terminological knowledge representation system with complete inference algorithms. In M. Richter and H. Boley, editors, *Proceedings of the First International Workshop on Processing Declarative Knowledge (PDK-91)*, volume 567 of *Lecture Notes in Computer Science*, pages 67–85, Kaiserslautern (Germany), 1991. Springer–Verlag.
16. F. Baader and B. Hollunder. Embedding defaults into terminological representation systems. *J. Automated Reasoning*, 14:149–180, 1995.
17. F. Baader and B. Hollunder. Priorities on defaults with prerequisites, and their application in treating specificity in terminological default logic. *J. Automated Reasoning*, 15:41–68, 1995.
18. F. Baader and A. Laux. Terminological logics with modal operators. In C. Mellish, editor, *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pages 808–814, Montréal, Canada, 1995. Morgan Kaufmann, San Francisco.
19. F. Baader and H.-J. Ohlbach. A multi-dimensional terminological knowledge representation language. *Journal of Applied Non-Classical Logics*, 5(2):153–197, 1995.
20. F. Baader and U. Sattler. Description logics with symbolic number restrictions. In W. Wahlster, editor, *Proceedings of the Twelfth European Conference on Artificial Intelligence (ECAI-96)*, pages 283–287. John Wiley & Sons Ltd, 1996.
21. F. Baader and U. Sattler. Number restrictions on complex roles in description logics. In L. C. Aiello, J. Doyle, and S. Shapiro, editors, *Proceedings of the Fifth*

International Conference on the Principles of Knowledge Representation and Reasoning (KR-96), pages 328–339, Cambridge, Mass., 1996. Morgan Kaufmann, San Francisco.
22. F. Baader and U. Sattler. Description logics with concrete domains and aggregation. In H. Prade, editor, *Proceedings of the 13th European Conference on Artificial Intelligence (ECAI-98)*, pages 336–340. John Wiley & Sons Ltd, 1998.
23. C. Baral and M. Gelfond. Logic programming and knowledge representation. *Journal of Logic Programming*, 19–20:73–148, 1994.
24. S. Bergamaschi and D. Beneventano. Incoherence and subsumption for recursive views and queries in object-oriented data models. *Data and Knowledge Engineering*, 21(3):217–252, 1997.
25. S. Bergamaschi, C. Sartori, D. Beneventano, and M. Vincini. ODB-tools: a description logics based tool for schema validation and semantic query optimization in object oriented databases. In M. Lenzerini, editor, *Proceedings of the 5th Congress of the Italian Association for Artificial Intelligence: Advances in Artificial Intelligence (AI*IA-97)*, volume 1321 of *Lecture Notes in Artificial Intelligence*, pages 435–438. Springer–Verlag, 1997.
26. E. W. Beth. *The Foundations of Mathematics*. North-Holland, Amsterdam, 1959.
27. N. Bidoit and C. Froidevaux. Negation by default and unstratified logic programs. *Theoretical Computer Science*, 78:85–112, 1991.
28. A. Borgida. On the relative expressiveness of description logics and predicate logics. *Journal of Artificial Intelligence*, 82(1–2):353–367, 1996.
29. A. Borgida, R. J. Brachman, D. L. McGuinness, and L. A. Resnick. CLASSIC: A structural data model for objects. In *Proceedings of the 1989 ACM SIGMOD International Conference on Management of Data*, pages 59–67, Portland, Oreg., 1989.
30. A. Borgida and P. Patel-Schneider. A semantics and complete algorithm for subsumption in the CLASSIC description logic. *Journal of Artificial Intelligence Research*, 1:277–308, 1994.
31. R. J. Brachman. What's in a concept: Structural foundations for semantic networks. *International Journal of Man-Machine Studies*, 9:127–152, 1977.
32. R. J. Brachman. Structured inheritance networks. In W. A. Woods and R. J. Brachman, editors, *Research in Natural Language Understanding*, Quarterly Progress Report No. 1, BBN Report No. 3742, pages 36–78. Bolt, Beranek and Newman Inc., Cambridge, Mass., 1978.
33. R. J. Brachman, D. L. McGuinness, P. F. Patel-Schneider, L. A. Resnick, and A. Borgida. Living with CLASSIC: When and how to use a KL-ONE-like language. In J. Sowa, editor, *Principles of Semantic Networks*, pages 401–456. Morgan Kaufmann, San Francisco, 1991.
