necessity be largely manual, requiring a major investment on somebody's part. You've got to think that the system is going to get a lot of use to justify the expense.

For many systems, the nature of the thing that we would have be explained is not a static and well-understood body of knowledge. For instance, all this helps me very little in generating an explanation of the behaviour of any expert system, unless the expert system knowledge representation contains deep knowledge of the domain. the information I would need is unknown ahead of time, changes with each run of the system, and is, in general, just too complex and insufficiently understood. There would not be enough time to generate the knowledge base, assuming it could be done at all, and the whole thing would become obsolete so fast that it would be hard to rationalize the effort if it could be done.

Next, these explanation systems are just that—systems that deliver given explanations. Not question and answer systems wherein by asking a question you can elicit an explanation. These systems will take the time and effort to see that you get their message, but it is only the message that they have for you that you can get. These systems require that the system initiate the dialogue.

In summary, Cawsey's book is strong on background and motivation, somewhat weak on analysis and conclusions. The book presumes the reader to have a basic familiarity with rule-based systems, but requires no specific technical or programming expertise. It takes only a few hours to read, and it surveys a large corpus of interesting work in the quest to determine a good information model for human discourse.

Reviewed by Joshua Glasser, Clarkson University, New York, USA

Representing uncertain knowledge by Paul Krause and Dominic Clark, Intellect Books, UK, 1993, pp 277, £14.50, ISBN 1-871516-17-X.

This book provides a broad overview of the most important representational approaches and imbeds these approaches in an interesting taxonomy of various types of uncertainty. The topics covered include Bayesian probability, MYCIN certainty factors, Dempster–Shafer theory, possibility theory, non-monotonic logics, and endorsement/argumentation approaches.

The style of presentation is understandable, mostly informal, and example based. The book emphasizes the goals and overall structure of the various uncertainty representation approaches, rather than presenting "how to do it" details. As such, it is excellent reading for one who is new to uncertain reasoning in AI systems, or for one who is knowledgeable in some approaches, but would like to know more about other approaches. With its excellent set of references, both to seminal work and recent developments, this book can also serve as a starting point for one beginning research in the area.

At its core this book is part analysis and part journalistic reporting. It is an analysis of the goals of various research programs, but only a report of claims made for achievement of these goals. Indeed, a major caveat in reading this book is to differentiate its analytic content from its reportorial content. Most approaches are discussed from the point of view of each approach's advocates rather than from the point of view of a detached, unbiased observer. This method of discussion serves the authors well as they relate the goals of the various research programs to their excellent taxonomy of uncertainty. The method succeeds in this aspect because researchers tend to be clear in articulating their goals and in differentiating their work from that of others. But the method fails in the aspect of providing any substantive analysis or comparison of the actual achievements. Researchers are not known to dwell on the shortcomings of their favorite theories and of their own work.

One illustration will serve to show the type of analyses and comparisons which are missing from this book. In the discussions both of probability theory and of formal logic, the underlying semantics (sample spaces in probability, models in logic) are raised as significant issues relevant to uncertainty representation. In the discussions of the non-additive quantitative techniques (e.g.

possibility theory and fuzzy logic), the issue of underlying semantics is ignored. Semantic issues do rarely appear in such papers, because, in my view, these approaches have little or no semantic foundation. I am sure that had these issues been raised in the papers of the non-additive advocates, they would have appeared in this book. Nonetheless, it seems that if the issue of semantics is important in other approaches, the authors might have pointed out the absence of such discussion in the non-additive literature.

Another result of the reportial approach of this book is that the reader might well conclude that probability theory is the most discredited of all the representations of uncertainty. Though they do claim in the beginning of the book that probabilistic techniques represent the "gold standard" of uncertainty, in the remainder of the book probability is simply a target of criticism. Again, this is because of the nature of the uncertainty literature. In this literature, it is fashionable for one to show the value of his or her technique by discussing the aspects in which it overcomes perceived shortcomings of probability. Rarely does one mention the aspects in which probability may be superior, or provide comparisons of one's own theory to other, non-probabilistic, techniques. Likewise in the literature (though hardly in private discussions!), the fashion of probability advocates is to ignore competing theories.

Despite these shortcomings, I heartily recommend this book to those interested in a broad survey of the field. The shortcomings are primarily reflective of the state of the literature in the field. The strength of the book lies in its broad coverage and its understandable presentation, well strewn with good examples.

Received by John F Lemmer, Clinton, NY, USA

From animals to animats 2 edited by Jean-Arcady Meyer, Herbert L Roiblat and Stewart W Wilson, MIT Press, 1993, pp 523, $55, ISBN 0-262-63149-0.

It is time animats lived side by side of animals. They can live, learn, build maps, evolve and be our domesticated pets, too. In the intersection of AI, robotics, philosophy and psychology, traditions of building ever more complex systems that try to exhibit intelligent behaviour or solve puzzles are giving way to systems that live simple lives, without pretending to reason deeply, or to act as if they were real animals. Instead, the system designer builds simple agents and endows them with basic shallow abilities to interact with the world. The point is that the designer doesn't solve all the problems that the system may encounter beforehand. Nor does the designer build a system that acts as if it is trying to talk to itself outloud. Many of the traditional symbolic systems are designed so that they can be made to generate a transcript of human understandable steps they take to decide on performing each and every action. Once deployed in a complex world, the animat (just as a device would) uses its repertoire of interactive behaviours as they become needed, e.g. a microwave oven responds to buttons pushed. The interactions that follow may appear intelligent. Such agents can also be made to learn from these interactions and enrich their range of capabilities.

This volume is an excellent collection of papers to represent this shift in paradigm. Due to space limitations in this review, I will refrain from commenting on each paper. Instead, I will comment section by section. This field hasn't reached maturity and must allow for various approaches. On the whole, the editors have done a good job of setting up sections with some confusion and overlap among action selection and behaviourial sequences, evolution and learning.

The four introductory papers in the animat approach to adaptive behaviour herald the theme of the paradigm shift through the rest of the papers. There is some level of publicising terms and jargon. For example, Maes calls the paradigm shift "behaviour based AI". The paradigm shift is characterized by learning, adaptation and lack of goals directedness.

One of the themes of the paradigm shift is the need to model immediacy of sensing to acting or close knitting of sensations to motor control. Papers in perception and motor control present methods for making these relatively close connections. Presentations here are varied but mostly