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

large knowledge-based systems. Currently there are no tools available for this task, with the exception of primitive tools such as knowledge-base editors, e.g., KREME, etc.

This book is divided into two parts, and investigates two important aspects of practical knowledge representation systems:

- *reasoning* with represented knowledge; and
- *revising* represented knowledge.

The first half of the book deals with the dynamic aspects of knowledge representation. It contains a comprehensive survey of knowledge representation formalisms, the architecture of various services of knowledge based systems, and their management. The emphasis is however on *hybrid* systems, including semantic networks, frame systems, structural inheritance networks, etc. The book demonstrates a deep understanding of these systems and of the underlying theories. The analysis is further constrained to a typical member of the family of hybrid representation systems based on KL-ONE.

The second half of the book is devoted to the study of the problems associated with the revision of knowledge bases. It is often assumed, when discussing inferences, that knowledge bases grow monotonically. This excludes the possibility that a fact or definition has to be retracted or modified once entered. Here, instead of proposing a solution, the book takes a more general view on the problem of revision, which is then used as a yardstick for the effectiveness of previous solutions such as network editing approaches, knowledge base editing, deletion and additions of definitions systems, etc.

The target audience of the book falls into three groups: teachers/students, researchers, and other practitioners in the field. Overall the book is a popularized version of the author's PhD thesis. As such it contains a considerable amount of theory and formalism which is of limited value to the first and last groups. It is however valuable reading for prospective PhD students and most serious researchers in the field. And it is still a good PhD thesis.

Review by Dr. Paul Refenes.

KARDIO: a study in deep and qualitative knowledge for expert systems by Ivan Bratko, Igor Mozetic and Nada Lavrac, The MIT Press, Cambridge, MA, 1989, pp 260.

Disorders in the electrical control system of the heart are called cardiac arrhythmias. There are about 30 basic arrhythmias, each of which causes some characteristic change in the normal ECG. Simple arrhythmias combine to give complex combinations of multiple arrhythmias. Although the ECG features corresponding to simple arrhythmias are well documented in the medical literature, there is no systematic description of ECG features for multiple arrhythmias. This knowledge simply does not exist to render the application of a dialogue-based knowledge acquisition technique a possibility. If the expertise does not exist, the only viable knowledge acquisition technique is to "learn" the expertise from the domain first principles. This was the motivation behind the KARDIO project which started in 1982. The book by Bratko, Mozetic and Lavrac is an account of the research conducted in the period 1982–1987. The initial intention of KARDIO was ECG diagnosis, directly from the actual ECG curve; however, in the present system the ECGs are manually translated to a symbolic description. Prediction is also provided, since it is intrinsically related to the automatic derivation of the relevant expertise: "Given a particular legal combination of simple arrhythmias what are the possible ECG descriptions?"

The terms "deep" and "qualitative" included in the book title are currently very much in vogue within the knowledge engineering community. Through its title the book should therefore have an instant appeal to many knowledge engineers, who will benefit if they actually decide to invest in the book. From the title, however, one may not immediately guess that a central theme in KARDIO is automatic knowledge acquisition through rule induction based on techniques which have been established in the early 1980s. The success of the KARDIO project and the wide dissemination of