34. R. J. Brachman and J. G. Schmolze. An overview of the KL-ONE knowledge representation system. *Cognitive Science*, 9(2):171–216, 1985.
35. S. Brass and J. Dix. Disjunctive semantics based upon partial and bottom-up evaluation. In L. Sterling, editor, *Proceedings of the 12th International Conference on Logic Programming*, pages 199–213. MIT Press, 1995.
36. P. Bresciani, E. Franconi, and S. Tessaris. Implementing and testing expressive description logics: A preliminary report. In *Proceedings of the International Symposium on Knowledge Retrieval, Use, and Storage for Efficiency, KRUSE-95, Santa Cruz, USA*, 1995.
37. G. Brewka. *Nonmonotonic Reasoning*. Cambridge University Press, Cambridge, UK, 1991.

38. G. Brewka, J. Dix, and K. Konolige. *Nonmonotonic Reasoning: An Overview.* CSLI Publications, Center for the Study of Language and Information, Stanford, Cal., 1997.
39. F. M. Brown, editor. *The Frame Problem in Artificial Intelligence: Proceedings of the 1987 Workshop.* Morgan Kaufman, 1987.
40. M. Buchheit, F. M. Donini, and A. Schaerf. Decidable reasoning in terminological knowledge representation systems. *Journal of Artificial Intelligence Research*, 1:109–138, 1993.
41. M. Buchheit, M. Jeusfeld, W. Nutt, and M. Staudt. Subsumption of queries to object-oriented databases. *Information Systems*, 19(1):33–54, 1994.
42. M. Buchheit, R. Klein, and W. Nutt. Configuration as model construction: The constructive problem solving approach. In *Proceedings of the Third International Conference on Artificial Intelligence in Design, AID'94*, Lausanne, Switzerland, 1994.
43. D. Calvanese. Finite model reasoning in description logics. In L. C. Aiello, J. Doyle, and S. Shapiro, editors, *Proceedings of the Fifth International Conference on the Principles of Knowledge Representation and Reasoning (KR-96)*, pages 292–303, Cambridge, Mass., 1996. Morgan Kaufmann, San Francisco.
44. B. F. Chellas. *Modal Logic: An Introduction.* Cambridge University Press, Cambridge, UK, 1980.
45. K. Clark. Negation as Failure. In *Logic and Data Bases*, pages 293–322. Plenum Press, New York, 1978.
46. G. De Giacomo. *Decidability of Class-Based Knowledge Representation Formalisms.* PhD thesis, Università degli Studi di Roma "La Sapienza", Italy, 1995.
47. G. De Giacomo. Eliminating "converse" from converse PDL. *Journal of Logic, Language, and Information*, 5:193–208, 1996.
48. H. de Nivelle. A resolution decision procedure for the guarded fragment. In C. Kirchner and H. Kirchner, editors, *Proceedings of the 15th International Conference on Automated Deduction (CADE-15)*, volume 1421 of *Lecture Notes in Artificial Intelligence*, pages 191–204, Lindau, Germany, 1998. Springer-Verlag.
49. L. Fariñas del Cerro. Molog: A system that extends Prolog with modal logic. *New Generation Computing*, 4:35–50, 1986.
50. P. Devanbu, R. J. Brachman, P. G. Selfridge, and B. W. Ballard. LaSSIE: A knowledge-based software information system. *Communications of the ACM*, 34(5):34–49, 1991.
51. J. Dix. A classification-theory of semantics of normal logic programs: I. Strong properties. *Fundamenta Informaticae*, XXII(3):227–255, 1995.
52. J. Dix. A classification-theory of semantics of normal logic programs: II. Weak properties. *Fundamenta Informaticae*, XXII(3):257–288, 1995.
53. J. Dix. Semantics of logic programs: Their intuitions and formal properties. An overview. In A. Fuhrmann and H. Rott, editors, *Logic, Action and Information – Essays on Logic in Philosophy and Artificial Intelligence*, pages 241–327. DeGruyter, 1995.
54. J. Dix, U. Furbach, and A. Nerode, editors. *Logic Programming and Nonmonotonic Reasoning, Fourth International Conference*, volume 1265 of *Lecture Notes in Artificial Intelligence*. Springer–Verlag, 1997.
55. J. Dix, L. Pereira, and T. Przymusinski, editors. *Non-Monotonic Extensions of Logic Programming, Proceedings*, volume 927 of *Lecture Notes in Artificial Intelligence*. Springer–Verlag, 1995.

