engineers, computer programmers, engineers, business professionals, psychologists, and others who need to keep up with the field of AI". The dictionary is divided into four sections.

The first and second sections are the dictionary itself, providing detailed entries for a large number of expressions. The author claims that there are more AI definitions than any other dictionary currently available, and this seems believable. The definitions vary from a mere cross-reference to another expression, through a short paragraph, to a whole page, and are often divided into sub-entries for the various meanings of a word. They tend to be clearly written, though the style is sometimes a little telegraphic. Cross-references to related terms are abundant, and a particularly welcome feature (and one that is sadly missing from some other such dictionaries) is that they contain pointers to items in the bibliography for the reader who wishes to find out more about the subject.

The third section contains detailed references to various expert system shells available in the market, and is perhaps the most "perishable" part of the book. It contains a brief description of the most popular products available, together with information about vendors, and sometimes prices.

The final section is an extensive and up-to-date bibliography, with some 400 entries, that should prove very useful in helping the reader to learn more about specific entries.

In any book of this nature, which tries to cover the very broad range of subjects which make up AI, any reviewer is bound to find that the covering of their own branch is somewhat sparse, and this one is no exception. My feeling is that the dictionary is very heavily slanted towards expert systems, a reflection of the author's background, at the expense of the other subjects. For instance, in my opinion the coverage for natural language processing is not only very short, but more importantly it is seriously out of date, ignoring many of the advances of the last decade. A solution to this problem would have been perhaps to share the work between people from various backgrounds, or to have them on the editorial board.

In general, the entries are well written, but the editorial work occasionally leaves something to be desired. In addition to the many annoying typos and occasional mistakes, there are a few editorial quirks and inconsistencies, like cross-references to non-existing entries, or entries for expressions which are supposed to be synonymous (such as those for *Language Processing* and *Natural-Language Processing*), but which seem to be so unrelated that they could have been written by different people. Many of these problems should have been spotted by a careful proof-reader. No doubt they are a consequence of the tight deadlines imposed by a rapidly changing field, and will be put right in future editions of the dictionary.

In short, I would say this dictionary is a reference tool well worth having, but its heavy slant towards expert systems probably means that you will need another more general dictionary alongside it. Still, it is as complete as any of the other dictionaries that come to mind, and should be a welcome addition to any reference shelf.

Knowledge representation: an approach to artificial intelligence by TJM Bench-Capon, Academic Press, London, 1990, pp 220, £19.50.

Reviewed by: John Washbrook, Department of Computer Science, University College, London, UK.

When I first saw the contents pages of this book, including chapters on Logic, Search, Production Rules, Nets and Frames, Prolog and Expert Systems, I expected a substantial tome. When it arrived I was surprised to find that I was right–the book is substantial–but not in the way I had anticipated. At 200 or so pages it is a fairly slim volume, which it is quite possible to read thoroughly in a couple of evenings (depending on one's knowledge of the area). But it is substantial in its intellectual content. Topics are introduced assuming little prior background and examples are given, so that the text stands on its own, but it is not really an introductory text. The introductions of the topics serve more to establish the notation and concepts. It lacks the detail of, say, Jackson or Rich, and instead gives a meta-level view of the area. The impression it gives is of a text which has been produced as a summary of a course—well, not a summary, but a discussion of some of the

important concepts and results which it is possible to appreciate at the end of an introductory course. This makes it valuable both to the student revising for an examination, and the practitioner who is familiar with logic, production rules, nets, Prolog, etc., and who needs freshening up on the power of and relationship between the different knowledge representation formalisms and inferencing mechanisms. Topics such as soundness and completeness, the problems of inheritance and defaults, negation as failure and the closed world assumption, non-monotonic reasoning and truth-maintenance, the limitations of Prolog, all of which have subtleties which often seem to be elusive, are dealt with briefly but with clarity and authority. There are pointers to the more esoteric aspects of the subject. I can recall a few occasions on which it would have been a useful book to press into someone's hand.

So much for the paeans of praise! The author discusses the various representation formalisms from a neutral point of view, but a predisposition for logic would seem apparent. In the discussion of semantic networks, the "is-a" relation is introduced, but not the "ako" (a kind of) relation. There is belatedly a discussion of the problems of the confusion between the class-instance relation and the class-subclass relation, but given the concise nature of the text it would seem to have been better to avoid this by using ako where appropriate and then, possibly, have a discussion of the problems which arose in the early days from using solely the is-a link. This is my only major criticism of the book. It is a pity that the figures of nets have arcs which are not directed, and there is at least one arc which is misnamed. In general, though, the examples in the book suffer little from such typographical errors. The table of contents is good (I found it more useful than the index), and there is a useful list of references at the end of each chapter.

I am pleased to recommend this book as one which gives a fair discussion of logic compared to other formalisms for knowledge representation. I like it for its breadth, its clarity, its conciseness, and for its readability: for me the text has that quality, rare in a technical book, of communication—the feeling that the author really wants to share his understanding.

References

Jackson, P, 1990 *Introduction to Expert Systems*, (2nd edn) London: Addison-Wesley.
Rich, E and Knight, K, 1991 *Artificial Intelligence*, (2nd edn), New York, NY: McGraw-Hill.

An introduction to neural computing by Igor Aleksander and Helen Morton, Chapman and Hall, London, 1990, pp 255, £15.95.

Reviewed by: S Parfitt, Research Fellow, IBM UK Scientific Centre, Winchester, UK.

AI specialists interested in getting to grips with the basics and background of neural computing will find this introductory text very useful. It covers the main network types while at the same time charting the history of the subject in such a way as to give the reader a clear understanding of how it has developed and reached its present state.

The introduction gives an exposition of some of the reasons for "the fuss" about artificial neural networks (ANN), contrasting conventional computing (algorithms) with neural computing (learning by experience), and also gives a good summary of the historical aspects which the rest of the book goes on to detail. Chapter 1 leaves nothing assumed, describing a simple network node and various definitions of network shape. The authors start as they mean to go on—rigorously—by talking of network firing patterns in terms of truth tables. This, plus a discussion of "Where is the knowledge stored?", should be enough to make the AI specialist feel at home.

Chapter 2 describes the early attempts at building small ANNs: the McCulloch and Pitts model and the Perceptron. These are looked at briefly from an engineering stance, and then the delta (learning) rule is introduced, along with an algebraic interpretation. As many readers will be aware, activity in ANNs received a heavy blow in 1969 in the form of Minsky and Papert's book on Perceptrons, in which they showed that these were not capable of "hard learning".