

cognitive science, some of the results reported here appear a little fragile, often being based on small subject samples and on only a very few experimental tasks.

However, my overall view is that this is a very worthwhile collection which deserves close reading by all students and researchers concerned with medical cognition in particular and expert cognition in general.

Reviewed by Dr Ken Gilhooly, Department of Psychology, University of Aberdeen, King's College, Aberdeen AB9 2UB.

Knowledge based management support systems edited by GI Doukidis, F Land and G Miller. Ellis Horwood Ltd, Chichester 1989, pp 356, £39.95.

The stated aim of Georgios Doukidis, Frank Land and Gordon Miller, the editors of this book, is to explain to business managers throughout the world of commerce and industry, key concepts of and relationships between knowledge-based systems and decision support systems. A knowledge-based or expert system can be loosely defined as a computer system which stores and uses problem solving expertise to support the solution of a well-defined, well-understood, recurrent business problem which would otherwise have required the intervention of a human expert. A decision support system, by contrast, is concerned with the facilitating of a decision-making process, which may not be repeated, often by modelling the principal factors influencing the decision and their interactions rather than by modelling the human approach to solution. The techniques of operational research are often used to underpin the construction of decision support systems: this is seldom the case with expert systems.

The book itself is a collection of 22 papers, mainly from a conference held in 1987 at the London Business School, though papers from other sources are also present. There are three sections of which the first covers Knowledge Based Management Support Systems (KBMSS)—the result of combining DSS and expert systems; the second, DSS; and the third, expert systems.

The section dealing with KBMSS has perhaps the most difficult task, since its subject matter is a type of system of which few examples exist. Nevertheless, there is some interesting work here: a paper by Rick Hayes demonstrates (though does not appear to recognize) the economic infeasibility of using protocol analysis as a knowledge acquisition technique for any but the most intractable problems; Dian Kjaergaard's paper, despite its title, is a fairly pragmatic analysis of the interaction between business, theoretical and political issues which complicate the assessment of commercial customer accounts at an anonymous Danish Bank; while Sandra Brown provides a detailed case-study of an organization's first use of expert systems.

Section 2, dealing with DSS, argues that it is important to consider not merely the relationship of DSS to decision-maker, but the more complex interaction between business problems and wider business context, decision makers and decision making process, and DSS. Papers by Colin Eden and Fran Ackerman—describing the use of computers to assist in the development and management of strategic vision—and Lawrence Phillips, on group decision support, make this point especially clearly. The last paper in this section, by Chris Martin, has a different slant: it analyses the use of DSS by top management and the reasons for lack of such use.

The final section of the book is devoted to expert systems. The first paper, by Anthony Wensley, lists a number of possible research directions for expert systems; while those suggested are clearly valid, there are many less fundamental—but in the short term more important—issues for research which are not listed: for example, the serious business issue of how to estimate the costs of an expert system. Hamish Stevenson provides an excellent analysis of the hype which has surrounded expert systems for the last 5 years; the second part of his paper, in which he attempts to explain this hype, however, reveals a lack of understanding of the business problems of expert systems developers: he criticizes expert systems groups for spending their time on investment appraisals instead of the development of technical prototypes. The truth is that a bit more time spent analysing business benefits, and rather less on unfocused building of prototypes would have prevented many of the

problems Stevenson highlighted so clearly in the earlier part of his paper. Enid Mumford provides an interesting history of the development of the XSEL system drawing a number of lessons from that project—many of which are of general applicability. Lynda Davies, by contrast, considers principles of expert system design in so abstract a way that it is hard to draw conclusions of relevance to any specific development project.

Managers reading this book to understand the business significance of KBMSS will be disappointed: being a collection of papers it is not systematic in its presentation of ideas, it does not address the immediate business issues and it provides very little in the way of practical guidance; it contains too few case-studies and those it does contain have too little analysis of the costs and benefits of the systems involved. As a work of reference for an academic library, it contains several papers of interest and perhaps importance. Its greatest contribution, however, may be to the creation of a focus of interest around the central thesis of the book: that integration of the ideas of DSS and expert systems will produce a significant new class of systems.

Reviewed by Mark Thomas, Ernst Young, Management Consultants, Rolls House, 7 Rolls Building, Fetter Lane, London EC4A 1NH.

Causal AI models: steps toward application edited by Werner Horn. Hemisphere Publishing Corporation, New York, 1990, pp 319, hardcover £36.00.

Life is full of consequences, or so it seems—one thing leads to another. This causality is at the heart of the way we learn and think about the world. Day to day, we easily accept that events are connected over time. More exact notions of causality have intrigued scientists and philosophers for generations, and have now become critical issues for AI.

Causality offers interesting but difficult theoretical AI problems. Researchers attempt to mimic the deceptively simple judgements with and about causality that humans make every day. The hope is that we can develop robust programs capable of reasoning about novel problems. For example, a program designed to diagnose component failures in a satellite might have to deal with unforeseen events. Several components may interact to produce a new system failure, or a component may fail in an unforeseen way because of the unusual environment it is placed in. It is simply not possible to create a set of rules to cover such circumstances. A deeper understanding of how mechanisms operate seems essential for such skills, and at the heart of this problem lies the conundrum of causality.

Causal AI Models presents a collection of 15 original papers which first appeared as an issue of *Applied Artificial Intelligence*. The paper authors are drawn from both the USA and Europe. The collection was motivated by the increasing interest in both causal modeling and qualitative physics within the AI community. The volume's editor notes that in the past, the emphasis within the literature has been to describe the fundamental theoretical aspects of causal modeling. With this book, he aims to demonstrate how one could extend upon and utilize these methods to build applied systems, as well as giving examples of the technology in application.

If the technology of causal reasoning had now reached a sufficiently advanced level appropriate for system development in the real world, then we could expect that many genuine applications would be reported here. However the snapshot of expert systems work presented in this volume shows that it is still largely an experimental field (hence the "steps toward" in the book's title). Looking through the papers, I counted 6 that dealt with purely theoretical issues using toy examples, 5 that reported prototype systems, and 5 that discussed what I felt were genuine applications. None of the latter seemed to be in routine use.

The applications reported include expert systems for vibration-based monitoring of turbine generators, rheumatology, human gait analysis, and heart failure, intelligent tutoring systems for electrical circuits, and a monitoring system for the JPL Space Simulator. Many of the projects reported from the USA have been around for a while and previously documented in the literature,