56. J. Dix, L. Pereira, and T. Przymusinski, editors. *Non-Monotonic Extensions of Logic Programming, Proceedings*, volume 1216 of *Lecture Notes in Artificial Intelligence*. Springer–Verlag, 1997.
57. F. M. Donini, B. Hollunder, M. Lenzerini, A. M. Spaccamela, D. Nardi, and W. Nutt. The complexity of existential quantification in concept languages. *Journal of Artificial Intelligence*, 53:309–327, 1992.
58. F. M. Donini, M. Lenzerini, D. Nardi, and W. Nutt. The complexity of concept languages. In J. Allen, R. Fikes, and E. Sandewall, editors, *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR-91)*, pages 151–162, Cambridge, Mass., 1991. Morgan Kaufmann, San Francisco.
59. F. M. Donini, M. Lenzerini, D. Nardi, and W. Nutt. Tractable concept languages. In J. Mylopoulos and R. Reiter, editors, *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, pages 458–463, Sydney, Australia, 1991. Morgan Kaufmann, San Francisco.
60. F. M. Donini, M. Lenzerini, D. Nardi, W. Nutt, and A. Schaerf. Adding epistemic operators to concept languages. In B. Nebel, Ch. Rich, and W. Swartout, editors, *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR-92)*, pages 342–356, Cambridge, Mass., 1992. Morgan Kaufmann, San Francisco.
61. F. M. Donini, M. Lenzerini, D. Nardi, and A. Schaerf. Deduction in concept languages: From subsumption to instance checking. *Journal of Logic and Computation*, 4(4):423–452, 1994.
62. F. M. Donini, M. Lenzerini, D. Nardi, and A. Schaerf. Reasoning in description logics. In G. Brewka, editor, *Principles of Knowledge Representation*, pages 191–236. CSLI Publications, Stanford (CA), USA, 1996.
63. W. A. Emerson. Temporal and modal logic. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science, Vol. B*. Elsevier, Amsterdam, The Netherlands, 1990.
64. M. Fattorosi-Barnaba and F. de Caro. Graded modalities I. *Studia Logica*, 44:197–221, 1985.
65. C. Fernmüller, A. Leitsch, T. Tammet, and N. Zamov. *Resolution Methods for the Decision Problem*, volume 679 of *Lecture Notes in Artificial Intelligence*. Springer–Verlag, 1993.
66. K. Fine. In so many possible worlds. *Notre Dame Journal of Formal Logics*, 13:516–520, 1972.
67. M. J. Fischer and R. E. Ladner. Propositional modal logic of programs. In *Ninth Annual ACM Symposium on Theory of Computing*, pages 286–294, New York, N.Y., 1977. ACM.
68. M. Fitting. Basic modal logic. In D. M. Gabbay, C. J. Hogger, and J. A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 1*. Oxford University Press, Oxford, UK, 1993.
69. D. M. Gabbay. Theoretical foundations for non-monotonic reasoning in expert systems. In K. R. Apt, editor, *Proceedings of the NATO Advance Study Institute on Logics and Models of Concurrent Systems*, La Colle-sur-Loup, France, 1985. Springer–Verlag.
70. D. M. Gabbay. Modal and temporal logic programming. In A. Galton, editor, *Temporal Logics and Their Applications*. Academic Press, London, UK, 1987.
71. D. M. Gabbay, C. J. Hogger, and J. A. Robinson, editors. *Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 3: Nonmonotonic Reasoning and Uncertain Reasoning*. Oxford University Press, Oxford, UK, 1994.

72. G. Gargov and V. Goranko. Modal logic with names. *J. Philosophical Logic*, 22:607–636, 1993.
73. M. Gelfond. Logic programming and reasoning with incomplete information. *Annals of Mathematics and Artificial Intelligence*, 12(1–2):89–116, 1994.
74. M. Gelfond and V. Lifschitz. The stable semantics for logic programs. In R. Kowalski and K. Bowen, editors, *Proceedings of the 5th International Symposium on Logic Programming*, pages 1070–1080, Cambridge, MA., 1988. MIT Press.
75. M. Gelfond and V. Lifschitz. Logic programs with classical negation. In D. Warren and P. Szeredi, editors, *Proceedings of the 7th International Conference on Logic Programming*, pages 579–597, Cambridge, MA, 1990. MIT Press.
