Representation and Reasoning, volume 810 of *Lecture Notes in Artificial Intelligence*,
p. 1–12. Springer-Verlag, Berlin, Heidelberg, New York 1994

Foundations of Knowledge Representation and Reasoning

A Guide to this Volume

Gerhard Lakemeyer[1] and Bernhard Nebel[2]

[1] Institute of Computer Science
University of Bonn
Römerstr. 164
D-53117 Bonn, Germany
gerhard@cs.uni-bonn.de

[2] Department of Computer Science
University of Ulm
D-89069 Ulm, Germany
nebel@informatik.uni-ulm.de

1 Introduction

Knowledge representation (KR) is the area of Artificial Intelligence that deals with the problem of representing, maintaining, and manipulating knowledge about an application domain. Since virtually all Artificial Intelligence systems have to address this problem, KR is one of the central subfields of Artificial Intelligence.

Main research endeavors in KR are

- representing knowledge about application areas (e.g., medical knowledge, knowledge about time, knowledge about physical systems),
- developing appropriate representation languages,
- specifying and analyzing reasoning over represented knowledge, and
- implementing systems that support the representation of knowledge and reasoning over the represented knowledge.

While knowledge about an application domain may be represented in a variety of forms, e.g., procedurally in form of program code or implicitly as patterns of activation in a neural network, research in the area of knowledge representation assumes an *explicit* and *declarative* representation, an assumption that distinguishes KR from research in, e.g., programming languages and neural networks. *Explicitness* of representation means that the represented knowledge is stored in a *knowledge base* consisting of a set of formal entities that describe the knowledge in a direct and unambiguous way. *Declarativeness* means that the (formal) meaning of the representation can be specified without reference to how the knowledge is applied procedurally, implying some sort of logical methodology behind it.

Although the above two points seem to be almost universally accepted by researchers working in KR, this consensus has been achieved only recently. Brachman and Levesque mentioned in the Introduction to a collection of papers in 1985 that the "research area of Knowledge Representation has a long, complex, and as yet non-convergent history," [Brachman and Levesque, 1985, p. xiii] an impression that is indeed confirmed by the papers in this collection. A large portion of the papers contain meta-level discussions arguing about the right methods for representing knowledge or they present approaches that are completely incompatible with a logical, declarative point of view.

Nowadays, the picture has completely changed, however. Logical methods predominate and methodological problems are hardly discussed any longer [Brachman *et al.*, 1989; Allen *et al.*, 1991; Nebel *et al.*, 1992; Brachman, 1990]. Instead, research papers focus on particular technical representation and reasoning problems and address these problems using methods from logic and computer science.

While this development indicates that KR has become a mature scientific discipline, it also leads to the situation that research results in KR appear to be less accessible to the rest of the Artificial Intelligence community. As a matter of fact, it is often argued that the foundational results that are achieved in the KR field are not relevant to Artificial Intelligence at all.

We concede that a large amount of KR research probably does not have any immediate impact on building Artificial Intelligence systems. However, this is probably asking for too much. Foundational KR research aims at providing the *theoretical foundations* on which we can build systems that are useful, comprehensible, and reliable, i.e., it aims at providing the *logical* and *computational* foundations of knowledge representation formalisms and reasoning processes. Results in foundational KR often "only" provide explanations why a particular approach works or how an approach can be interpreted logically. Additionally, the borderlines of what can be represented are explored and it is analyzed how efficiently a reasoning process can be. While this may not be of central concern when building Artificial Intelligence systems, such results are nevertheless important when we want to understand such systems, and when we want to guarantee their reliability.

Perhaps the main motivation and driving force behind most research in KR has been the desire to equip artifacts with *commonsense*. This is literally true of a paper by John McCarthy, first published in 1958 and republished as [McCarthy, 1968], which started the whole KR enterprise, and it is still true, if only implicitly, of the papers in this book. In fact, work on the foundations of KR can largely be indentified with work on the foundations of commonsense reasoning, a point of view which we will follow throughout this brief survey.

In the following sections, we touch on some basic logical and computational aspects of commonsense reasoning. The reader is warned that this is not a comprehensive overview of the field, which would be far outside the scope of this book. Instead we confine ourselves mainly to those areas that are actually covered by papers in this book.

