

Knowledge Representation and Reasoning – The Theoretical Side of AI

Bernhard Nebel¹

Abstract. Representing knowledge and reasoning with it is at the heart of most AI systems. Nevertheless, the field of Knowledge Representation and Reasoning (KR&R) does not seem to have much impact on practical work in AI. I will try to point out the reasons for the discrepancy and I will argue that KR&R can be understood as Theoretical Artificial Intelligence. Further, I point out a number of fruitful applications of KR&R techniques.

1 INTRODUCTION

In almost all AI systems it is necessary to represent knowledge about the world. Even in robotics, almost everybody seems to agree nowadays that it is not enough to have the world as its own representation.

From this observation one might conclude that the field of *Knowledge Representation and Reasoning* (KR&R) must be one of the most important sub-fields. As it is often with research, however, most of the research results are not directly relevant for practice. The reason for that is simple. In KR&R research we try to explore the border line of what is still representable and what is still efficiently inferable. These two questions are, however, often only marginally interesting in practical systems. Nevertheless, KR&R research provides two important lines of results that influence the practice of KR&R. First, complexity results provide us with guarantees of what can be done efficiently – and point out where heuristic methods should start. Second, new representation formalisms together with algorithms can help to solve problems more easily.

2 A BRIEF HISTORY

KR&R started as a sub-activity of more problem-specific issues such as natural language understanding. In those pioneering days a typical KR&R paper was on how to distinguish between *de re* and *de dicto* readings in some semantic network formalism. According to Brachman [1], this changed with the beginning of the eighties, and KR&R became formal and technical. Now questions such as “what is the *formal semantics* of your formalism” and “what is the *computational complexity* of reasoning in your formalism” became important. It is this part of KR&R that builds some of the theoretical foundations of AI and may become *Theoretical AI*. However, there was and is also the more practically oriented part of KR&R concerned with building systems and creating domain models.

3 WHERE DOES KR&R STAND NOW?

While there was a big change from the seventies to the eighties, the nineties did not establish a revolutionary new KR&R trend. Very much as predicted by Brachman [1], *logic and rigor, nonmonotonic reasoners, probability and statistics*, as well as *ontologies* are still important issues.

One interesting new development is that theoretically oriented researchers start to implement systems! At a workshop on *Logic-Based AI* organized by Jack Minker [3], all evenings up to 10 o'clock pm were spent on system demonstrations! There are very fast propositional reasoners, nonmonotonic reasoners such as SMOBELS [4], cognitive robotics reasoners, etc. Some are even *used*, for instance as a backend engine for a planner. Furthermore, we now see even industrial applications of formalisms and techniques developed by the KR community. For example, description logics [2] are now used by Lucent in a configuration system [6], and Allen's interval calculus is used to describe document layouts in order to support automatic mail sorting [5].

4 WHERE ARE WE HEADING?

KR&R is definitely the sub-field of AI which exercises the most technical rigor; and one may hope that some of that spirit spills over to other sub-fields of AI. Of course, some of the KR&R work may also need some engineering spirit. As mentioned, these things tend to become better (both ways). I see two grand challenges for KR&R in the near-time future. One is to demonstrate that cognitive robotics is in fact part of robotics. The second challenge is the application of KR&R techniques to the web (but then, this is a generic challenge for all Computer Science disciplines).

REFERENCES

- [1] R. J. Brachman, ‘The future of knowledge representation’, in *Proc.AAAI-90*, pp. 1082–1092, Boston, MA (1990).
- [2] F. M. Donini, M. Lenzerini, D. Nardi, and W. Nutt, ‘Tractable concept languages’, in *Proc. IJCAI-91*, pp. 458–465, Sydney, Australia (1991).
- [3] *Logic-Based Artificial Intelligence*, ed., J. Minker, Kluwer, Dordrecht, Holland. To appear.
- [4] I. Niemelä and P. Simons, ‘Smodels - an implementation of the stable model and well-founded semantics for normal logic programs’, in *Proc. ICLP/NMR-97*, pp. 420–429, Dagstuhl, Germany (1997). Springer-Verlag.
- [5] H. Walischewski, ‘Learning regions of interest in postal automation’, in *Proc. ICDAR-99*, Bangalore, India (1999).
- [6] J. R. Wright, E. S. Weixelbaum, G. T. Vesonder, K. E. Brown, 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’, *The AI Magazine*, **14**(3), 69–80. (1993).

¹ Albert-Ludwigs-Universität Freiburg, Institut für Informatik, Georges-Köhler-Allee, Geb. 52, D-79110 Freiburg, Germany, email: nebel@informatik.uni-freiburg.de