76. M. Gelfond and V. Lifschitz. Classical negation in logic programs and disjunctive databases. *New Generation Computing*, 9:365–385, 1991.
77. G. De Giacomo and M. Lenzerini. Boosting the correspondence between description logics and propositional dynamic logics. In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94)*, pages 205–212. AAAI-Press/The MIT-Press, 1994.
78. G. De Giacomo and M. Lenzerini. Concept languages with number restrictions and fixpoints, and its relationship with mu-calculus. In A. G. Cohn, editor, *Proceedings of the Eleventh European Conference on Artificial Intelligence (ECAI-94)*, pages 411–415. John Wiley and Sons, 1994.
79. G. De Giacomo and M. Lenzerini. TBox and ABox reasoning in expressive description logics. In L. C. Aiello, J. Doyle, and S. Shapiro, editors, *Proceedings of the Fifth International Conference on the Principles of Knowledge Representation and Reasoning (KR-96)*, pages 316–327, Cambridge, Mass., 1996. Morgan Kaufmann, San Francisco.
80. M. L. Ginsberg, editor. *Readings in Nonmonotonic Reasoning*. Morgan Kaufmann, San Francisco, 1987.
81. L. Giordano and A. Martelli. Structuring logic programs: A modal approach. *Journal of Logic Programming*, 21(2):59–94, 1994.
82. E. Giunchiglia, F. Giunchiglia, R. Sebastiani, and A. Tacchell. More evaluation of decision procedures for modal logics. In A. Cohn, L. Schubert, and S. Shapiro, editors, *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR-98)*, pages 626–635, Trento, Italy, 1998. Morgan Kaufmann, San Francisco.
83. F. Giunchiglia and R. Sebastiani. A SAT-based decision procedure for $\mathcal{ALC}$. In L. C. Aiello, J. Doyle, and S. Shapiro, editors, *Proceedings of the Fifth International Conference on the Principles of Knowledge Representation and Reasoning (KR-96)*, pages 304–314, Cambridge, Mass., 1996. Morgan Kaufmann, San Francisco.
84. E. Grädel. On the restraining power of guards. *J. Symbolic Logic*, 1998. To appear.
85. E. Grädel, P. G. Kolaitis, and M. Y. Vardi. On the decision problem for two-variable first-order logic. *The Bulletin of Symbolic Logic*, 3(1):53–69, 1997.
86. E. Grädel, M. Otto, and E. Rosen. Two-variable logic with counting is decidable. In G. Winskel, editor, *Proceedings of the Twelfth Annual IEEE Symposium on Logic in Computer Science (LICS-97)*, pages 306–317, Warsaw, Poland, 1997. IEEE Computer Society Press.
87. V. Haarslev, C. Lutz, and R. Möller. Foundations of spatioterminological reasoning with description logics. In A. Cohn, L. Schubert, and S. Shapiro, editors, *Proceedings of the Sixth International Conference on Principles of Knowledge*

Representation and Reasoning (KR-98), pages 112–123, Trento, Italy, 1998. Morgan Kaufmann, San Francisco.
88. J. Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logic of knowledge and belief. *Journal of Artificial Intelligence*, 54:319–379, 1992.
89. P. Hanschke. Specifying role interaction in concept languages. In B. Nebel, Ch. Rich, and W. Swartout, editors, *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR-92)*, pages 318–329, Cambridge, Mass., 1992. Morgan Kaufmann, San Francisco.
90. P. Hanschke. *A Declarative Integration of Terminological, Constraint-based, Data-driven, and Goal-directed Reasoning*, volume 122 of *DISKI*. Infix, Sankt Augustin, Germany, 1996.
91. P. J. Hayes. In defence of logic. In *Proceedings of the 5th International Joint Conference on Artificial Intelligence, IJCAI'77*, pages 559–565, Cambridge, Mass., 1977.
92. P. J. Hayes. The logic of frames. In D. Mentzing, editor, *Frame Conceptions and Text Understanding*. Walter de Gruyter and Co., Berlin, Germany, 1979.
93. B. Hollunder. Hybrid inferences in KL-ONE-based knowledge representation systems. In *14th German Workshop on Artificial Intelligence*, volume 251 of *Informatik-Fachberichte*, pages 38–47, Ebingerfeld, Germany, 1990. Springer–Verlag.