2 Logical Foundations of Commonsense Reasoning

As mentioned already in the beginning, the main assumption that distinguishes knowledge-based systems from other approaches is that knowledge is represented declaratively in some logic-like language. This is one part of what Brian Smith has called *the knowledge representation hypothesis* [Smith, 1982]. The other part postulates that these representations play a causal role in engendering the behavior of the system. While this causal connection is present in one form or another in every knowledge-based system, it is fair to say that so far there are very few, if any, theoretical results that explain this connection.

Hence most foundational research in KR, including the work reported in this book, deals with problems that arise from the first part of the KR hypothesis and which can be dealt with independently from the second part. In this context, one can identify three fundamental questions:

1. What is the right representation language?
2. What inferences should be drawn from a knowledge base?
3. How do we incorporate new knowledge?

In the rest of this section, we will address each question in turn with an emphasis on the relevant papers in this book.

2.1 The Right Representation Language

While there is little disagreement any more about the assumption that a representation language is one of logic, where the sentences can be interpreted as propositions about the world,[3] designing an adequate language is not an easy task, since the various desirable features are often incompatible. In particular, very expressive languages usually have poor computational properties, an issue that has drawn considerable interest since a seminal paper by Brachman and Levesque [1984] and which is discussed in more detail in the next section. At this point we only mention that computational considerations have led to the development of languages that are far less expressive than full first-order logic, most notably the so-called *concept languages* or *terminological logics*. Four of the papers in this collection are devoted to this topic [Baader and Hollunder, 1994; Bettini, 1994; Allgayer and Franconi, 1994; Donini *et al.*, 1994]. From the point of view of expressiveness, it often seems useful to have special epistemological or ontological primitives built into the language. Shoham and Cousins [1994] survey work in AI on a whole range of mental attitudes like beliefs, desires, goals, or intentions. The need for making such notions explicit is probably most convincing in multi-agent settings, where agents need to reason about each other's mental attitudes in order to communicate and cooperate successfully. [Gottlob, 1994; Kalinski, 1994; Niemelä and Rintanen, 1994] consider the specific case of *belief*,

[3] Until the late seventies, many so-called representation languages actually violated this fundamental assumption and led to vivid discussions such as [Hayes, 1977; McDermott, 1978].

which, together with *knowledge*, is probably the best understood among the attitudes. In these papers a very specific aspect of belief is considered, namely the ability to model certain forms of defeasible reasoning by referring explicitly to the system's own epistemic state (see Section 2.2 below). Bettini and Lin [Bettini, 1994; Lin, 1994], on the other hand, are concerned with adding explicit notions of time to the language. While Bettini considers incorporating an existing interval-based concept of time to a temporal logic, Lin proposes a new axiomatization of time, where time instances are defined on the basis of events.

2.2 The Right Inferences

Having explicit representations of knowledge alone is not very useful in general. Instead one wants to reason about these representations to uncover what is *implied* by them. After all, we use the term commonsense reasoning and not commonsense representation. Until the early seventies, *deduction* was the main focus of attention as far as inference mechanisms are concerned. It became clear, however, that a lot of commonsense reasoning is not deductive in nature.

In particular, many inferences humans draw all the time are uncertain in some sense and may therefore be defeasible if new information becomes available. The prototypical example is the assumption that birds normally fly and if someone tells me about a bird called Tweety, then, knowing nothing else, I conclude that Tweety flies. Later on, if I find out that Tweety is indeed a penguin, I withdraw my earlier conclusion without hesitation. There are essentially two main research fields that try to formalize such reasoning, one which is based on probability theory (see, for example, [Pearl, 1988]) and another which directly models *nonmonotonic reasoning* by modifying classical logic in one way or another (see, for example, [Brewka, 1991]). While probabilistic methods are not dealt with at all in this volume, nonmonotonic reasoning receives a fairly broad coverage [Baader and Hollunder, 1994; Gottlob, 1994; Kakas, 1994; Kalinski, 1994; Niemelä and Rintanen, 1994; Weydert, 1994]. Except for McCarthy's [1980] Circumscription, the main formalisms on nonmonotonic reasoning are represented in this volume. Baader and Hollunder [1994] discuss extending terminological logics using Reiter's [1980] Default Logic (DL). Kakas extends DL by applying ideas from abductive logic programming to it. Gottlob [1994] relates DL and Moore's [1985] Autoepistemic Logic (AEL) by showing how to faithfully translate DL theories into AEL theories. Both Kalinski [1994] and Niemelä and Rintanen [1994] are concerned with complexity issues, the former by considering a weaker form of AEL and the latter by considering only AEL theories of a special form (with applications to other nonmonotonic formalisms as well). Finally, Weydert [1994] presents results on nested conditionals. This work is in the tradition of modeling nonmonotonic inferences on the basis of conditional logics such as [Lewis, 1973; Adams, 1975].

