

# Formalizing knowledge and expertise: where have we been and where are we going?

JOHN FOX

*Department of Engineering Science, University of Oxford, Oxford, UK;*  
*e-mail: john.fox@eng.ox.ac.uk*

The first issue of the *Knowledge Engineering Review* (KER) was published privately in 1984 to demonstrate the kind of material that the new journal would cover, and as a prototype for discussion by the editorial board and potential publishers. The first article in the issue was ‘A short account of knowledge engineering’ in which my aim was to set out a view of the general nature of knowledge engineering and how it might develop into an engineering discipline with a firm science base. The other two papers in the issue reflected some other priority goals for the field as I saw them; ‘Logic programming in the fifth generation’<sup>1</sup> by Bob Kowalski of Imperial College and ‘The commercial applications of expert systems technology’ by Tim Johnson, founder of Ovum Ltd. These papers were not peer-reviewed for the issue, but I chose them because they represented to me three important perspectives on knowledge engineering that needed to be addressed if we were to turn the new and somewhat *ad hoc* field of expert systems into a ‘proper’ engineering specialty. The three perspectives were: what *is* ‘knowledge engineering’ exactly? How can we give it sound theoretical foundations? How can we ensure its relevance to the needs of practical users whose concerns are rather different from those of academics and theoreticians?

I am grateful to the present editors of the *Review* for the opportunity to make a contribution to the 25th anniversary issue. In this short piece, I would like to revisit those early questions and discuss how knowledge engineering has developed with respect to them.

## 1 What is knowledge engineering?

Oddly, my original editorial did not give a definition of knowledge engineering, but the consensus at the time was probably pretty clear. Engineering, it would be agreed, is ‘the practical application of science to commerce or industry’ and knowledge is ‘the expertise and skills acquired by a person through experience or education, and the theoretical or practical understanding of a subject’. Therefore, knowledge engineering is ‘The engineering discipline that involves integrating knowledge into computer systems in order to solve complex problems normally requiring a high level of human expertise’.

For me the science, which is being applied is cognitive science. This has various definitions including a human-centred meaning such as ‘the study of human intelligence and of the symbol-processing nature of cognition’ and a more neutral one such as ‘the scientific project dedicated to

<sup>1</sup> Younger readers may not be aware that the announcement of the ‘Fifth Generation’ computer project in Japan in 1982 created an unprecedented level of interest in AI and, its application wing, knowledge-based expert systems ([http://en.wikipedia.org/wiki/Fifth\\_generation\\_computer](http://en.wikipedia.org/wiki/Fifth_generation_computer)). The establishment of the British Computer Society Specialist Group in Expert Systems (SGES) was a UK manifestation of this interest, and the proposal to establish a review journal aimed at fostering the development of an engineering discipline to underpin developments in the field was developed under the auspices of SGES and with the support of several of the members of the group’s management board, notably its first two Chairmen Donald Michie and Alex d’Agapayeff.

understanding the processes and representations underlying intelligent action' which accommodates insights from both artificial and natural intelligence.

## 2 Where have we been?

A lot has happened since 1984, so inevitably any attempt to summarize where we have been over the last 25 years will be somewhat selective. I believe there have been significant advances though perhaps a few retreats as well. To get an overview of these developments I have sampled papers from the first 6 years of the KER's publication as a Cambridge University Press title, and similarly for the last 6 years. In the table below, I have attempted to group the topics under headings that relate them to the three perspectives introduced above: namely the *themes* of the field (what distinguishes the knowledge engineering approach from other branches of, say, software engineering?); *applications* (what new classes of application have been facilitated by knowledge-based systems?), and *foundations* (what novel concepts and theories have emerged as a consequence of practical experience or with the aim of taking the field forward?). The sampling process was of course limited only to the KER. It was also not particularly rigorous and there is room for debate about the specific headings I have used. However, there are a few patterns in the data, which I believe are indicative of real trends in the field.

Under the heading of *themes*, the characteristic signature of knowledge engineering is that it uses explicit, symbolic representations of knowledge in order to emulate human cognitive capabilities like decision-making, planning and learning. This seems to have been sustained over the 25 years of publication.

