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,

e.g. Long's heart failure program, MDX2, and Koton's CASEY. I personally found more value in the presentation of European work, most of which I had not seen before.

The major disappointment with the material presented in this volume is the lack of sophistication in the way many of the papers approach causal reasoning. One could cite several recent advances in the theoretical handling of causality—for example Shoham's "Reasoning About Change" or the Causal Ordering methods of Iwasaki and Simon which have passed largely unnoticed in these papers. To be fair, this may partly be a reflection of the time it has taken to assemble the volume. Further, perhaps because of the applications emphasis in the volume, there is rarely significant discussion about the competing methods that the authors might have used in representing causality.

Most of the papers represent causality with methods that have been around for a while. For example causal probabilistic networks are used in several of the projects, despite the frequent difficulty in obtaining the probabilities. Often these have to be estimated and interpreting inferences based upon them is consequently often less than satisfactory. Most of the remaining papers represent domain knowledge with models that consist of device components or functional descriptions. This reflects the emphasis on device modeling evident in theoretical work over the last 5 years.

Causal AI Models is not, nor I suspect was intended, to be a text on causal reasoning in AI. The volume will be of value to those interested in developing expert systems with deeper reasoning mechanisms than afforded by rule-based systems. It presents a wide variety of applications, and one can get a feel for the difficulties and benefits of model-based or causal reasoning methodologies. It is an intriguing snapshot of work in progress, and offers glimpses of things to come.

References

Iwasaki, Y and Simon, HA, 1986. "Causality in device behaviors". *Artificial Intelligence* **29** 3–82.

Reviewed by Dr Enrico Coiera, New South Wales, Australia.

Analysis for knowledge-based systems: a practical guide to the KADS methodology by FR Hickman, JL Killin, L Land, T Mulhall, D Porter and RM Taylor, 1989. Ellis Horwood: Chichester, UK, pp 190, £24.95.

Although there can be few readers who have never heard of the KADS methodology, probably even fewer have had an opportunity to read a detailed exposition of what it entails. Increasingly often it is cited in workshop and conference papers, and several relevant book chapters exist, but it has always been difficult to obtain most of the documentary deliverables (a "KADS handbook" was withdrawn). A large number of these have been generated since July 1983 by the two successive ESPRIT projects responsible for formulating the methodology.

This, the first book devoted to introducing KADS to a general knowledge engineering audience, will no doubt be eagerly read, as KADS is widely recognized as a potential standard for KBS development. Two basic questions may be posed: is it an accurate, sufficiently comprehensive and (as the title claims) practical account of KADS? And may the reader be convinced by the book that KADS is indeed important for the future of knowledge-based systems?

Esprit Project 12 started with just two partners, the University of Amsterdam (particularly represented by Bob Wielinga and Joost Breuker) and The Knowledge-Based Systems Centre at the Polytechnic of the South Bank (London). This was expanded in 1985 (shortly after Frank Hickman joined KBSC) to six partners who sponsored Phase 1 of Project 1098 "A Methodology for the development of knowledge-based systems"; Phase 2 ran subsequently from March 1988 to January 1990. It is believed a new "KADS 2" project will be launched in autumn 1990, with a major aim of commercializing the methodology. The KBSC is now located at Touche Ross Management