Apart from probabilistic and nonmonotonic reasoning, there are many other forms such as fuzzy, inductive, abductive or analogical reasoning. Of those the latter two are represented here with one paper each. Console and Dupre [1994] address abduction, which is concerned with finding plausible explanations for a given observation. In particular, they address the problem of finding explanations at different levels of abstraction. Myers and Konolige [1994] discuss reasoning with analogical representations such as maps. They are particularly concerned with integrating both analogical and symbolic (sentential) representations.

2.3 Evolving Knowledge

Since knowledge bases are hardly ever static, devising methods for incorporating new information into a knowledge base is of great importance in KR research. This problem, often referred to as *belief revision*, is particularly challenging if the new information conflicts with the contents of the old knowledge base. Over the past decade, substantial progress has been made on the topic of belief revision, particularly since the ground-breaking work by Alchourrón, Gärdenfors, and Makinson [1985], who propose postulates which any rational revision operator should obey (now referred to as AGM-postulates). Later, Katsuno and Mendelzon [1991] introduce an important distinction between revising a knowledge base, which refers to incorporating new information about a static world, and updating it, where the new information reflects changes in the world. They also propose a set of rationality postulates for update operators. In this volume, Boutilier [1994] and Nejdl and Banagl [1994] present new results following this line of research. Nejdl and Banagl define subjunctive queries for knowledge bases in the case of both update and revision. In particular, they show that their query semantics for revision and update satisfies precisely the AGM-postulates and the Katsuno-Mendelzon-postulates, respectively. Boutilier shows that, in the context of conditional logic, belief revision and nonmonotonic reasoning have precisely the same properties, further substantiating the claim that the two areas are closely related.

Witteveen and Jonker[1994] address revision from a somewhat different angle. Here the emphasis is on finding plausible expansions of logic programs, which are incoherent under the well-founded semantics, such that the revised programs are no longer incoherent.

3 Commonsense Reasoning as Computation

Once a knowledge representation scheme together with its associated commonsense reasoning task has been formalized logically, we can immediately make use of the computational machinery associated with logic. For instance, once we have identified that a particular representation formalism is "simply" a subset of standard first-order logic, we know that resolution (or any other complete proof method) is a method to compute all the valid consequences of a knowledge base. In other words, in such a case, commonsense reasoning could be reduced to a well-known computation technique.

However, this point of view is over-simplifying. First of all, often one deals with *non-standard* logics, e.g., non-monotonic or modal logics, for which standard techniques do not work. Secondly, even in the case that one only has a subset of standard first-order logic, it does not make sense to use general proof methods if specialized reasoning techniques, tailored to the restricted language, turn out to be much more *efficient*. In particular, one might be able to specify methods that always terminate, i.e., *inference algorithms*.

Efficiency is indeed one of the major problems when we turn logical formalization into computation. As is well-known, even propositional logic requires already significant computational resources – reasoning in propositional logic is NP-hard.[4] On the other hand, commonsense reasoning appears to be quite fast when humans perform it, and, moreover, should work reasonably fast on computers if the system is required to be of any use [Levesque, 1988]. In particular, if it is required that the reasoning process is *computationally tractable*, we are often forced to restrict the expressiveness of the representation language or to give up on the accuracy of the answer [Levesque and Brachman, 1987].

Research questions coming up in this context are:

1. Can we specify an inference algorithms for the reasoning task?
2. What is the computational complexity of the reasoning task?
3. How can we achieve tractability?

3.1 Inference Algorithms

As is evident from most papers, the formalization of a commonsense reasoning task as a form of logical inference is usually not overwhelmingly difficult, provided appropriate formal techniques and tools are employed. For instance, the semantics of a terminological logic extended by operators to express collective entities and relations [Allgayer and Franconi, 1994] can be specified on less than half a page. What appears to be much more involved is the specification of an appropriate reasoning technique.