Perhaps not surprisingly there has been a substantial growth in the number of papers about agents and particularly multi-agent systems in recent years. Given that agent technology is so fashionable in mainstream artificial intelligence (AI) and computer science this is exactly what one would expect, though given the engineering focus of the *Review* I would have liked more discussion of the practical successes (and failures?) of agent technology than I have seen in the general literature. I would imagine that there is a strong appetite for a comprehensive survey of this field, and how it is being taken up in industry, commerce and elsewhere, and the *Knowledge Engineering Review* is perhaps the ideal sponsor of such a survey.

Machine learning research seems to have grown considerably in prominence, though even in 1984 machine learning and related topics were attracting great interest; so this might mean that we just did not cover this topic well enough. Recent *Review* content seems to have a particularly strong body of work in case-based reasoning and case-based learning, which was only one of several 'big' topics in machine learning in the 1980s. The growth of interest is striking, even allowing for the fact that a special issue on this subject included 17 papers!

Turning to the *Applications* theme, the commitment to developing practical results in the form of applications and application development tool has been fairly well sustained though there has been some drop-off. From 1985 to 1990, there were approximately 29 papers concerned with applications and application-oriented surveys, while in the period 2005–2010 there were about 20 papers. Bearing in mind that the sampling and interpretation processes were somewhat unsystematic this difference should not be overstated, however it is notable that this fall has occurred in spite of new applications areas emerging (deductive spreadsheets and knowledge management systems). A statistical analysis yields a  $\chi^2$  value of 36.571 and a probability of this change occurring on a chance basis  $<0.00001^2$ .

<sup>2</sup> One reviewer conjectures quite reasonably that this may be because 'applications are no longer worthy of publication unless there is some novelty in them'. Given the excitement of 25 years ago, I am surprised by the implication that knowledge engineers have run out of ideas for new kinds of applications and novel techniques. My suspicion is that a new generation of knowledge engineering researchers simply has different interests. This is a pity because building challenging applications creates great opportunities to test current theoretical ideas and force the discovery of new techniques by confronting the complexities of the real world.

This suggests a real reduction in the amount of work on applications, which were at the heart of knowledge engineering research 25 years ago. It is also interesting, perhaps worrisome, that there were eight articles with a focus on the commercial aspects of expert systems and knowledge-based systems in the first survey period, but there have been none published in the recent period.

	1985–1990	2005–2010
<i>Themes</i>		
Symbolic knowledge representations (e.g. rules, logic and related)	8	7
Cognitive tasks and capabilities (e.g. decision making, planning, design)	4	4
Machine learning, knowledge discovery, data mining	2	12 (CBR 17)
Autonomous systems (e.g. agents, cognitive robotics)	1	7
<i>Distributed AI and multi-agent systems</i>	1	26
<i>Applications</i>		
Engineering (general)	7	2
Software and hardware (e.g. programming, system configuration)	1	5
Biomedicine (e.g. diagnosis treatment, biomedical informatics)	2	3
Legal and regulatory applications	1	1
Mathematics and statistics	3	0
Natural language and text processing	2	1
Commercial applications	8	0
Development tools	5	0
<i>Deductive spreadsheets</i>	0	5
<i>Knowledge management</i>	0	3
<i>Foundations</i>		
Architectures for knowledge-based systems	3	2
Methodology, knowledge acquisition/elicitation	5	2
Explanation and user interface	3	4
Control, reflection and metalogic	4	0
Uncertainty management, logical, probabilistic and other methods	5	4
Cognitive science foundations (psychology, AI)	5	3
Formal and mathematical foundations	3	1
<i>Semantics, ontologies, semantic web</i>	0	10

The final theme of *foundations* is concerned with theoretical questions and techniques for developing, using and evaluating knowledge-based systems. This also seems to indicate a change in the interests of contributors to the *Knowledge Engineering Review*, since the 28 foundational papers published in the first period falls to 16 in the second. However, the numbers involved are relatively small so despite the apparently substantial falloff, it is not a statistically significant change ( $\chi^2 = 8.8, p = 0.185$ ).

One substantial change reflects the phenomenal growth of the World Wide Web, and perhaps a recognized need for firm technical foundations for the *semantic web*, such as ontological modelling and associated formalisms like description logics. Even in the first review period, some leading research teams were building large application knowledge bases (notably at Stanford and Carnegie-Mellon Universities where the expert and knowledge-based systems paradigms first emerged). They were already finding significant benefits in organizing these knowledge bases ontologically (e.g. MOLGEN, EURISKO). However, the ‘O word’ and terms like ‘ontology engineering’ did not become mainstream until well after the first review period.

