The Knowledge Engineeering Review, No. 4: 3, 1989, 249–253

Book reviews

The role of information scientists in the development of expert systems: interim report reviewed by R J Harley, Department of Information and Library Studies, University College of Wales, Aberystwyth, Dyfed SY23 3AS, Wales.
On being a machine, Vol. 1: Formal aspects of artificial intelligence reviewed by Dr Gerhard Brewka and Dr Joachim Hertzberg, GMD (German National Research Center for Mathematics & Data Processing), St Augustin 1, Federal Republic of Germany.
Progress in machine learning—Proceedings of EWSL 87: 2nd European working session on learning reviewed by Dr Lee McCluskey, Department of Computer Science, The City University, London EC1V 0HB, UK.
Professional judgement: A reader in clinical decision making reviewed by Dr Phil Slatter, Telecomputing plc, Oxford, UK.

The role of information scientists in the development of expert systems: interim report by Margaret O'Neill and Anne Morris 1989, British Library Publications Sales Unit, 1989, pp 103, £5.00.

Library and information scientists have been interested for some time in the potential application of expert systems in their work. Research projects have examined the use of expert systems in the application of cataloguing rules, as intelligent front-ends to information retrieval systems and as referral sources amongst others. One project has resulted in the subsequent development of a commercial product. It was thus logical for information scientists to look beyond their own domain and examine their potential role in the development of expert systems in other domains. The research reported in this report, funded by the British Library and executed at Loughborough University of Technology, has sought to examine this potential as a contribution to research within the library and information science community about the exploitation of their skills within a new emerging information marketplace. The bulk of the document reports a survey of the actual use of expert systems within industry, commerce and academia. This survey reveals the areas in which expert systems are being developed—largely engineering fault diagnosis and configuration tasks and financial and business decision support. It notes that shells are by far the most popular development environment and that remarkably few expert systems appear to be in daily use. The survey also performs a useful function in revealing a diversity of opinion and practice in expert system development techniques. In the light of this information, suggestions are made about the potential contribution of library and information specialists to expert systems development. These suggestions reflect a mixture of the opinions of the authors and the knowledge engineers who were interviewed as a part of the project.

The report has two potential audiences namely information scientists and knowledge engineers. Information scientists should be interested in the findings of the project about the potential application of their skills as part of expert systems development teams and the views of knowledge engineers about how information scientists should develop their skills and credibility in expert systems development.

Knowledge engineers should find the report of interest because it provides a snapshot of the actual state of expert systems development and application within the United Kingdom which cuts through all the hype surrounding the subject. Furthermore it provides useful information about the methods actually used in the development of expert systems rather than academic accounts of the methods which should be used. This section of the report reveals a diversity of methods and opinions amongst knowledge engineers which reflects the infancy of the subject and, if encountered by knowledge engineers in experts consulted as a part of their work, would cause them considerable

problems. Knowledge engineers may also be directed to long established methods of knowledge representation used by information scientists which may be of value in their own work. The modest investment of £5.00 may repay both audiences with some useful insights.

On being a machine. Vol. 1: Formal aspects of artificial intelligence by A. Narayanan 1988, Ellis Horwood/John Wiley, pp 200, £25.00.

Why write and publish a book? We can imagine several reasons for doing so, let aside the more political ones: an author may wish to discuss some new ideas of his or her own; one may wish to discuss old ideas in a new, more coherent or more popular and yet correct way; in particular, when teaching a course one may feel the need to have some material covering its very content.

The purpose of Narayanan's book is to explore the thesis that AI "has not developed significantly as far as its theoretical and philosophical aspects are concerned, because from the very beginning AI has been miscategorized by theoreticians, philosophers, and even AI researchers" (p. 9). The author has split his exposition into two volumes, the first of which (reviewed here) contains some formal and philosophical background. In particular, Turing's Imitation Game is discussed at length (Ch. 1); and there are introductions to concepts of theoretical computer science such as computability, automata theory and formal languages (Ch. 2) and logic (Chs. 3 and 4). Besides being presented formally, these theoretical concepts and their implications are related to the more philosophical questions raised by Turing's paper. The book aims at "philosophers (who) may want an introduction to formal theory" and experts "to remind expert formal theoreticians ... that at the roots of their subjects are very many definitions and assumptions which they may have forgotten about ..." (p. 10).

Unfortunately, the introductory presentation is very often imprecise, misleading or even incorrect (not to mention the incredible number of misprints throughout the book). The beginner in one of the fields covered will, at best, be confused. Take as a deterring example Chapter 3, the introduction to predicate calculus. You will find there, without any explanation, two different definitions of the syntax of predicate calculus; a third definition, with an absolutely unnecessary and unexplained shift in terminology, is given in Chapter 4. On the other hand, the author never presents a complete list of inference rules, rather surprising for a more than 50-page introduction to predicate calculus. Formal systems are defined as usual but then mixed up with logical calculi having a model-theoretic semantics. Properties of specific logical systems, e.g. that a formula is derivable if and only if its negation is unsatisfiable, are presented as general properties of formal systems. Even such basic notions as satisfiability and entailment are mixed up at various places (pp. 106f, p. 110).

And how about the following statement:

> ... Gödel showed that it is not possible for a formal system (of a sufficiently powerful kind) to be both complete and consistent *at the same time*. That does not mean that a formal system is not both complete and consistent; only that if it is required that completeness and consistency be demonstrated simultaneously, this cannot be done. (p. 143)

We have no idea what the author wants to say with these remarks, and what they have to do with Gödel's theorems.

Just as some of the remarks in Narayanan's book are amazing so is what is *not* in it. What does he take AI to be? It is just classical symbolic AI or are subsymbolic issues included? There is no reference to the Physical Symbol Systems Hypothesis! Searle's Chinese Room argument is presented, but there is no reference to the lively discussion it stimulated. The author takes philosophical views (which might simply be described as "mechanism is wrong") but the arguments in favour of them are never discussed in detail; you may either believe it or not. As an example, take the remark "that ascription of mental predicates to computers is logically impossible, simply because the ascription of mental predicates to an entity implies, or presupposes, that the entity in question has consciousness" (p. 191). Is this a dogma or an axiom? We thought this should be one of the questions to be *discussed* in a book *On Being a Machine*!