

Cognitive science and knowledge engineering

Many commentators talk about the aims of cognitive science and the achievements of knowledge engineering in the same breath. Even ignoring premature, even silly, claims that expert systems are “intelligent” or have capabilities which are comparable with those of human experts, the link between the scientific study of intelligence and the pursuit of new technological capabilities has been emphasized far too strongly. It gives the impression of a coherent and integrated community with common problems and interests. Nothing could be further from the truth.

Cognitive science is normally defined as the intersection of a number of disciplines, including artificial intelligence (AI), cognitive psychology, linguistics and philosophy. As figure 1 illustrates, a number of technological disciplines which are often loosely referred to as “AI” are more naturally viewed as parts of a more general field of symbolic information processing. These technological disciplines share with cognitive science an interest in symbolic computation, but do not usually have a strong interest in the phenomena of natural intelligence.

It is true that cognitive scientists are frequently interested in technological developments in symbolic information processing. These offer tools, techniques and sometimes even clarifying concepts for their work, but the development of smarter computer based tools is not what their own research is aimed at. This is of course understanding the phenomena and underlying mechanisms of natural human intelligence and, for some, an interest in how those phenomena might be expressed by machines.

It is also true that knowledge engineering has historical roots in AI, and that the techniques of symbolic information processing were primarily developed as part of the technical equipment needed for the computational approach to cognitive theory. However, the field has changed dramatically. Over the last few years many people have joined the ranks of so-called AI, including machine architects, programming language designers, and innumerable traditionally trained programmers working on applications. Relatively few of them, however, have any serious interest in the scientific wing of AI.

In fact the trend in the five years since Japan’s Fifth Generation report has been unmistakable. In topic after topic of traditional AI—automated reasoning, decision making and problem solving, planning, language and vision—the research trend has been towards a greater emphasis on the formalization which is so strongly prized in traditional computing, and away from the empirical, what we used to call “scruffy”, end of AI.




Figure 1 AI viewed as part of the scientific study of intelligence, but distinct from technological disciplines often referred to as AI which are better viewed as disciplines engaged in engineering advanced symbol manipulating systems

Why is empirical cognitive science important to knowledge engineering?

Although the aims of the scientific and technological communities are different, the apparent loss of interest by the latter in basic cognitive science is worrying. Let us recall two important examples of how basic cognitive science research has moulded applied AI.

Among the handful of pioneers of AI were Allen Newell and Herbert Simon who began work on human and machine problem solving in the 1950s taking a fundamentally computational approach. Their work reached a high point with the publication of their book "Human problem solving" in 1972 [1]. At about the same time Newell designed a "production rule interpreter" called PSG (Production System version G) specifically for use in simulating human problem solving. By some accounts their work on rule based systems at Carnegie-Mellon directly influenced the adoption of rule-based technology at Stanford, and hence the architecture of the first practical expert systems. Whether or not this is true it is certain that PSG led directly on to the influential OPS series of rule based cognitive modelling systems which, in later versions, were adapted for the development of expert systems. OPS5 was used by John McDermott and his colleagues to develop R1, the rule based system for configuring DEC's VAX computers which was probably the most significant commercial expert system ever built, since it appears to have been the first commercially successful one.

Another example of the importance of cognitive science is in the area of human and computer vision. The ability to see the world is a major capability, and one of the central topics of psychology and neurophysiology during this century. AI had attempted various formal approaches to computer vision during the sixties (the perceptron of Minsky and Papert at MIT; pattern recognition and syntactical approaches to vision at MIT and Universities of Maryland, Edinburgh, Sussex, and elsewhere) but after a promising start appeared stuck by 1973 or so. The computers at that time were so limited that scientists were being distracted by unprincipled "tricks" in their attempts to overcome speed problems. The field was unjammed by a scientist who was primarily interested in developing theoretical models of brain function, the mathematician and neurophysiologist David Marr. He asked a new kind of question, about what the brain was designed to "compute". Under the influence of experimental work in visual neurophysiology and perceptual psychology he defined a new direction for AI vision in a paper in 1975 (see [2]). Marr's work profoundly affected visual science, and the specialized hardware and image understanding programs which are now set to appear still show the clear contribution of Marr's reformulation of the visual computation problem.

Implications for practical knowledge engineering

Why did these enquiries, initially purely scientific, lead on to technical innovation with such obvious practical implications?

An observation that has been made many times but is worth repeating is that higher animals represent the only systems we know of that embody intelligence. They are the systems that can solve problems, see, make decisions etc. Anyone interested in achieving similar capacities in machines would do well to pay attention to the ways in which brains and minds meet evolutionary requirements.

This is not, as Slatter points out in his paper in this issue (see also [3]), a reason to try to mimic natural intelligence in detail or uncritically, but it is a reason to seek inspiration from it, as Newell and Simon and their colleagues did. After all cognitive science is in large part an experimental science. Disciplines which exploit practical experiments as well as theoretical armchairs have two features which are important for engineering innovation; experiments keep experimenters honest, and they frequently produce surprises which challenge received theories and force change.

While the growth of formalization in AI and knowledge engineering is an important trend, perhaps signifying a new maturity in the field, formalization is not by itself enough. If we ignore natural intelligence our very use of terms like knowledge and intelligence are likely to be fraudulent. The whole agenda of AI derives from common observation of the linguistic, perceptual, problem solving and learning capabilities of human beings. Computational studies of those capabilities in cognitive science should surely be part of the training and regular reading of knowledge engineers.

The Knowledge Engineering Review welcomes contributions from those wishing to develop the formal foundations of knowledge engineering and applied AI, but at the same time is concerned to promote the contribution of cognitive science. We shall always try to find space in its pages for presentation and discussion of the findings of cognitive science as well as for surveys and analysis of engineering topics.

References

- [1] Newell, A. and Simon, H. *Human Problem Solving*, Engelwood Cliffs, New Jersey: Prentice Hall, 1972.
- [2] Marr, D. *Vision*, San Francisco: Freeman, 1982.
- [3] Slatter, P. *Building expert systems: cognitive emulation*, Chichester: Ellis Horwood, 1987.