### 3 Notable absences

In 1984, when knowledge engineering was thought to be ‘the political wing of Artificial Intelligence’, most of us expected that most of the core subjects of theoretical AI would be translated into practical technologies that would depend heavily on formalized knowledge and knowledge engineering. Prominent among these would be speech and natural language systems, and computer

vision systems. This has not in fact occurred. The reason for this is probably that these domains are associated with very large volumes of real time data which are best handled with conventional numerical computing and to which knowledge-based techniques have, so far, added little value.

The challenges of moving from signal processing to symbolic representation and processing have proved to be notoriously and unexpectedly difficult. We have had pretty good commercial word recognition and speech transcription products for quite a few years now, but to my knowledge no commercial language systems can claim to understand the *meaning* of what is being said or make use of domain knowledge in order to reason, recommend decisions and explain conclusions within natural language dialogues. In fact, there are precious few experimental systems with these capabilities even though it seems likely that the use of domain knowledge can help to improve understanding, and knowledge of conversational norms and pragmatics would improve naturalness of human interactions with expert systems, agents, avatars and the rest<sup>3</sup>.

The position with computer vision is similar: the raw power of modern hardware combined with a large range of computational techniques and algorithms has led to the emergence of wonderful image processing devices and systems, typified for me by the stunningly impressive capabilities of the modern digital camera. However, although these cameras can configure themselves in real time and real space, they still do not know much about what they are looking at. In my own field of medical informatics, in which it seems rather obvious why we might want image processing systems to be able to apply background medical knowledge to improve diagnosis and other decisions<sup>4</sup> we just do not know how to combine symbolically encoded knowledge with image processing in a general way. Many have tried to combine knowledge engineering techniques with image processing, but generally, the methods have been *ad hoc* and the results disappointing. A quick google for 'knowledge-based interpretation of medical images' suggests that after a flurry of activity in the 1980s and 1990s, most of this work had petered out by the turn of the century. The *signals to symbols* problem is scientifically, technically and commercially important but it remains unsolved and is a key challenge for new generations of researchers in knowledge engineering and allied fields.

#### 4 Development of the field

Looking back over the 25 years, there have been obvious advances and changes in the field shown by recent content of the KER, such as the substantial growth of interest in agents and ontologies. However, these are also topics that are of great interest in mainstream computer science; I am on the lookout for developments that distinguish knowledge engineering from mainstream computing rather than simply apply it in knowledge-based fields.

One distinctive feature of the knowledge systems paradigm that has been sustained is the commitment to symbolic representation of knowledge. In my paper in the first issue of the *Review*, I attempted to summarize a number of manifestations of this paradigm in a systematic way. This started from the idea that symbolic representations of concepts, relationships between objects, etc. differ from the kinds of representation which were the norm for computing in 1984 because knowledge representation builds on explicit *ontological commitments* which are available to the system when it is carrying out whatever task the system is designed for. Ontologies can capture the meaning of concepts in a humanly natural way (a concept such as a vehicle, a disease or a piece of equipment is located in a class hierarchy from which it inherits properties). Concepts can be composed into *descriptions* of complex objects and situations (e.g. Concept1 Relation Concept2), and descriptions can be further composed into *rules*, which can be used directly for reasoning

<sup>3</sup> As an example, I would draw the reader's attention to work by Martin Beveridge who developed a medical expert system with a spoken language user interface, which addresses these requirements. More information about this promising technology, including demonstrations and links, can be found at [http://www.openclinical.org/dm\\_homey.html](http://www.openclinical.org/dm_homey.html).

<sup>4</sup> Clinical decision support systems are a very hot topic in the medical world (see [www.openclinical.org](http://www.openclinical.org)).

(e.g. SituationDescription1 *and* SituationDescription2 *implies* SituationDescription3). When a symbolic view of domain knowledge is combined with the expressiveness of first-order logic, as advocated by Bob Kowalski in the first issue, we acquire some very powerful and versatile tools for practical problem solving. Countless groups in Europe and Japan, though less in the United States, became strongly committed to the importance of declarative representation of knowledge and the techniques of logic programming under the influence of Kowalski and others.