As pointed out above, one could employ standard proof techniques if the formalism under consideration is (a notational variant of) a subset of standard first-order logic. However, usually we do not want an arbitrary method, but an *algorithm* that is as efficient as possible – a problem that is addressed by most of the papers in this volume.

Allgayer and Franconi [1994], for instance, showed in their paper that it is possible to extend the tableau-based technique introduced by Schmidt-Schauß and Smolka [Schmidt-Schauß and Smolka, 1991] to terminological logics containing operators for collective entities, providing us with a sound, complete, and terminating method for reasoning in this language.

Baader and Hollunder [1994] also start with terminological logics, but extend these by incorporating default logic [Reiter, 1980], i.e., in this case it is not possible to use standard first-order logic methods. However, as they are able to show, it is possible to combine the tableau-based reasoning techniques for terminological logics with reasoning techniques developed for default logics [Junker

[4] Consult, e.g., [Garey and Johnson, 1979] for an introduction to computational complexity theory.

and Konolige, 1990; Schwind and Risch, 1991] in an almost straightforward way, leading to an inference algorithm for the combined formalism. It should be noted that in order to guarantee decidability, it is necessary to use a somewhat non-standard interpretation of open defaults, though. Since, as shown by Baader and Hollunder, the standard interpretation of open defaults not only leads to undecidability but also to counter-intuitive results, giving up this interpretation does not seem to be much of a sacrifice.

3.2 Computational Complexity of Reasoning

An inference algorithm for a particular commonsense reasoning task demonstrates that that there is one way to turn this task into computation. However, it does not answer the question whether this is the most efficient way. In order to answer this question, computational complexity theory can be used for analyzing the inherent difficulty of the problem. Such an analysis can guide the search for more efficient algorithms or for a reformulation of the reasoning problem in a way that renders reasoning more efficient. Finally, a computational complexity analysis can be used to compare and contrast different reasoning problems.

For instance, Donini et al [1994] study the extension of terminological logics by an epistemic operator and show that this operator does not increase the computational complexity of reasoning in one of the standard terminological logics (the so-called $\mathcal{ALC}$ language [Schmidt-Schauß and Smolka, 1991]). Furthermore, Donini et al [1994] are able to show that in some relevant special cases the complexity goes even down from co-NP-hardness to polynomial time.

Kautz and Selman [1994] analyze the computational problems arising when approximating arbitrary propositional theories by Horn theories. They show that such a Horn theory may sometimes be of exponential size and that it is unlikely that a dense representation can be found in all cases.

A final example for the use of computational complexity theory is the paper by Gottlob [1994]. Although this paper is not by itself a paper on computational complexity analysis of commonsense reasoning, it makes use of computational complexity results [Gottlob, 1992] that show that the three main forms of nonmonotonic reasoning all have the same complexity, which implies that there must exist (polynomial) translations between these formalisms. Based on this observation, Gottlob develops a translation from default logic to autoepistemic logic that is quite interesting.

3.3 The Expressiveness vs. Efficiency Tradeoff

If a reasoning problem can be shown to require time that is not polynomial in the size of the problem description (under the assumption that NP$\neq$P), this implies that in the *worst case* we will not get an answer in tolerable time when the problem description grows beyond a certain (usually moderate) size. Of course, if the problem descriptions are almost always small, such computational complexity results are irrelevant. However, we usually want to deal with more than 20 concepts or 10 default rules. So we should consider the possibility of worst cases for moderately sized problem descriptions.

One way to exclude worst cases is to restrict the expressiveness of the representation language the reasoning task has to deal with. Brachman and Levesque, for example, showed that excluding a particular operator from a terminological logic reduces the complexity of reasoning from NP-hardness to polynomiality [Brachman and Levesque, 1984; Levesque and Brachman, 1987]. Subsequent investigations along this line [Donini *et al.*, 1991] have shown that requiring polynomiality of the inference algorithm leads to a severe restriction on the possible constructs one can use.

Although there have been strong arguments about the usefulness of achieving efficiency by restricting the expressiveness [Doyle and Patil, 1991], there seems to be nevertheless a consensus that it is useful to analyze special cases of general reasoning patterns that can be solved more easily than the general problem, provided the special cases are relevant. Moreover, *restricting* the expressiveness can mean a number of things that are quite different from, for example, excluding a particular operator from a representation language.

