

Book reviews

Bryant N 1988, John Wiley and Sons, Chichester. Pp. 214

MANAGING EXPERT SYSTEMS

Reviewed by John Barber, Unilever Research Port Sunlight

The market place badly needs books which address the topic of how expert systems will impinge on management practice, how they should be integrated into the work environment, and what management issues they directly raise. For those who agree with this statement, do not be misled by the title of this book.

Managing Expert Systems is, in reality, an introduction to expert systems for managers. As such it is a legitimate enterprise but perhaps this rather less sexy (but more accurate) title, would have been judged a less satisfactory precursor to impressive sales figures. Having said this, at £11.95 the book is sensibly priced for a likely audience of casual readers.

The terms of reference are clearly stated in the preface. The author claims to have written an easy-to-read book (yes), with chapters that can be read independently (more or less), and at a level that will not intimidate the general manager and students of management (certainly). He also warns that a lot of computing publicity tends to appear in the form of promises, and that it should be remembered that expert systems are the current vogue in a long standing tradition of industry "hype". This statement sets the general tone of the book, which on the whole strikes a good balance between strong publicity and words of caution. As such it avoids the major pitfall associated with this sort of text, namely that of damaging the greatest potential market for expert systems by creating overexpectations of what is available to managers from today's software. Perhaps the only blemish on this aspect of the volume is the author's decision to place expert systems in the context of AI right from page one. There is much to be said, given the projected readership for this book, for keeping the heavy AI gun in reserve rather than pointing it full in the reader's face at a stage when he can still be frightened off.

The topics dealt with are ordered in a sensible progression, beginning with an attempted rationale of expert systems, through to structural descriptions of software, criteria for choosing suitable applications, development stages, updating and maintenance. There then follows a section on choosing a shell. The author rightly concentrates on pc shells for the audience he addresses, but having made this decision devoted 50% of his space to discussion of languages. There is some onus to say why a comparative discussion of languages is important, but somehow this does not properly emerge. This reader certainly felt that there was no clear explication of the language/shell controversy in this book, and that the uninitiated would not be enlightened on the issue.

In general however, the book's structure is its best feature. Most things that one feels ought to be addressed are there, and are explained in simple terms.

Unfortunately the actual content of the book is not so satisfactory. Economies of truth in a superficial discussion of complex subjects are to some extent forgivable, particularly if clarity is attained at the expense of verbosity. Hence one might overlook the implication that the terms relation and tuple are interchangeable; that microSYNICS is just a means of writing front-ends to BASIC programs; that KEE has PROLOG-style programming, and one or two other minor distortions of reality. However to describe deep knowledge as "knowledge about knowledge" does not fall into the category of forgivable misdemeanour, neither does the claim that problems requiring "commonsense" are the ones best handled by expert systems. Sadly the reader will meet many such passages in this book which will grate on his better judgement, and hint at the absence of deep (in the more usual sense) knowledge.

However the most unusual feature of this publication is the inclusion of lists of software descriptions together with details of prices. The excuse for including prices is that they are only there

to facilitate, comparative judgements. Nevertheless they will become rapidly out-of-date, and the decision to state absolute figures rather than provide some category scale as an aid to relative judgement was surely a mistake. Indeed some of the descriptive details provided are already obsolete. For instance KEE is stated to be available only on Xerox machines. Far from it: not only has it been on Symbolics and Texas machines for a while, but it will shortly, we hear, not be supported on Xerox at all.

This book left me with mixed feelings. Anything that guides managers toward beneficial usage of expert system software helps the industry advance further into what is potentially its most valuable market. When addressing this market we must be sure that our voice not only has the quality of clarity, but is pitched at a level to inspire confidence, and has behind it the authority necessary for delivering state-of-the-art distillations. This book succeeds on the first count, but fails thereafter. Given its size, structure and aims, it might have done much better as a series of articles in one of the popular management journals. In this way the author's stated aim of increasing awareness amongst managers might have been more effectively met.

ALL THE RIGHT MOVES – A VLSI ARCHITECTURE FOR CHESS

Carl Ebeling, 1987. The MIT Press, Massachusetts. Pp. 145

Reviewed by Don Beal, Queen Mary College, London University

This book is a publication of Carl Ebeling's Doctoral Thesis which merited selection as an ACM distinguished dissertation. It is remarkably well written and easy to read. Although the title includes "VLSI", hardware design details only occur in parts of the text, and the book is accessible to a wide audience.

The first two chapters survey some of the performance achieved and techniques used by chess-playing programs up to 1985. A crucial operation is that of generating the set of legal moves for a given board position. Chapter 3 discusses the details of doing this by a new approach used in the chess machine "Hitech" built by Ebeling and others at Carnegie Mellon University. The new approach uses parallelism in a different way from hitherto, replicating circuitry on a per-square basis, but in addition with each square replicating some knowledge of the global board state on a relevant-to-this-square basis. Much of the circuitry is devoted to achieving a desirable ordering of the moves.

Chapter 4 discusses position evaluation and describes the recently introduced concept of an "oracle" and a hierarchy of first-order, second-order and higher-order evaluations. The oracle is a collection of chess knowledge separated from the search mechanisms, and which serves to choose the coefficients for the evaluation function during a particular search. The other side of the coin is that the evaluation terms are deliberately simplified down to what can, with affordable hardware, be computed in a single machine cycle. This turns out to be a good trade despite the possibility that the search will encounter deep positions in which the coefficients chosen as appropriate at the root will be inappropriate. This occurs, but is sufficiently infrequent, or sufficiently minor, to be acceptable. The oracle knowledge can be as complex as desired, because it is only executed once per move computation, whereas the position evaluation is executed millions of times. This splitting of knowledge into two extreme forms has been successful enough to make Hitech currently the best chess-playing machine, having achieved a chess rating approaching 2400 against human opposition in regular human tournaments.

Chapter 5 contains hardware details of the rest of the Hitech machine. Chapter 6 contains a survey of the performance results and testing of Hitech, showing how the various hardware components, move ordering, and the hash table contribute to the overall performance. It is