My attempt to informally map knowledge engineering also sought to accommodate higher-order knowledge structures, such as ‘tasks’, which attracted great interest in the early years. Alongside the development of symbolic representation and logical reasoning methods was a second key discovery; despite the enormous differences between application domains like engineering, medicine, law and business, there were also some common ‘tasks’ which shared many abstract features and to which general solutions could be applied. Clancey (1985) was among the first to exploit this idea in his model of *heuristic classification*<sup>5</sup>. Chandrasekaran (1986) and his colleagues had an influential research programme which suggested that there are many ‘generic tasks’ which can be viewed as standard components for assembling expert systems, including components for classification, hypothesis evaluation, planning, scheduling and design. The generic task was a key concept in the *KADS: Methodology for Building Knowledge-Based Systems Development* (Schreiber *et al.*, 1993).

Despite the potential importance of the idea that knowledge-based systems can be built out of standard components, which has been applied successfully in so many other engineering disciplines, it does not seem to be much discussed these days. A cursory search on the web suggests that it may simply not have delivered on its promise, though I suspect that it has just dropped out of fashion rather than failed. If it is the case that this classic engineering strategy is no longer being pursued it would be helpful to know why, and if it is not then it would be valuable to have an up-to-date survey of work in the area. In either case, the *Review* is the natural home for such a discussion.

My informal overview of how knowledge is structured has been superseded by more formal ways of thinking, though it had the virtue of being simple to understand and covered a surprising amount of ground (useful things in a manifesto that was intended to appeal to publishers as well as colleagues). I also found that it had practical value in *mapping* the field, helping me as the editor to see what had been covered and what had not, and how ancillary research topics such as reasoning under uncertainty, learning, vision and so on fitted into the ontological ‘backbone’ of the subject. I think that a new and updated map would be extremely valuable for understanding the current terrain and identifying promising directions and opportunities.

My broad conclusion from this short discussion of knowledge engineering is that the *Knowledge Engineering Review* is a unique publication, which maintains the role that was the objective of its founders as the ‘*primary secondary source*’ of information about trends in knowledge engineering and expert systems. The KER’s original goals of disseminating an understanding of the technical themes and scientific foundations of the field, in order to encourage the adoption of its methods, are also as valid as ever. After a quarter of a century of regular publication it has the position, perhaps even the authority, to articulate a new consensus about what the field is today and how it differs from mainstream computer science and cognate fields like applied AI.

When the idea of starting the *Review* was first mooted, its champions saw it as a contribution to establishing a professional organization of engineers, concerned with quality, safety, standards, accreditation and all the other things such organizations concern themselves with. I was to find out that this is not an easy challenge as in the end the editors are wholly dependent on what their authors want to write! Somewhere during my tenure as editor, I came across the following little poem. I hope it reassures members of the present editorial board as they approach their next publications deadline, and amuses the readers of this 25 years special issue as much as it did me.

<sup>5</sup> The essential idea here is a *situation* of interest (such as a fault in a device or a medical symptom that needs to be diagnosed) can be classified by linking it to a node in a suitable domain ontology, and then this node can be mapped to a node in another ontology of possible *causes* (e.g. failures, diseases).

## 5 The Editor's passport

The Editor stood 'fore the Heavenly Gate,  
His features pinched and cold,  
He bowed before the Man of Fate,  
Seeking admission to the fold  
'What have you done,' St. Peter asked  
'To gain admission here?'  
'I was the Journal's editor, Sir,  
For many a weary year,'  
The Pearly gates swung open wide  
As Peter pressed the bell,  
'Come in and choose your harp,' he cried;  
'You've had your share of hell!'

Anonymous:

Published in an editorial in *The Journal of the Irish Medical Association*,  
Vol. 42, No 247, Jan 1958, pp. 31–32.

## References

- Chandrasekaran, B. 1986. Generic tasks in knowledge-based reasoning: high-level building blocks for expert systems design. *IEEE Expert* **1**(3), 23–30.
- Clancey, W. J. 1985. Heuristic classification. *Artificial Intelligence* **27**(3), 289–350.
- Fox, J. 1984. A short account of knowledge engineering. *The Knowledge Engineering Review* **1**(1), 4–14.
- Johnson, T. 1984. The commercial applications of expert systems technology. *The Knowledge Engineering Review* **1**(1), 15–25.
- Kowalski, R. 1984. Logic programming in the fifth generation. *The Knowledge Engineering Review* **1**(1), 26–38.
- Schreiber, G. Wielinga, B. & Breuker, J. (eds). 1993. *KADS: A Principled Approach to Knowledge-based System Development*. Elsevier.