94. B. Hollunder and F. Baader. Qualifying number restrictions in concept languages. In J. Allen, R. Fikes, and E. Sandewall, editors, *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, pages 335–346, Cambridge, Mass., 1991. Morgan Kaufmann, San Francisco.
95. B. Hollunder, W. Nutt, and M. Schmidt-Schauß. Subsumption algorithms for concept description languages. In L. C. Aiello, editor, *Proceedings of the 9th European Conference on Artificial Intelligence (ECAI-90)*, pages 348–353, Stockholm, Sweden, 1990. Pitman Publishing, London.
96. T. Hoppe, C. Kindermann, J. Quantz, A. Schmiedel, and M. Fischer. BACK V5: tutorial and manual. KIT Report 100, Technical University of Berlin, Berlin, Germany, 1993.
97. I. Horrocks. *Optimizing Tableaux Decision Procedures for Description Logics*. Ph.D. thesis, University of Manchester, Manchester, UK, 1997.
98. I. Horrocks. Using an expressive Description Logic: FaCT or fiction? In A. Cohn, L. Schubert, and S. Shapiro, editors, *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR-98)*, pages 636–647, Trento, Italy, 1998. Morgan Kaufmann, San Francisco.
99. I. Horrocks and G. Gough. Description logics with transitive roles. In *Proceedings of the 1997 International Workshop on Description Logics (DL-97)*, pages 25–28, Gif sur Yvette, France, 1997. LRI, Université Paris-Sud.
100. I. Horrocks and P. F. Patel-Schneider. Optimising propositional modal satisfiability for description logic subsumption. In J. Calmet and J. Plaza, editors, *Proceedings of the International Conference on Artificial Intelligence and Symbolic Computation (AISC-98)*, volume 1476 of *Lecture Notes in Artificial Intelligence*, pages 234–246. Springer–Verlag, 1998.
101. I. Horrocks and U. Sattler. A description logic with transitive and converse roles and role hierarchies. In *Proceedings of the 1998 International Workshop on Description Logics (DL-98)*, pages 72–81, Povo - Trento, Italy, 1998. ITC-irst Report 9805-03.
102. G. E. Hughes and M. J. Cresswell. *A Companion to Modal Logic*. Methuen & Co., London, 1984.

103. U. Hustadt and R. A. Schmidt. On evaluating decision procedures for modal logics. In M. Pollack, editor, *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI'97)*, volume 1, pages 202–207. Morgan Kaufmann, 1997.
104. U. Hustadt and R. A. Schmidt. Simplification and backjumping in modal tableau. In H. Swart, editor, *Proceedings of TABLEAUX'98*, volume 1397 of *Lecture Notes in Computer Science*, pages 187–201. Springer–Verlag, 1998.
105. T. Imielinski. Results on translating defaults to circumscription. *Journal of Artificial Intelligence*, 32(1):131–146, 1987.
106. A. Kobsa. First experiences with the SB-ONE knowledge representation workbench in natural language applications. *SIGART Bulletin*, 2(3):70–76, 1991.
107. J. Koehler. An application of terminological logics to case-based reasoning. In J. Doyle, E. Sandewall, and P. Torasso, editors, *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning (KR-94)*, pages 351–362, Bonn, Germany, 1994. Morgan Kaufmann, San Francisco.
108. K. Konolige. On the relation between default and autoepistemic logic. *Journal of Artificial Intelligence*, 35(3):343–382, 1988.
109. K. Konolige. On the relation between circumscription and autoepistemic logic. In N. S. Sridharan, editor, *Proceedings of the 11th International Joint Conference on Artificial Intelligence, IJCAI-89*, pages 1213–1218, Detroit, MI, 1989. Morgan Kaufmann, San Francisco.
110. R. A. Kowalski. Logic for knowledge representation. In M. Joseph and R. Shyamasundar, editors, *Proceedings of the Fourth Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 1–12, 1984.
111. S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Journal of Artificial Intelligence*, 44(1–2):167–207, 1990.
112. R. Küsters. Characterizing the semantics of terminological cycles in $\mathcal{ALN}$ using finite automata. In A. Cohn, L. Schubert, and S. Shapiro, editors, *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR-98)*, pages 499–510, Trento, Italy, 1998. Morgan Kaufmann, San Francisco.
113. A. Laux. Beliefs in multi-agent worlds: A terminological logics approach. In A. G. Cohn, editor, *Proceedings of the Eleventh European Conference on Artificial Intelligence (ECAI-94)*. John Wiley and Sons, 1994.
