

Methodologies for knowledge engineering

JOHN FOX

The ups and downs of expert systems are well known. In the early eighties they were up and coming, commentary was upbeat and the ‘conventional computing community’ was upset. However, the late eighties saw a downturn in the credibility of expert system technology, a downgrading in the priority given by companies and individuals to knowledge-based system development, and a downturn in the fortunes of the fledgling expert systems industry. And there, one might have thought, things would rest; a period of excitement and inflated expectations followed by disappointment, lessons learned and the recognition that neural nets (or whatever) are the technology that is *really* going to change everything.

Apparently, things have not rested there. Although there has been a revision of expectations and a shakeout in the AI industry, evidence is now growing that knowledge-based systems are a more successful technology than they seemed a couple of years ago. A recent report by Touche Ross*, sponsored by the UK Government’s Department of Industry, arrives at some interesting conclusions. Analysis of some 200 telephone interviews with user organizations, suggests that 60% of the UK organizations questioned are so convinced of the benefits of knowledge-based systems technology that they are developing or planning to develop further applications. One in ten of the organizations are apparently working on five or more applications!

Five years ago we were all anticipating an AI winter, and it came, but now Spring seems to be in the air. This is all the more surprising, since the computer industry generally is presently enduring one of the coldest winters on record. No doubt we shall all have our pet explanations for the apparent acceptance of knowledge technology. Mine has two parts, and is perhaps shared by many readers. First, knowledge technology introduced some *really new and powerful ideas*, but these were temporarily eclipsed by misunderstanding and misuse. Second, we had to *assimilate these new ideas into main-stream computing*; we had to learn how to combine knowledge based systems and AI tools with existing software and systems, to resist reinventing wheels and build on the considerable (painfully acquired) expertise of the software industry.

Knowledge engineering is truly becoming an engineering discipline; there is now as much interest in *design* as there is in the traditional, but relatively uncontrolled, procedures of *knowledge acquisition*. In this issue of the *Review* we publish a paper by Preece *et al.* on ‘Principles and practice in verifying rule-based systems’. This paper reflects, and reflects on, the rapid growth of interest in verification, validation and test of knowledge based systems. Over the last eight years the contents of *The Knowledge Engineering Review* have naturally mirrored the maturing of knowledge engineering. One can see in the discussions a move away from *ad hoc* compilations of knowledge bases towards a more careful and professional approach to engineering knowledge, and an increasingly formal approach to design. The aims of the *Review* have always been ‘to foster the development of knowledge engineering as a technical discipline . . .’. I hope that its mix of authoritative, foundational surveys; critical reviews of developing technologies; case studies, and the timely provision of information about current literature (such as the bibliography on validation of knowledge bases in this issue) will contribute to the maturing of the field.

Having welcomed this growing maturity we should, however, sound a small note of warning. A significant part of the progress in the field is due, it must be said, to the adoption of ideas from conventional software engineering. Indeed many of the ‘methodologies’ of knowledge engineering

*‘Knowledge based systems: Survey of UK application’, London: Department of Industry, 1992. Available from Touche Ross & Co, 133 Fleet Street, London, UK.

that are touted were directly inspired and strongly influenced by software engineering methodologies, and the Touche Ross report firmly advises managers thinking of adopting KBS technology to 'follow software engineering disciplines'. But the above puff about fostering knowledge engineering as a technical discipline continues '... with strong roots in Artificial Intelligence and cognitive science'. I am delighted that knowledge-based systems are an accepted part of the software specialists' armoury, but we should remember that the technology really is different from conventional software. If we go too far in emphasizing conventional practices then another round of misunderstanding and misuse may set in. We would run the risk of being slow to exploit, and even to recognize, the new possibilities created within AI.