

Knowledge representation and reasoning: The view from cognitive psychology

JOHN FOX and PAUL KRAUSE

The knowledge engineering community has been working on the design of schemes for knowledge representation and reasoning for more than two decades. Much of this work, particularly work on the development of expert systems, explicitly or implicitly assumes that artificial knowledge-based systems emulate to some degree the natural knowledge representation and reasoning methods of human problem solvers and decision makers.

Experimental psychologists and other cognitive scientists have been studying the properties of natural cognition for even longer, indeed for much of this century. Their findings indicate not only that some of our engineering assumptions about the nature of human expertise may be a little simplistic, and also that one needs to be very careful about those aspects of human knowledge and expertise one should attempt to emulate.

This special issue of the *Review* presents papers by experimental psychologists who have worked extensively on expertise, decision making and reasoning under uncertainty, all topics that overlap strongly with the interests of expert systems and AI researchers and developers. Their reviews of work on these topics are instructive for those of us who are interested in the natural counterparts of the artificial mechanisms and techniques we use.

The first paper, "Cognitive expertise research and knowledge engineering" by Fergus Bolger, provides an overview of psychological studies of expertise, drawing attention to weaknesses in our criteria for defining an "expert". He summarises our current understanding of the cognitive processes that underpin expertise and identifies some implications for knowledge engineers. In "Bias in human judgement under uncertainty?", Peter Ayton and Eva Pascoe focus on an aspect of expertise which is a major issue for expert systems designers; uncertain inference methods. They discuss important doubts about the competence of human judgement when compared with the behaviour prescribed by normative mathematical theory, and many subtleties of human understanding which are not well reflected in current knowledge technologies. Finally, in a related paper reviewing "Human linguistic probability processing", Tom Wallsten and David Budescu look at issues concerning the intuitive representation of uncertainty, and, in particular, how we use natural language concepts to reason with and communicate uncertainty about our beliefs and inferences. They close with some principles that summarize the cognitive processes that underly human uncertain reasoning and decision making. It is interesting to consider these principles in the context of the probabilities, certainty factors and non-monotonic logics which are the foci of AI research.

Although one frequently comes across remarks in the AI and expert systems literature to the effect that "expert systems emulate human expertise", such claims are rarely examined in depth. The findings of psychologists described in these papers indicate both that their validity can be seriously questioned, and the desirability of a design strategy based on emulation is open to considerable debate.