114. D. Lehmann. What a conditional knowledge base entails. In R. J. Brachman, editor, *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning*, pages 212–222, Toronto, Ont., 1989.
115. H. J. Levesque and R. J. Brachman. Expressiveness and tractability in knowledge representation and reasoning. *Computational Intelligence*, 3:78–93, 1987.
116. A. Levy and M.-C. Rousset. Verification of knowledge bases based on containment checking. *Journal of Artificial Intelligence*, 101:227–256, 1998.
117. A. Levy and M.C. Rousset. CARIN: A representation language combining Horn rules and description logics. In W. Wahlster, editor, *Proceedings of the Twelfth European Conference on Artificial Intelligence (ECAI-96)*, pages 323–327. John Wiley & Sons Ltd, 1996.
118. V. Lifschitz. Between circumscription and autoepistemic logic. In R. J. Brachman, editor, *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning*, pages 235–244, Toronto, Ont., 1989.

119. J. Lobo, J. Minker, and A. Rajasekar. *Foundations of Disjunctive Logic Programming*. MIT-Press, 1992.
120. W. Lukaszewicz. *Non-Monotonic Reasoning*. Ellis Horwood Series in Artificial Intelligence, 1991.
121. R. MacGregor. Inside the LOOM classifier. *SIGART Bulletin*, 2(3):88–92, 1991.
122. R. MacGregor and R. Bates. The LOOM knowledge representation language. Technical Report ISI/RS-87-188, University of Southern California, 1987.
123. D. Makinson. General theory of cumulative inference. In M. Reinfrank, J. de Kleer, M. L. Ginsberg, and E. Sandewall, editors, *Proceedings of the 2nd International Workshop on Non-monotonic Reasoning*, volume 346 of *Lecture Notes in Artificial Intelligence*, pages 1–18, Grassau, Germany, 1989. Springer-Verlag.
124. D. Makinson. General patterns in nonmonotonic reasoning. In D. M. Gabbay, C. J. Hogger, and J. A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 3: Nonmonotonic and Uncertain Reasoning*, pages 35–110. Oxford University Press, Oxford, UK, 1994.
125. V. Marek, A. Nerode, and M. Truszczynski, editors. *Logic Programming and Nonmonotonic Reasoning, Third International Conference*, volume 928 of *Lecture Notes in Artificial Intelligence*. Springer–Verlag, 1995.
126. V. M. Marek and M. Truszczyński. *Nonmonotonic Logics; Context-Dependent Reasoning*. Springer–Verlag, Berlin-Heidelberg-New York, 1993.
127. E. Mays, C. Apté, J. Griesmer, and J. Kastner. Experience with K-Rep: An object-centered knowledge representation language. In *Proceedings of IEEE CAIA-88*, pages 62–67, 1988.
128. E. Mays, R. Dionne, and R. Weida. K-Rep system overview. *SIGART Bulletin*, 2(3):93–97, 1991.
129. J. McCarthy. Circumscription – A form of non-monotonic reasoning. *Journal of Artificial Intelligence*, 13(1–2):27–39, 1980.
130. J. McCarthy and P. J. Hayes. Some philosophical problems from the standpoint of AI. *Machine Intelligence*, 4:463–502, 1969. Reprinted in *Readings in Knowledge Representation*, R. J. Brachman and H. J. Levesque (editors), Morgan Kaufman, 1985.
131. D. McDermott. Nonmonotonic logic II: Nonmonotonic modal theories. *Journal of the ACM*, 29(1):33–57, 1982.
132. D. McDermott and J. Doyle. Non-monotonic logic I. *Journal of Artificial Intelligence*, 13(1–2):41–72, 1980.
133. D. L. McGuinness, L. Alperin Resnick, and C. Isbell. Description Logic in practice: A CLASSIC application. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence, IJCAI'95*, pages 2045–2046, Montréal, Canada, 1995. Morgan Kaufmann, San Francisco. Video Presentation.
134. J. Minker. An overview of nonmonotonic reasoning and logic programming. *Journal of Logic Programming*, 17:95–126, 1993.
135. M. Minsky. A framework for representing knowledge. In P. Winston, editor, *The Psychology of Computer Vision*. McGraw-Hill, New York, 1975.
136. R. C. Moore. Semantical considerations on nonmonotonic logic. *Journal of Artificial Intelligence*, 25(1):75–94, 1985.
