Editorial

The problem of knowledge elicitation

An earlier issue of *The Knowledge Engineering Review* dealt with a number of aspects of technology transfer (vol. 2, issue 3). In putting that special issue together I wanted to know how research in expert systems, and AI more generally, was actually translated into commercial engineering and industrial practices. Knowledge engineering was, after all, born in an explosion of hype and many things, both thoughtful and ill-considered, were said about its nature, its revolutionary implications and its commercial potential. However, as I remarked in the editorial for that issue, no one actually *knew* how research ideas were evolving in engineering practice, though evolving they certainly are.

One area which has seemed critical to the success of knowledge based systems since the beginning, is knowledge elicitation. Notwithstanding its importance the processes and effects of transferring research results in this field seem, to this observer at least, more poorly understood than in any other area of applied AI. The early slogan that "knowledge acquisition is the principal bottleneck for expert systems development" has motivated an army of research workers looking at methods and methodologies for knowledge elicitation, but exactly how, or if, industry as a whole is using the results of their research is very unclear.

In this issue we have two papers which seem almost perfect poles on the research-application dimension. Ian Neale's paper is a comprehensive, and I think very valuable, overview of current research in knowledge elicitation until about the middle of 1988. Neale covers an enormous amount of ground, showing the variety of thinking and proposals for systematic knowledge acquisition. What is missing however (though this is not a criticism of Neale's paper but rather a comment on the field) is a picture of how the ideas are being used in practice. This is in contrast to Alex D'Agapeyeff's commentary which is almost all about the practices which have evolved in one company and the worry that research is not reflecting the realities and constraints of business. d'Agapeyeff is acutely conscious that he is not a professional research worker but, as he puts it "our experience constructing some 2,000 knowledge bases in a single family of software is a potent asset". His paper is an appeal for straightforwardness and common sense in thinking about knowledge engineering methods, and a suggestion that some of the evolution in this field may not represent clear technical progress.

Reading these two papers it is obvious that the preoccupations of many in the research community and at least some of those trying to forge engineering practice in a new branch of software are very different. They are not, it seems to me, necessarily incompatible. But the gulf between those communities is so large as to be a serious concern. Researchers might perhaps complain that their findings and techniques are not being fully understood or properly exploited by industry (we certainly do in most other areas of AI) but we might also consider the possible dangers of getting out of touch with industrial reality and being at times just a wee bit too cute.