For instance, instead of considering a representation language with less constructs, it makes sometimes sense to use a language with more constructs but with restrictions on the structure of allowed expressions. Donini et al [1994] show that enlarging a terminological logic with an epistemic operator for building concepts that are used as queries and restricting the forms of the query can indeed lead to a more natural reasoning task which is also more efficient.

The work by Myers and Konolige [1994] also extends the representational framework (first-order logic) in order to achieve efficiency. In this case, however, the aim is not to guarantee worst-case efficiency in all cases, but to provide special means for representing knowledge about one particular domain – spatial knowledge – that can be more naturally represented and more efficiently reasoned about using analogical representations, which are also much more restricted than general propositional representations. The main problem Myers and Konolige identify and solve is the integration of analogical reasoning with the general framework of reasoning in first-order logic.

The paper by Niemelä and Rintanen [1994] aims again at guaranteeing polynomial runtime in all cases by restricting expressive power. As in the cases above, however, they do not restrict the expressive power by disallowing logical operators in AEL theories, but they consider restrictions on the form of the theories. In particular, they show that reasoning in stratified AEL Horn theories can be done in polynomial time.

3.4 The Accuracy vs. Efficiency Tradeoff

If the expressiveness of a representation cannot be restricted, other means for getting timely answers are called for. Usually, one gives up on the quality or accuracy of an answer, for example, by restricting the processing time or by employing *incomplete* reasoning methods. While this may lead to the desired runtime behavior, it raises the question as to how far we can still trust answers from a representation and reasoning system. In other words, we are seeking a *principled* description of the reasoning capabilities of an incomplete reasoner.

Kautz and Selman [1991] addressed this problem by a "knowledge compilation" technique. They propose to compute (off-line) Horn theories that approximate the logical contents of a given arbitrary theory. As mentioned above, this approximation can lead to computational problems in itself [Kautz and Selman, 1994]. Kautz and Selman show that the approximating theory can become very large, and although there are sometimes ways around this problem, they can show that it is very unlikely that dense representations of a approximating Horn theory exist in all cases. Nevertheless, their approximation scheme appears to be interesting since instead of general Horn theories one may aim for more restricted forms of such theories which can be polynomially bounded in size.

Greiner and Schuurmans [1994] address the *multiple extension* problem of default reasoning [Reiter, 1987], which is known to be one source of computational complexity in default reasoning [Gottlob, 1992; Nebel, 1991]. They propose to approximate default reasoning by ordering the defaults linearly, where the particular order chosen is intended to be "optimally correct." As they show, it is not possible to compute such an ordering in polynomial time, but they approximate such an ordering by computing a locally optimal ordering.

The paper by Witteveen and Jonker [1994] applies a similar method to achieve tractability for revising logic programs. They show that a globally minimal revision cannot be computed in polynomial time, but a *locally* minimal revision can well be computed in polynomial time.

4 Outlook

The collection of papers in this book does certainly not give a complete overview of the research going on at providing foundations for knowledge representation and reasoning. For instance, probabilistic approaches are not represented at all. Nevertheless, the set of papers in this book covers a wide range of topics in the area of foundational KR&R research and highlights the common research methodology, namely, to analyze representation and reasoning tasks from a *logical* and *computational* perspective. As already mentioned in the Introduction, this research methodology does most probably not lead to any immediate benefit in the sense that we can build faster or better reasoning systems. However, by providing the theoretical underpinning for KR&R systems, this research will help us understand where and what the limits of representation and reasoning are and how we can guarantee a reasonable behavior of KR&R systems.

References

[AAAI-90, 1990] *Proceedings of the 8th National Conference of the American Association for Artificial Intelligence*, Boston, MA, August 1990. MIT Press.

[Adams, 1975] E. W. Adams. *The Logic of Conditionals*. Reidel, Dordrecht, Holland, 1975.

[Alchourrón *et al.*, 1985] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, June 1985.

[Allen *et al.*, 1991] J. A. Allen, R. Fikes, and E. Sandewall, editors. *Principles of Knowledge Representation and Reasoning: Proceedings of the 2nd International Conference*, Cambridge, MA, April 1991. Morgan Kaufmann.