137. M. Mortimer. On languages with two variables. *Zeitschr. f. math. Logik und Grundlagen d. Math.*, 21:135–140, 1975.
138. B. C. Moszkowski. *Executing Temporal Logic Programs*. Cambridge University Press, Cambridge, UK, 1986.
139. B. Nebel. Computational complexity of terminological reasoning in BACK. *Journal of Artificial Intelligence*, 34(3):371–383, 1988.

140. B. Nebel. Terminological cycles: Semantics and computational properties. In J. Sowa, editor, *Principles of Semantic Networks*, pages 331–362. Morgan Kaufmann, San Francisco, 1989.
141. B. Nebel. *Reasoning and Revision in Hybrid Representation Systems*, volume 422 of *Lecture Notes in Computer Science*. Springer–Verlag, 1990.
142. B. Nebel. Terminological reasoning is inherently intractable. *Journal of Artificial Intelligence*, 43(2):235–249, 1990.
143. B. Nebel and K. von Luck. Hybrid reasoning in BACK. In Z. W. Ras and L. Saitta, editors, *Methodologies for Intelligent Systems*, volume 3, pages 260–269. North-Holland, 1988.
144. L. Pacholski, W. Szwast, and L. Tendera. Complexity of two-variable logic with counting. In G. Winskel, editor, *Proceedings of the Twelfth Annual IEEE Symposium on Logic in Computer Science (LICS-97)*, pages 318–327, Warsaw, Poland, 1997. IEEE Computer Society Press.
145. R. Parikh. Propositional dynamic logics of programs: A survey. In E. Engeler, editor, *Proceedings of the Workshop on Logic of Programs*, volume 125 of *LNCS*, pages 102–144, Zürich, Switzerland, 1979. Springer–Verlag.
146. P. F. Patel-Schneider. A four-valued semantics for terminological logics. *Journal of Artificial Intelligence*, 38(3):319–351, 1989.
147. P. F. Patel-Schneider. Undecidability of subsumption in NIKL. *Journal of Artificial Intelligence*, 39(2):263–272, 1989.
148. P. F. Patel-Schneider. DLP system description. In *Proceedings of the 1998 International Workshop on Description Logic (DL-98)*, pages 87–89, Trento, Italy, 1998. ITC-irst Report 9805-03.
149. C. Peltason. The BACK system – an overview. *SIGART Bulletin*, 2(3):114–119, 1991.
150. T. Przymusinski. Non-monotonic formalisms and logic programming. In *Proceedings of the 6th International Conference on Logic Programming*, pages 655–674, Cambridge, MA, 1989. MIT Press.
151. T. Przymusinski. Stable semantics for disjunctive programs. *New Generation Computing*, 9:401–424, 1991.
152. J. Quantz, G. Dunker, and V. Royer. Flexible inference strategies for DL systems. In *Proceedings of the 1994 International Workshop on Description Logic (DL-94)*, pages 27–30, Bonn, Germany, 1994. DFKI Document D-94-10, DFKI, Saarbrücken (Germany).
153. J. Quantz and V. Royer. A preference semantics for defaults in terminological logics. In B. Nebel, Ch. Rich, and W. Swartout, editors, *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR-92)*, pages 294–305, Cambridge, Mass., 1992. Morgan Kaufmann, San Francisco.
154. J. Quantz and B. Schmitz. Knowledge-based disambiguation for machine translation. *Minds and Machines*, 4:39–57, 1994.
155. M. Quillian. Semantic memory. In M. Minsky, editor, *Semantic Information Processing*, pages 216–270. MIT Press, Cambridge, Mass., 1968.
156. R. Reiter. On closed world data bases. In *Logic and Data Bases*, pages 55–76. Plenum Press, New York, 1978.
157. Raymond Reiter. A logic for default reasoning. *Journal of Artificial Intelligence*, 13:81–132, 1980.
158. N. Rychtyckyj. DLMS: An evaluation of KL-ONE in the automobile industry. In L. C. Aiello, J. Doyle, and S. Shapiro, editors, *Proceedings of the Fifth International Conference on Principles of Knowledge Representation and Reasoning*

(KR-96), pages 588–596, Cambridge, Mass., 1996. Morgan Kaufmann, San Francisco.
