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:

"1. The differentiation of different kinds of knowledge.
2. The explicit rendition of strategic and structural knowledge that tends to get compiled into program code.
3. The modelling of individual user's knowledge or levels of skill."

All of this is applied to explanation in the context of expert systems. Chapters 6, 7, 8 and 9 are related to this first paper under consideration, since they too are concerned with the provision of explanation to their human users.

The first two papers in the book are both concerned with verbal reports and knowledge acquisition. P Pollard and R Crozier are the authors of the very first paper entitled "The validity of verbal reports: unconscious processes and biases in judgement." It is lengthy and detailed, and is concerned primarily with the collection of physiological articles on verbal reports. This leads to a consideration of how expert decision processes might be modelled.

Chapters 3 and 4 deal with expertise, and the first of these deals with the nature of human intelligence. They both take up the matter of machine intelligence, and refer especially to the famous work of Alan Turing; some thought is also given to the different features that make up human intelligence.

The last three papers (Chapters 10, 11 and 12) are primarily concerned with the acceptability of an IKBS to the user. Chapter 10 is called "Lay Explanations of Behaviour", and is concentrated on four main points. The intelligent system must care whether its explanation is to be believed, and this means that it must be sensitive to just what kind of "why" the questioner means: whether cause or warrant, what the user does and does not know, and what the user will accept as an explanation. Chapter 11 is concerned with the "The Expert Controller", and Chapter 12 is called "Expert systems acceptability: human and organisational context of expert systems."

The core of the book—Chapters 6 to 9 inclusive—is, as the book's title implies, concerned with various aspects of explanation. An especially interesting article entitled "A language to allow expert systems to have beliefs about their users" written by Anthony Maida amd Mingqui Deng from Penn State University, deals with "the problem of representing subtle aspects of another agent's belief states." A more lengthy statement by the authors helps to clarify their purpose:

"unfortunately, simply knowing the contents of another agent's database is insufficient to characterize that agent's belief state. We must also know how that agent interprets his representational symbols."

This all takes us close to the extremely interesting frontier between science and philisophy, and rather especially the meaning and semantic aspects of philosophy: an excellent contribution to the subject.

Another exciting paper comes directly after the one just discussed, by A. C. Zimmer from the Universitat Regensburg in Germany. He thinks of the Gricean model of pragmatics as a general standard for the design of human–computer dialogue. The concluding remarks start with the statement:

"It has been demonstrated how cognitive and linguistic theories of dialogue processing can help to derive criteria for a principled development of human-computer interaction with an eye to the needs and skills of the user."

All this leads on to consideration of computer language use, and is a very important part of the book's overall topic.

This is an outstanding book which makes a major contribution to the very excellent series of Ellis Horwood books on information science. It is very well produced and well written, and deals with a subject which is of central importance to information science. It is one aspect of what to the reviewer is a central theme of the subject, which is the science-philosophy relationship. It is, of course, from the computerized point of view, that the approach is made, and this sort of approach needs to be followed up by an approach which is couched more in philosophical, logical and semantic terms, all of which are essential to the future development of information science.