

Logics for Knowledge Representation

Bernhard Nebel

Albert-Ludwigs-Universität Freiburg, Germany

1 Introduction

Knowledge representation and reasoning plays a central role in Artificial Intelligence. Research in Artificial Intelligence (henceforth AI) started off by trying to identify the *general mechanisms* responsible for intelligent behavior. However, it quickly became obvious that general and powerful methods are not enough to get the desired result, namely, intelligent behavior. Almost all tasks a human can perform which are considered to require intelligence are also based on a huge amount of knowledge. For instance, understanding and producing natural language heavily relies on knowledge about the language, about the structure of the world, about social relationships etc.

One way to address the problem of representing knowledge and reasoning about it is to use some form of logic. While this seems to be a natural choice, it took a while before this “logical point of view” became the prevalent approach in the area of knowledge representation. Below, we will give a brief sketch of how the field of knowledge representation evolved and what kind logical methods have been used.

In particular, we will argue that the important point about using formal logic is the *logical method*.

2 Logic-Based Knowledge Representation: A Historical Account

McCarthy (1968) stated very early on that mathematical, formal logic appears to be a promising tool for achieving human-level intelligence on computers. In fact, this is still McCarthy's (2000) vision, which he shares with many researchers in AI. However, in the early days of AI, there were also a number of researchers with a completely different opinion. Minsky (1975), for example, argued that knowledge representation formalisms should be flexible and informal. Moreover, he argued that the logical notions of *correctness* and *completeness* are inappropriate in a knowledge representation context.

While in those days heated arguments of the suitability of logic were exchanged, by the end of the eighties, the logical perspective seem to have gained the upper hand (Brachman 1990). During the nineties almost all research in the area of knowledge representation and reasoning was based on formal, logical methods as demonstrated by the papers published in the bi-annual international conference on *Principles of Knowledge Representation and Reasoning*, which started in 1989.

It should be noted, however, that two perspectives on logic are possible. The first perspective, taken by McCarthy (1968), is that logic should be used to *represent* knowledge. That is, we use logic as the representational and reasoning tool inside

the computer. Newell (1982) on the other hand proposed in his seminal paper on the *knowledge level* to use logic as a formal tool to *analyze* knowledge. Of course, these two views are not incompatible. Furthermore, once we accept that formal logic should be used as a tool for analyzing knowledge, it is a natural consequence to use logic for representing knowledge and for reasoning about it as well.

3 Knowledge Representation Formalisms and Their Semantics

Saying that logic is used as the main formal tool does not say which kind of logic is used. In fact, a large variety of logics (Gabbay, Hogger and Robinson 1995) have been employed or developed in order to solve knowledge representation and reasoning problems. Often, one started with a vaguely specified problem, developed some kind knowledge representation formalism without a formal semantics, and only later started to provide a formal semantics. Using this semantics, one could then analyze the complexity of the reasoning problems and develop sound and complete reasoning algorithms. I will call this the *logical method*, which proved to be very fruitful in the past and has a lot of potential for the future.

3.1 Description Logics

One good example for the evolution of knowledge representation formalisms is the development of *description logics*, which have their roots in so-called *structured inheritance networks* formalisms such as KL-ONE (Brachman 1979). These networks were originally developed in order represent word meanings. A *concept*

node connects to other concept nodes using *roles*. Moreover, the roles could be structured as well. These networks permits for, e.g., the definition of the concept of a bachelor.

Later on, these structured inheritance networks were formalized as so-called *concept languages*, *terminological logics*, or *description logics*. Concepts were interpreted as unary predicates, roles as binary relations, and the connections between nodes as so-called *value restrictions*. This leads for most such description logics to a particular fragment of *first-order predicate logic*, namely, the fragment L_2 . In this fragment only two different variable symbols are used. As it turns out, this is a *decidable fragment* of first-order logic.

