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

Consultants, London (for whom all the authors work), where Tim Bingham acts as ESPRIT P1098 librarian and supplier of KADS public domain deliverables.

The fundamental principle behind KADS is that KBS development involves modelling knowledge at a series of levels of abstraction. At the highest level, the epistemological level, four different "layers" of knowledge are distinguished in the KADS model of expertise. The domain layer includes concepts, relations and structures. The inference layer describes the domain layer in terms of the inferences that can be made on the basis of the knowledge it contains. The task layer applies the knowledge in the two lower layers in an ordered way, or task, to achieve a goal. Finally, the strategic layer selects appropriate tasks for problem-solving in plan execution.

Strong emphasis is placed on a rigorous and extensively documented analysis phase, to be completed before any design and implementation; in stark contrast to the formerly-dominant "rapid prototyping" paradigm. KADS has clear origins in the conventional waterfall life cycle model—and also relates well to Boehm's more recent spiral model (Chapter 4 of this book includes discussion of Taylor's CONCH model, which incorporates both): the difference is one of complexity rather than fundamentals. This complexity arises from the "internal analysis" of the knowledge experts draw upon, in addition to the "external analysis", or specification of the external requirements of the KBS. There is also, however, a place for prototyping (as many practitioners today will agree), in a supporting exploratory role; for example, in developing and refining the user interface.

KADS has traditionally espoused a top-down approach to knowledge engineering, in which an "interpretation model" (incorporating one or more "generic task models" [note here, incidentally, that KADS does not share Chandrasekaran's meaning of this term]) is selected as a template to guide knowledge acquisition. When this is "fleshed out", via a process of incremental refinement, the resulting conceptual model is said to capture at a high level of abstraction—the epistemological level—both the underlying expertise and the cooperation required between the problem solver and the user in order to complete the task in question. Elemental categories (or "primitives") at this level include types or roles of concepts ("meta-classes" in KADSpeak), types of inferences ("knowledge sources"), structuring relations and types of strategies.

One problem with this approach is that the knowledge engineer has to recognize, in advance and at an early stage in knowledge acquisition, a mapping from the apparent characteristics of the problem task to the components and structure of this high-level interpretation model. The authors of this book acknowledge that it may not be possible to recognize an underlying generic task (such as heuristic classification) or combination of tasks in the application domain. In such cases a task model must be *created*, "in accordance with the emerging characteristics of the domain" (p. 56), from primitives: essentially, knowledge sources and meta-classes. Even here, however, there is a gap in epistemological support as KADS does not seem to provide adequate guidance on identification and naming of these abstract primitives.

The final chapter of the book, whilst a welcome attempt to tackle this problem of how to construct an interpretation model from primitives (the author, Jonathan Killin, ingenuously acknowledges that "little guidance about construction is available" (p. 154) !), falls short of convincing this reader that the practicalities are fully understood. On page 147 we read this summary:

"What are we modelling, when we build interpretation models? Tasks, which are a kind of organized behaviour, which can be broken down into a set of entities (let us call them meta-classes for the sake of argument) and the way that these entities can relate to each other epistemologically, which is no more than the valid inferences that can be made to transform the components from one state to another."

Several models are shown in the book, some in successive stages of refinement, and the process of building an interpretation model, as discussed by Lise Land in chapter six and Jonathan Killin in chapter seven, is in general comprehensible. However, what is absent is that vital first step in abstraction from the data by which knowledge sources and meta-classes are created. The reader's confidence is not engendered by a sentence like "It was by using this holistic blend of heuristic fudge and domain knowledge that the first version of incremental configuration was constructed" (p. 152).

Recognizing that it is far easier to pick a template "off the shelf" than it is to create one from scratch, the KADS workers have produced a library of interpretation models (each composed of one or more generic task models). Two quite radically different classifications of these generic tasks are presented in the book (by Lise Land on p. 53 and by Tim Mulhall on p. 92): the latter, it is said, "reflects a more recent classification" ("Notes", p. 167). The accompanying claim that this conflict "poses no practical problems for the reader" does not dissipate the feeling that six collaborating authors should demonstrate greater certainty and unanimity; or at least, that the one invested with the additional title of "editorial collaborator" (Robert Taylor) lacked sufficient authority in the role. Another illustration of the degree of terminological flux still recognizable in the project (as represented by this book) is the confusing description of the relationship of generic models to interpretation models on p. 52 (where, at one point at least, they are apparently the same thing); an alternative treatment is given on p. 91. Despite a good (and essential) glossary, references to different types of model ("initial", "generic", "interpretation", "conceptual", etc) in the text would benefit from clearer distinction and more unified and disciplined interpretation.