[Allgayer and Franconi, 1994] Jürgen Allgayer and Enrico Franconi. Collective entities and relations in concept languages. In Lakemeyer and Nebel [1994].

[Baader and Hollunder, 1994] Franz Baader and Bernhard Hollunder. Computing extensions of terminological default theories. In Lakemeyer and Nebel [1994].

[Bettini, 1994] C. Bettini. A formalization of interval-based temporal subsumption in first order logic. In Lakemeyer and Nebel [1994].

[Boutilier, 1994] Craig Boutilier. Normative, subjunctive and autoepistemic defaults. In Lakemeyer and Nebel [1994].

[Brachman and Levesque, 1984] Ronald J. Brachman and Hector J. Levesque. The tractability of subsumption in frame-based description languages. In *Proceedings of the 4th National Conference of the American Association for Artificial Intelligence*, pages 34–37, Austin, TX, 1984.

[Brachman and Levesque, 1985] Ronald J. Brachman and Hector J. Levesque, editors. *Readings in Knowledge Representation*. Morgan Kaufmann, Los Altos, CA, 1985.

[Brachman *et al.*, 1989] R. Brachman, H. J. Levesque, and R. Reiter, editors. *Principles of Knowledge Representation and Reasoning: Proceedings of the 1st International Conference*, Toronto, ON, May 1989. Morgan Kaufmann.

[Brachman, 1990] Ronald J. Brachman. The future of knowledge representation. In AAAI-90 [1990], pages 1082–1092.

[Brewka, 1991] Gerhard Brewka. *Nonmonotonic Reasoning: Logical Foundations of Commonsense*. Cambridge University Press, Cambridge, UK, 1991.

[Console and Dupré, 1994] Luca Console and Daniele Dupré. Abductive reasoning with abstraction axioms. In Lakemeyer and Nebel [1994].

[Donini *et al.*, 1991] Francesco M. Donini, Maurizio Lenzerini, Daniele Nardi, and Werner Nutt. Tractable concept languages. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, pages 458–465, Sydney, Australia, August 1991. Morgan Kaufmann.

[Donini *et al.*, 1994] Francesco Donini, Maurizio Lenzerini, Daniele Nardi, Andrea Schaerf, and Werner Nutt. Queries, rules and definitions as epistemic sentences in concept languages. In Lakemeyer and Nebel [1994].

[Doyle and Patil, 1991] Jon Doyle and Ramesh S. Patil. Two theses of knowledge representation: Language restrictions, taxonomic classification, and the utility of representation services. *Artificial Intelligence*, 48(3):261–298, April 1991.

[Garey and Johnson, 1979] Michael R. Garey and David S. Johnson. *Computers and Intractability—A Guide to the Theory of NP-Completeness*. Freeman, San Francisco, CA, 1979.

[Gottlob, 1992] Georg Gottlob. Complexity results for nonmonotonic logics. *Journal for Logic and Computation*, 2(3), 1992.

[Gottlob, 1994] Georg Gottlob. The power of beliefs or translating default logic into standard autoepistemic logic. In Lakemeyer and Nebel [1994].

[Greiner and Schuurmans, 1994] Russell Greiner and Dale Schuurmans. Learning an optimally accurate representation system. In Lakemeyer and Nebel [1994].

[Hayes, 1977] Patrick J. Hayes. In defence of logic. In *Proceedings of the 5th International Joint Conference on Artificial Intelligence*, pages 559–565, Cambridge, MA, August 1977.

[Junker and Konolige, 1990] Ulrich Junker and Kurt Konolige. Computing extensions of autoepistemic and default logics with a truth maintenance system. In AAAI-90 [1990], pages 278–283.

[Kakas, 1994] A. C. Kakas. Default reasoning via negation as failure. In Lakemeyer and Nebel [1994].

[Kalinski, 1994] Jürgen Kalinski. Weak autoepistemic reasoning and well-founded semantics. In Lakemeyer and Nebel [1994].

[Katsuno and Mendelzon, 1991] Hirofumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. In Allen et al. [1991], pages 387–394.

[Kautz and Selman, 1994] Henry Kautz and Bart Selman. Forming concepts for fast inference. In Lakemeyer and Nebel [1994].

