The second part of this chapter dealing with expert systems in practice was a bit of a letdown. There was too much emphasis placed on rule-based programming with scarcely a mention of the other AI programming paradigms. Also the examples used to illustrate various points were not taken from an engineering domain and the commercial products used to code up the examples are not particularly representative of what is around in the market place at the moment. In addition to this, sample code was presented in the body of the paper without adequate supporting explanation in the text.

Of the other introductory chapters I particularly liked the paper on blackboard architectures. The blackboard model and architecture was illustrated using a jigsaw model, with each component of the blackboard and its purpose being described using further examples. With all the chapters in part one of the book the reader was brought from a position of no knowledge of the topic to a position of understanding not only the basic concepts behind the topic, but also where the current research interest lies in a number of easy and interesting steps.

Part two concentrated on applications of artificial intelligence within engineering. Some of these applications were very interesting, and clearly illustrated the concepts described in part one. In many of the applications the AI techniques have been used to improve on existing conventional processing techniques. A blackboard architecture was used in one application to enhance traditional signal processing algorithms, producing a system capable of intelligent spectral estimation with relation to foetal heart signals, and in another application in relation to sonar interpretation. Foetal heart sounds are analysed in another paper too, this time using a combination of signal processing algorithms and rule-based reasoning to estimate the duration of various phases of the cardiac cycle.

The use of AI techniques in automatic speech recognition is another application area covered. The problem of speech recognition has been tackled using both a mainstream applied AI approach, known as the "strong knowledge" approach, and using a more mathematical approach, known as the "strong algorithms" approach. Here the emphasis is given to research into the use of neural nets and multi-layer perceptrons.

Machine learning features strongly in the remaining papers. One paper contrasts aspects of a diagnostic system built using both a production rule system and a machine learning system. The other paper considers machine learning in relation to tuning waveguide filters used for tailoring filters before and after transmission.

The applications are readable and interesting, probably as a result of the background ideas having been explained in the first part of the book. A criticism however is that the examples are drawn almost entirely from electrical engineering, and from the signal processing area within that discipline. The applications described illustrate the use of AI techniques to tackle a wide range of problem-types which could be relevant to a number of engineering disciplines. It is a shame that other engineering disciplines are not represented in any way.

I enjoyed reading the book and I learned a lot. I would recommend it to engineers having some knowledge of artificial intelligence who have reached the limitations of toolkits and shells. The techniques described in this book will help these engineers to understand and use more sophisticated techniques alongside other knowledge based or conventional programming approaches. Although I see it primarily aimed at academic engineers carrying out research projects requiring the use of artificial intelligence, this book has additional value in making the AI community aware of the level of activity going on in the engineering area and the continuing need for AI research to overcome the problems which are still outstanding.

Reviewed by Terri Lydiard, Knowledge Engineering Group, AI Applications Institute, University of Edinburgh.

Envisioning information by E. R. Tufte, Graphics Press, London, 1990, pp 126, £30.

Information comes in all shapes and sizes, and can be complex or simple, dynamic or static. It is therefore not surprising that problems arise when trying to summarize information in a clear,

concise and elegant manner. For example, there are very few systems that depict in a straightforward way the detail as well as the general structure of a knowledge base in a diagrammatic or pictorial form.

This book dexcribes ways of breaking out of the "flatland" that is paper to illustrate complex and multi-dimensional information in fascinating and compelling ways. Tufte's previous book, *The Visual Display of Quantitative Information*, sold nearly 100,000 copies, and has become a coffee table classic. Both books succeed in communicating serious design principles in a very readable and engaging way. Whilst the previous book concentrated on statistical graphics, this one encompasses all types of information, from poet population maps to the life cycle of the Japanese beetle. There are six main chapters in the book which are described below.

The Escaping Flatland chapter draws on many different sources, from books on colour coordination to academic treatises on urban pollution, and illustrates many innovative design solutions to the problem of representing very complex information in an understandable way. The Micro/Macro Readings chapter investigates ways of condensing and summarizing data in such a way as to allow meaning to be easily perceived. Examples here include Tukey's stem and leaf plots of statistical data, and a set of amazing diagrams illustrating the computed orbit of man-made space litter. Tufte regards confusion and clutter not as inherent attributes of information, but as failures of design. In the chapter on Layering and Separation, methods for sifting out important information from unimportant background and secondary data are discussed and illustrated. Especially interesting here are the redesigned maps, diagrams and notations showing how intuitively good design solutions may be achieved, if only they are properly thought through. The chapter on Small Multiples shows examples of data that are generally repeated each time, with only details changing between diagrams. For example, the multiple views of Saturn from Earth as both orbit the sun by Huygens, and the differences in brain scan activity by different types of mental illness and brain wave. Perhaps the most interesting and useful chapter for me was the one on colour and information, examples of cartographic and computer displayed colour information illustrate persuasively the effectiveness of colour in information displays. The final chapter, Narratives of Space and Time, shows many examples of how spatial and time series data can be shown clearly and effectively, again several examples are redrawn to illustrate how small changes can make vast improvements in comprehensibility and aesthetics.

In summary, this book is probably more interesting to a general audience than his previous work, but is perhaps its academic inferior. Envisioning information does remain readable and avoids the trap of being fanatically prescriptive (unlike his previous book), and only suggests solutions through illustrating good and bad methods of depicting data. Perhaps after reading through this book I can forgive the overbearing nature of his previous book, and finally join the Tufte fan club.

Reviewed by D. Van Laar.

Reasoning and revision in hybrid representation systems by Bernhard Nebel, Springer-Verlag, Berlin, 1990, pp 270, DM 42.

It must be clear to anyone who has worked with knowledge representation systems dealing with real problems that knowledge is not merely a static entity. Practitioners are constantly confronted with new discoveries about the task domain, and a changing world. Many are often confronted with the problems of correcting previous errors in their knowledge representation. Thus, the task domain knowledge frequently has to be revised. This is not a trivial task as it involves several problems. Some of the problems are of a pragmatic nature, others are concerned solely with semantics. There are also problems associated with the efficient management of the knowledge bases.

Supporting the revision of knowledge-based systems is a crucial topic for building interesting