However, some of the more involved description logics go beyond L_2 . They contain, e.g., relational composition or transitive closure. As it turns out, such description logics can be understood as variants of multi-modal logics (Schild 1991), and decidability and complexity results from these multi-modal logics carry over to the description logics. Furthermore, description logics are very close to *feature logics* as they are used in unification-based grammars. In fact, description logics and feature logics can be viewed as members of the same family of representation formalisms (Nebel and Smolka 1990).

All these insights, i.e., determination of decidability and complexity as well as the design of decision algorithms (e.g. Donini, Lenzerini, Nardi and Nutt 1991), are based on the rigorous formalization of the initial ideas. In particular, it is not just

one logic that it is used to derive these results, but it is the *logical method* that led to the success. One starts with a specification of how expressions of the language or formalism have to be interpreted in formal terms. Based on that one can specify when a set of formulae *logically implies* a formula. Then one can start to find similar formalisms (e.g. modal logics) and prove equivalences and/or one can specify a method to *derive* logically entailed sentences and prove them to be *correct* and *complete*.

3.2 Nonmonotonic Logics

Another interesting area where the logical method has been applied is the development of the so-called *non-monotonic logics*. These are based on the intuition that sometimes a logical consequence should be retracted if new evidence becomes known. For example, we may assume that our car will not be moved by somebody else after we have parked it. However, if new information becomes known, such as the fact that the car is not at the place where we have parked it, we are ready to drop the assumption that our car has not been moved.

This general reasoning pattern was used quite regularly in early AI systems, but it took a while before it was analyzed from a logical point of view. In 1980, a special issue of the *Artificial Intelligence* journal appeared, presenting different approaches to non-monotonic reasoning, in particular Reiter's (1980) *default logic* and McCarthy's (1980) *circumscription* approach.

A disappointing fact about nonmonotonic logics appears to be that it is very difficult to formalize a domain such that one gets the intended conclusions. In particular, in the area of *reasoning about actions*, McDermott (1987) has demonstrated that the straightforward formalization of an easy temporal projection problem (the “Yale shooting problem”) does not lead to the desired consequences. However, it is possible to get around this problem. Once all underlying assumptions are spelled out, this and other problem can be solved (Sandewall 1994).

It took more than a decade before people started to analyze the computational complexity (of the propositional versions) of these logics. As it turned out, these logics are usually somewhat more difficult than ordinary propositional logic (Gottlob 1992). This, however, seems tolerable since we get much more conclusions than in standard propositional logic.

Right at the same time, the tight connection between nonmonotonic logic and *belief revision* (Gärdenfors 1988) was noticed. Belief revision – modeling the evolution of beliefs over time – is just one way to describe how the set of nonmonotonic consequences evolve over time, which leads to a very tight connection on the formal level for these two forms of nonmonotonicity (Nebel 1991). Again, all these results and insights are mainly based on the logical method to knowledge representation.

4 Outlook

The above description of the use of logics for knowledge representation is necessarily incomplete. For instance, we left out the area of qualitative temporal and spatial reasoning completely. Nevertheless, one should have got an idea of how logics are used in the area of knowledge representation. As mentioned, it is the idea of providing knowledge representation formalisms with formal (logical) semantics that enables us to communicate their meaning, to analyze their formal properties, to determine their computational complexity, and to devise reasoning algorithms.

While the research area of knowledge representation is dominated by the logical approach, this does not mean that all approaches to knowledge representation must be based on logic. Probabilistic (Pearl 1988) and decision theoretic approaches, for instance, have become very popular lately. Nowadays a number of approaches aim at unifying decision theoretic and logical accounts by introducing a qualitative version of decision theoretic concepts (Benferhat, Dubois, Fargier, Prade and Sabbadin 2000). Other approaches (Boutilier, Reiter, Soutchanski and Thrun 2000) aim at tightly integrating decision theoretic concepts such as Markov decision processes with logical approaches, for instance. Although this is not pure logic, the two latter approaches demonstrate the generality of the *logical method*: specify the formal meaning and analyze!