[Lakemeyer and Nebel, 1994] Gerhard Lakemeyer and Bernhard Nebel, editors. *Foundations of Knowledge Representation and Reasoning.* Springer-Verlag, Berlin, Heidelberg, New York, 1994.

[Levesque and Brachman, 1987] Hector J. Levesque and Ronald J. Brachman. Expressiveness and tractability in knowledge representation and reasoning. *Computational Intelligence*, 3:78–93, 1987.

[Levesque, 1988] Hector J. Levesque. Logic and the complexity of reasoning. *Journal of Philosophical Logic*, 17:355–389, 1988.

[Lewis, 1973] David K. Lewis. *Counterfactuals.* Harvard University Press, Cambridge, MA, 1973.

[Lin, 1994] Yuen Q. Lin. A common-sense theory of time. In Lakemeyer and Nebel [1994].

[McCarthy, 1968] John McCarthy. Programs with common sense. In M. Minsky, editor, *Semantic Information Processing*, pages 403–418. MIT Press, Cambridge, MA, 1968.

[McCarthy, 1980] John McCarthy. Circumscription—a form of non-monotonic reasoning. *Artificial Intelligence*, 13(1–2):27–39, 1980.

[McDermott, 1978] Drew V. McDermott. Tarskian semantics, or no notation without denotation! *Cognitive Science*, 2(3):277–282, July 1978.

[Moore, 1985] Robert C. Moore. Semantical considerations on nonmonotonic logic. *Artificial Intelligence*, 25:75–94, 1985.

[Myers and Konolige, 1994] Karen Myers and Kurt Konolige. Reasoning with analogical representations. In Lakemeyer and Nebel [1994].

[Nebel *et al.*, 1992] B. Nebel, W. Swartout, and C. Rich, editors. *Principles of Knowledge Representation and Reasoning: Proceedings of the 3rd International Conference*, Cambridge, MA, October 1992. Morgan Kaufmann.

[Nebel, 1991] Bernhard Nebel. Belief revision and default reasoning: Syntax-based approaches. In Allen et al. [1991], pages 417–428.

[Nejdl and Banagl, 1994] Wolfgang Nejdl and Markus Banagl. Asking about possibilities - revision and update semantics for subjunctive queries (extended report). In Lakemeyer and Nebel [1994].

[Niemelä and Rintanen, 1994] Ilkka Niemelä and Jussi Rintanen. On the impact of stratification on the complexity of nonmonotonic reasoning. In Lakemeyer and Nebel [1994].

[Pearl, 1988] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann, San Mateo, CA, 1988.

[Reiter, 1980] Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1):81–132, April 1980.

[Reiter, 1987] Raymond Reiter. Nonmonotonic reasoning. *Annual Review of Computing Sciences*, 2, 1987.

[Schmidt-Schauß and Smolka, 1991] Manfred Schmidt-Schauß and Gert Smolka. Attributive concept descriptions with complements. *Artificial Intelligence*, 48:1–26, 1991.

[Schwind and Risch, 1991] C. Schwind and V. Risch. A tableau-based characterisation for default logic. In R. Kruse and P. Siegel, editors, *Symbolic and Quantitative Approaches to Uncertainty, Proceedings of the European Conference ECSQAU*, pages 310–317, Marseilles, France, 1991. Springer-Verlag.

[Selman and Kautz, 1991] Bart Selman and Henry Kautz. Knowledge compilation using Horn approximations. In *Proceedings of the 9th National Conference of the American Association for Artificial Intelligence*, pages 904–909, Anaheim, CA, July 1991. MIT Press.

[Shoham and Cousins, 1994] Yoav Shoham and Steve B. Cousins. Logics of mental attitudes in AI — a very preliminary survey. In Lakemeyer and Nebel [1994].

[Smith, 1982] Brian C. Smith. *Reflection and Semantics in a Procedural Language.* PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 1982. Report MIT/LCS/TR-272.

[Weydert, 1994] Emil Weydert. Hyperrational conditionals - monotonic reasoning about nested default conditionals. In Lakemeyer and Nebel [1994].

[Witteveen and Jonker, 1994] Cees Witteveen and Catholijn M. Jonker. Revision by expansion in logic programs. In Lakemeyer and Nebel [1994].

This article was processed using the LaTeX macro package with LLNCS style