159. U. Sattler. A concept language extended with different kinds of transitive roles. In G. Görz and S. Hölldobler, editors, *20. Deutsche Jahrestagung für Künstliche Intelligenz, KI'97*, volume 1137 of *Lecture Notes in Artificial Intelligence*. Springer–Verlag, 1996.
160. A. Schaerf. On the complexity of the instance checking problem in concept languages with existential quantification. *Journal of Intelligent Information Systems*, 2:265–278, 1993.
161. K. Schild. A correspondence theory for terminological logics: Preliminary report. In J. Mylopoulos and R. Reiter, editors, *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, pages 466–471, Sydney, Australia, 1991. Morgan Kaufmann, San Francisco.
162. K. Schild. Combining terminological logics with tense logic. In M. Filgueiras and L. Damas, editors, *Progress in Artificial Intelligence – 6th Portuguese Conference on Artificial Intelligence, EPIA'93*, Lecture Notes in Artificial Intelligence, pages 105–120, Porto, Portugal, 1993. Springer–Verlag.
163. K. Schild. Terminological cycles and the propositional μ-calculus. In J. Doyle, E. Sandewall, and P. Torasso, editors, *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning (KR'94)*, pages 509–520, Bonn, Germany, 1994. Morgan Kaufmann, San Francisco.
164. R. A. Schmidt. Resolution is a decision procedure for many propositional modal logics. In M. Kracht, M. de Rijke, H. Wansing, and M. Zakharyaschev, editors, *Advances in Modal Logic*, volume 87 of *CSLI Lecture Notes*, pages 189–208. CSLI Publications, Stanford, Cal., 1998.
165. M. Schmidt-Schauß. Subsumption in KL-ONE is undecidable. In R. J. Brachman, editor, *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning*, pages 421–431, Toronto, Ont., 1989. Morgan Kaufmann, San Francisco.
166. M. Schmidt-Schauß and G. Smolka. Attributive concept descriptions with complements. *Journal of Artificial Intelligence*, 47:1–26, 1991.
167. L. K. Schubert, R. G. Goebel, and N. J. Cercone. The structure and organization of a semantic net for comprehension and inference. In N. V. Findler, editor, *Associative Networks: Representation and Use of Knowledge by Computers*, pages 121–175. Academic Press, 1979.
168. J. C. Shepherdson. Negation as failure: A comparison of Clark's completed database and Reiter's closed world assumption. *J. Logic Programming*, 1(15):51–79, 1984.
169. Y. Shoham. A semantical approach to nonmonotonic logics. In *IEEE Symposium on Logic in Computer Science*, pages 275–279, Ithaca, N.Y., 1987.
170. R. M. Smullyan. *First-Order Logic*, volume 43 of *Ergebnisse der Mathematik und ihrer Grenzgebiete*. Springer–Verlag, Berlin, Heidelberg, New York, 1968.
171. C. Stirling. Modal and temporal logic. In S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum, editors, *Handbook of Logic in Computer Science, Vol. 2*. Oxford University Press, Oxford, UK, 1993.
172. W. van der Hoek and M. de Rijke. Counting objects. *J. Logic and Computation*, 5(3):325–345, 1995.
173. A. van Gelder, K. A. Ross, and J. S. Schlipf. The well-founded semantics for general logic programs. *Journal of the ACM*, 38:620–650, 1991.
174. M. Y. Vardi and P. Wolper. Automata-theoretic techniques for modal logic of programs. *Journal of Computer and System Sciences*, 32:183–221, 1986.

175. F. Wolter and M. Zakharyaschev. Satisfiability problems in description logics with modal operators. In A. Cohn, L. Schubert, and S. Shapiro, editors, *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR-98)*, pages 512–523, Trento, Italy, 1998. Morgan Kaufmann, San Francisco.
176. W. A. Woods. What's in a link: foundations for semantic networks. In D. G. Bobrow and A. M. Collins, editors, *Representation and Understanding: Studies in Cognitive Science*, pages 35–82. Academic Press, London, 1975.
177. W. A. Woods and J. G. Schmolze. The KL-ONE family. *Computers and Mathematics with Applications*, 23(2–5):133–177, 1992.
178. J. R. Wright, E. S. Weixelbaum, K. Brown, G. T. Vesonder, S. R. Palmer, J. I. Berman, and H. H. Moore. A knowledge-based configurator that supports sales, engineering, and manufacturing at AT&T network systems. *AI Magazine*, 14(3):69–80, 1993.