Bibliography

- Allen, J. A., Fikes, R. and Sandewall, E. (eds): 1991, *Principles of Knowledge Representation and Reasoning: Proceedings of the 2nd International Conference (KR-91)*, Morgan Kaufmann, Cambridge, MA.
- Benferhat, S., Dubois, D., Fargier, H., Prade, H. and Sabbadin, R.: 2000, Decision, nonmonotonic reasoning and possibilistic logic, *in* Minker (2000), pp. 333–360.
- Boutilier, C., Reiter, R., Soutchanski, M. and Thrun, S.: 2000, Decision-theoretic, high-level agent programming in the situation calculus, *Proceedings of the 17th National Conference of the American Association for Artificial Intelligence (AAAI-2000)*, MIT Press, Austin, TX.
- Brachman, R. J.: 1979, On the epistemological status of semantic networks, *in* N. V. Findler (ed.), *Associative Networks: Representation and Use of Knowledge by Computers*, Academic Press, New York, NY, pp. 3–50.
- Brachman, R. J.: 1990, The future of knowledge representation, *Proceedings of the 8th National Conference of the American Association for Artificial Intelligence (AAAI-90)*, MIT Press, Boston, MA, pp. 1082–1092.
- Donini, F. M., Lenzerini, M., Nardi, D. and Nutt, W.: 1991, The complexity of concept languages, *in* Allen, Fikes and Sandewall (1991), pp. 151–162.
- Gabbay, D. M., Hogger, C. J. and Robinson, J. A. (eds): 1995, *Handbook of Logic in Artificial Intelligence and Logic Programming – Vol. 1–5*, Oxford University Press, Oxford, UK.

- Gärdenfors, P.: 1988, *Knowledge in Flux—Modeling the Dynamics of Epistemic States*, MIT Press, Cambridge, MA.
- Gottlob, G.: 1992, Complexity results for nonmonotonic logics, *Journal for Logic and Computation* **2**(3), 397–425.
- McCarthy, J.: 1968, Programs with common sense, in M. Minsky (ed.), *Semantic Information Processing*, MIT Press, Cambridge, MA, pp. 403–418.
- McCarthy, J.: 1980, Circumscription—a form of non-monotonic reasoning, *Artificial Intelligence* **13**(1–2), 27–39.
- McCarthy, J.: 2000, Concepts of logical AI, in Minker (2000), pp. 37–58.
- McDermott, D. V.: 1987, A critique of pure reason, *Computational Intelligence* **3**(3), 151–160.
- Minker, J. (ed.): 2000, *Logic-Based Artificial Intelligence*, Kluwer, Dordrecht, Holland.
- Minsky, M.: 1975, A framework for representing knowledge, in P. Winston (ed.), *The Psychology of Computer Vision*, McGraw-Hill, New York, NY, pp. 211–277.
- Nebel, B.: 1991, Belief revision and default reasoning: Syntax-based approaches, in Allen et al. (1991), pp. 417–428.
- Nebel, B. and Smolka, G.: 1990, Representation and reasoning with attributive descriptions, in K.-H. Bläsius, U. Hedtstück and C.-R. Rollinger (eds), *Sorts and Types in Artificial Intelligence*, Vol. 418 of *Lecture Notes in Artificial Intelligence*, Springer-Verlag, Berlin, Heidelberg, New York, pp. 112–139.

Newell, A.: 1982, The knowledge level, *Artificial Intelligence* **18**(1), 87–127.

Pearl, J.: 1988, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, San Francisco, CA.

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

Sandewall, E.: 1994, *Features and Fluents*, Oxford University Press, Oxford, UK.

Schild, K.: 1991, A correspondence theory for terminological logics: Preliminary report, *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91)*, Morgan Kaufmann, Sydney, Australia, pp. 466–471.