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".

Chapter 3 gives a summary of their arguments, which are then countered later in the book by a description of how Rumelhart, Hinton and Williams, in 1986, developed a network architecture—the multi-layer Perceptron—involving hidden nodes and a new learning algorithm (back propagation, a form of the delta rule), which succeeded in overcoming the limitations of the simple Perceptron. After some introductory material on various methods of building ANNs in hardware, there follows a description of WISARD, an adaptive pattern recognition machine in use in industrial settings since 1984, which exemplifies Professor Aleksander's preferred method of building hardware ANNs by employing RAMs. The ideas introduced at this point are expanded upon later in the book.

There are two chapters describing a class of neural networks with associative properties: the Hopfield net and the Boltzmann machine. An understanding of the latter provides insight into the benefits accruing from the addition of noise into certain networks: something that will come as a surprise to those new to ANNs and used to the idea that noise can only hinder!

Chapter 9 covers other classes of ANN such as Kohonen and competitive learning networks. In the final part of the book, the reader is given a taste of current ANN applications, such as speech, language and vision tasks, and of possible future applications in domestic, industrial, medical and financial settings.

The book as a whole is well laid out, with self-explanatory headings, plenty of helpful diagrams, and exercises at the end of each chapter. There is also a section of comments at the end of the book which go with the exercises. The text is almost always clear and concise, and there are frequent overviews of what's coming next to help you see where the authors are taking you. The bibliography covers all the established background texts, and the index is similarly thorough, except for the very occasional cases of items being positioned where you would not necessarily expect them. For example, a newcomer who has heard those mysterious words 'back propagation' will not find them there, and may not know to look under 'error propagation' instead.

The inclusion of carefully-formulated learning algorithms will arm the novice with the necessary basic information to begin building his or her own ANN. As a teaching aid, the book will be similarly invaluable, particularly given the availability of an accompanying software package, 'Cortex'. As well as simulating the network types described in the book, the package includes some demonstration networks which mirror examples in the book, and which can be both run and altered as a means of gaining hands-on experience. The manual is a little on the brief side, but the package itself is excellent in terms of ease of use and clarity. It does not provide an all-singing, all dancing development environment, but then its purpose is rather that of a learning tool. My early copies of the package and manual have one or two minor inconsistencies which, the distributors inform me, have now been ironed out.

Anyone expecting a hype of the 'magic and mystery' of ANNs will not find it here. It is therefore an ideal introductory text for the sceptic, as it deals rigorously with all the aspects of network architecture and learning upon which it touches. If you are looking for a one-step guide to neural computing, or a starting point for a serious study of the subject, this text and the accompanying software package are certainly to be recommended.

CORTEX: Neural Networks Demonstration System, written by Michael Reiss, is distributed in the UK by Unistat Ltd., PO Box 383, Highgate, London, N6 5UP. Price £100.

Expert knowledge and explanation, by Charlie Ellis (Ed.), Ellis Horwood, Chichester, 1989, pp 246, £39.95.

Reviewed by: Professor FH George, Bureau of Information Science, Seer Green, Bucks, UK.

The subtitle of the book is "The Knowledge Language Interface" and is made up of 12 papers by a variety of authors, including the book's editor. It is concerned with the communication of knowledge and of explanations in the context of intelligent knowledge-based systems (IKBS). Charlie Ellis himself has a paper in the collection entitled "Explanation in intelligent systems", and the critical issues are seen to be: