

# What Does Knowledge Representation Have to Say to Artificial Intelligence?

Invited Talk Abstract

**David W. Etherington**

Computational Intelligence Research Laboratory  
1269 University of Oregon  
Eugene, OR 97403-1269  
ether@cirl.uoregon.edu

## Abstract

In recent years, the subarea of Knowledge Representation and Reasoning (KR) has become more and more of a discipline unto itself, focusing on artificial problems while other areas of AI have tended to develop their own representations and algorithms. There are signs that this is changing, however. This talk will explore what the current state of KR has to offer to AI.

On the face of it, it would seem that nothing could be more central to Artificial Intelligence than knowledge representation and reasoning. It is somewhat paradoxical, therefore, to note that for many years KR has become more and more of a sub-discipline unto itself, with little commerce with the rest of AI. KR researchers often appear to care little for the representational and inferential problems faced by researchers in traditional “consumers” of KR technology (e.g., robotics, vision, and natural language processing), and progress in KR frequently appears to have little impact on those disciplines.

This divergence is partially explainable by the different goals of the research communities: much KR research has been focused on developing a general understanding of representation and reasoning for their own sakes. Based on an underlying assumption that “one theory fits all”, researchers have sought the grail of a universal theory of reasoning, largely ignoring the particulars of what tasks the reasoning will support. Given this perspective, and faced with the overwhelming complexity of the problems of representing and using arbitrary knowledge about the world, research has tended to choose clean, simple (a.k.a. “toy”) problems, and focus exclusively on what linguists call “competence models” (models that express what an idealized, omniscient, agent with infinite computational resources would do) at the expense of “performance models” (models that express the capabilities of real, finite,

agents with time pressure). As the underlying logics and mathematics have grown more sophisticated, many have been attracted by the beauty of the formal systems, leading to a further erosion of the connection to practical problems and systems.

The resulting focus on severely restricted problems and computationally unrealistic reasoning paradigms has led to the growing disconnection between KR and the rest of AI. At times, the term “vacuous theorizing” has circulated rather widely. While less universally applicable than some critics suggested, this critique has not been totally without merit. Some blame may lie with those who would prefer not to make the effort required to understand complicated formal theories, but much of the fault has been with KR researchers who have tended to focus on technical aspects of KR problems without ever returning to ground them in practical AI.

In some respects, this situation has been detrimental to the field, not only in reducing its practical impact, but in losing the vitality and sense of practical urgency that working on real problems imbues to a research effort. Nonetheless, there *has* been significant progress in KR, and many exciting technologies are developing, and there are signs that KR is ready to return to the fold. Both because of the growing maturity of the “competence” models in many areas, and because of a realization that practical computational concerns bring problems that are interesting in their own right, many in the field are trying to identify *real* representational problems that affect those trying to make progress in other areas, communicate ways that work in KR can have practical implications in other areas, and focus on KR research aimed at problems whose solution will be of interest and utility outside KR. In this talk, hopeful signs and concrete evidence will be discussed, in the hope of both reporting on and encouraging this rapprochement.

---

Copyright © 1997, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

This work was supported in part by NSF grant IRI-94-12205, and DARPA/Rome Labs contracts F30602-93-C-0031 and F30602-95-1-0023, and AFOSR contract F49620-96-1-0335.