fields in isolation, but it is her emphasis on empirical work, both to found various aspects of the EDGE system and to evaluate it as a whole, which is truly distinctive within AI.

In brief, the heart of the EDGE system is a central explanation planning subsystem. This subsystem takes input from three sources: dialogue planning rules as developed within conversation analysis; content planning rules designed to express the required semantic content; and a domain specific knowledge base. It also interacts with a model of the discourse and a model of the user. There is no real language processing within the system, which takes input via a system of menus and uses simple templates for generating text output. The system is developed within the domain of rudimentary electrical circuit analysis, and graphical interactions via circuit diagrams are also supported.

The user model is dynamic, being constantly revised on the basis of user's responses and interruptions. Explanation is goal driven, but the goal may change as the user model is revised. Planning and execution are consequently interleaved, with some reactive strategies also being employed. The result of planning is an agenda of partially ordered explanation goals, with the order of the components of the resultant explanation being determined by discourse context.

Cawsey begins with a brief introduction, outlining the context and goals of her work. Chapter 2 examines the structure of explanatory discourse, with special reference to transcripts of expert/novice tutorials. Chapters 3–5 then address the three major areas (text planning, dialogue planning and user modelling) which Cawsey draws together in the EDGE system. In each case, she begins by surveying the field before focusing on the relevant component of the EDGE system. None of the material is particularly ground-breaking in itself (except, perhaps, Cawsey's incorporation of graphical actions into the dialogue). It is the way which Cawsey pulls together the various fields that characterizes her work. This pulling together is demonstrated in Chapter 6, where detailed examples of the functioning of the complete EDGE system are given.

The concluding chapter frankly evaluates the system, noting its strengths but highlighting its weaknesses. Cawsey is harsh on her own system, listing many inadequacies noted by real users. Some of these she addresses, but others, Cawsey notes, require substantial modifications to the underlying explanation-generation architecture.

Four appendices are also included: a selection of complete human explanations (including expert/novice interactions); a description of the planning rules and domain knowledge; a selection of computer/novice dialogues generated with the system; and a selection of sample screen displays.

The book is clearly structured and well written. Though rather specialized, its appeal is broadened by the overviews given in each of the content chapters. Its target audience thus includes researchers in the fields of text content planning, organization of discourse and user modelling. The book provides sound introductory material for each of these fields, and it should be of special benefit to those willing to consider the relationships between them.

Reviewed by Richard Cooper, Department of Psychology, University College London, UK

Explanation and interaction: The computer generation of explanatory dialogues by Alison Cawsey, MIT Press, USA, 1993, pp 232, £26.95, ISBN 0-262-03202-3.

Explanation systems are commonly designed to present information at some fixed rate, or level of abstraction, with the assumption that the user possesses a sufficient level of prerequisite knowledge and expertise within the domain of the explanation, and so will be able to use well whatever explanation gets provided. But such explanations may be ill-adapted for communicating information to some subset of users because this assumption may either not hold for these users, or may not hold uniformly for these users over the target domain. Thus, bored experts and bewildered beginners (or *vise versa*), and users who follow some but not all of the explanation, with the frequent sense of being lost.

Many explanation systems try to minimize this ineffectiveness by allowing the user to identify his

or her level of expertise in an area as low, middle or high, and then to produce explanations targeted to the declared level. But explanations may still be ill-adapted for communicating the information, because the users may not hold the declared level uniformly over the target domain.

So, how can you know if the user is following the explanation from moment to moment? And how can you remediate to maximize the explanation's effectiveness? Many researchers, across several disciplines, have been investigating what goes on in actual human explanatory discourse, and many information models for explanatory discourse have been proposed in efforts to answer these very questions.

In this book, Alison Cawsey presents an extensive survey of the many information models for explanation and explanatory discourse that have arisen from these investigations, past and ongoing, into human explanatory discourse. Intertwined with this, the author also introduces and discusses the EDGE system, which is a prototype system that brings together a lot of the ideas, and which provides a powerful illustration of the potential and efficacy, of the theoretical and foundational work presented.

The book has been laid out in seven chapters, and focuses on the following four central divisions in the categorization of the research into the generating of user-sensitive explanatory dialogues.

First, 'Analysing human explanatory dialogues'. Here Cawsey presents a discussion of the fundamental research into developing schema and heuristic models for dialogue, based on the analysis of transcripts of actual human discourse. The grammar-based models, and especially those of McKeown, Paris and Sinclair and Coulthard, are highlighted in a thoughtful overview of the literature. These models seek to provide the overall structural framework for presenting the explanations.

Second, 'Planning explanatory context'. Here Cawsey presents a discussion of the research, such as the recent works of Moore and May, into developing models that order the sequence of ideas to be presented within an explanatory dialogue. How text planning is done within the EDGE system, and a discussion of the trade-offs considered in deciding between vying approaches is interleaved into this well-focused overview.

Third, 'Managing the dialogue with the user'. Here Cawsey brings together work from both sociology and artificial intelligence to present a summary of research into managing interaction with the user. Specifically, the issue discussed here is how a system can present a cohesive explanation given that the user may at anytime interrupt, requiring that the system digress and later resume the explanation. The focus of this section is a presentation and discussion of how dialogue with the user is managed within the EDGE system.

Fourth, 'Updating the user model'. Here Cawsey presents a discussion of the research into modelling the user's level of expertise and the degree to which the user is following an explanation. Such models of the user are central to generating user-sensitive dialogue. Work by Chin and Sleeman, and many others, is featured in this section, where the focus is on how the user is modelled and how that model is updated within the EDGE system.

The balance of the book pertains to the EDGE system. The book is rich in examples throughout, as well as providing several extended examples as Appendices, and the author concludes with suggestions for further research.

Here's what I think. The EDGE system is certainly a significant step down the road towards a good theoretical basis for explanation. But the prototype and the discussion of supporting theoretical framework both highlight what I see as a real hindrance to their widespread application as anything other than tutoring or intelligent help systems.

Primarily, these systems all require that the information over which an explanation is to range must first be thoroughly decomposed into some given tightly connected knowledge representation. All the interdependencies and relationships between all the aspects of the knowledge must be expressed. What's more, once the knowledge is all broken down, the given representation must also be liberally embedded with choice verbs, nouns and phrases for the explainer to find and use when assembling explanatory utterances.

Given even a limited domain, this is not easy. The construction of this representation must by

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.