The principal advance in knowledge engineering offered by KADS to date, aside from the important but incomplete work on epistemological analysis, is arguably the extremely comprehensive "feasibility assessment" represented by the KADS analysis life cycle model. A well-structured set of output documents is outlined in the chapter by Robert Taylor, which is the most useful part of the book from the point of view of practical project management. This draws upon one of the major P1098 deliverables, coded D3 "Requirements analysis in KBS development", by S Barthelemy et al.; like another key document, A1 "Model-driven knowledge acquisition: interpretation models", edited by J Breuker et al., this appeared in 1987. Phase 2 of P1098 was due to concentrate upon issues other than analysis, notably the problematic mapping from conceptual model to design model. We are promised a "sister volume" on KADS design. In his final concluding paragraph, Frank Hickman also acknowledges there is much work to be done on many other areas, including verification and validation (the "new bottleneck" in delivery of practical KBSs, in the opinion of this reviewer; the state of the art seems relatively similar to that of knowledge acquisition when KADS commenced in 1983); and also testing, maintenance, integration and development of a business-oriented tool set.

It is noteworthy that all these under-researched areas relate strongly to acceptance of the technology in business. At this point it is appropriate to return to the questions posed at the beginning of this review. The book does provide an accurate introduction to the greater part of the KADS methodology as reported up to July 1989 (the date of the Preface) and previously read by this reviewer. It is a good, comprehensive account of a rigorous approach to analysis which is designed to "get it right" (as far as possible) before design and implementation commitments are made, after which the error cost increases significantly. Rigour, of course, carries its own costs; for which reason it would seem apparent to the reviewer that KADS is suited to fairly large projects with a budgeted elapsed time of at least a year. Production of the full set of analysis documents alone represents a substantial overhead. The benefit of developing both an explicit model of the expertise and a full system requirements document before design and coding should nevertheless be realised by all non-trivial projects: it is not necessary to adopt KADS in its entirety to profit from the guidance offered in this book.

From a practical point of view, it is a meritorious presentation of an important contribution to analysis for KBS development, although it does not entirely succeed in clarifying the rather esoteric KADS terminology. The book would also have benefited from citation of more than just a couple of applications of the methodology. There is now sufficient evidence in European conference proceedings alone that KADS is being widely applied and publicly reported. Inclusion of these case studies in the list of references would encourage awareness that KADS is not still just a research programme.

Readers who have experience either of conventional systems development or of unsuccessful real-world KBS projects are likely to recognize the relevance of KADS and to welcome this book as the first readily available exposition of the methodology for analysis. We must hope that further volumes appear covering all the other stages and issues in KBS development. The publishers are to be congratulated on a well-produced and speedily published (November 1989) book, which has a good index and very few errors.

Reviewed by Ian Neale, ARIES at City, The City University, Northampton Square, London EC1V 0HB.

Abductive inference models for diagnostic problem-solving by Yun Peng and James A. Reggia, Springer-Verlag, Berlin, 1990, pp 268.

For many years diagnostic problem-solving has been, and still is, an active research area. The understanding and formalization of diagnostic reasoning, especially in medical domains, has intrigued researchers since before the advent of knowledge-based systems. This work has in fact significantly contributed to the technology of expert systems; diagnostic reasoning is certainly one of the best documented applications of this technology.

The book by Peng and Reggia is a notable addition to the existing literature on the subject. The authors by discussing abductive inference models for diagnostic problem-solving, to the exclusion of deductive inference models, on the one hand present an incomplete picture of diagnostic reasoning but on the other hand they emphasise the importance of abductive inference in diagnostic reasoning; in many discussions on diagnostic reasoning abduction takes second role to deduction. Abductive inference, or hypothesis generation, is intrinsically related to diagnostic reasoning. The authors informally define abduction as the process of generating plausible explanations for a set of observed manifestations. Deductive reasoning is also intrinsically related to diagnostic reasoning. Often a large number of equally plausible alternative explanations are generated through abduction and deductive reasoning is required to differentiate between these alternatives; hence the use of the term hypothetico-deductive reasoning for diagnostic reasoning. The "parsimonious covering theory" which is the generic name given to the various models presented in the book, does not therefore constitute a complete theory of diagnostic reasoning. This is not to say that it does not make a worthwhile contribution to the field of diagnostic reasoning, as explained below.

In general, there are two main schools of thought regarding diagnostic reasoning. One school advocates the use of formal models, which by and large are founded on probability theory: the other school advocates the use of heuristic models which emulate as closely as possible the elicited (articulated) thought processes of human diagnosticians. The formal models tend to be simpler in symbolic knowledge but richer in probabilistic knowledge. The heuristic models tend to be richer in symbolic knowledge, and often use qualitative rather than numeric models for handling uncertainty; heuristic models can be difficult to formalize. Formal models are domain-independent, applicable in any domain satisfying the model's underlying assumptions. Generic heuristic models normally constitute generalizations of domain specific models. Formal models and heuristic models can certainly be integrated. There are strengths and weaknesses for either type of model. Which approach is the better one (efficient and effective) for a specific diagnostic domain given the particular design objectives can only be proven in practice.

The above book belongs to the formal model school of thought – although some aspects of the models presented are justified by relating them to human diagnostic behaviour. Readers belonging to the heuristic model school may find this book rather heavy reading. It abounds in lemmas, theorems and proofs. Mathematical appendices can be skipped, but still there is much notation which needs to be understood to appreciate the discussion on the models.

Parsimonious covering theory generates plausible multi-element (i.e. multi-disorder) hypotheses for a set of manifestations. The elements of a hypothesis collectively cover all the manifes-