seem to be more useful as an illustration of the technique rather than as a serious alternative for implementing Prolog. Nevertheless, the project as reported here was still in its youthful stages, and one would not expect issues of efficiency to have been aggressively addressed.

Tim Flannagan's '*PR: a logic for practical reasoning*' made it into the book. The motivation behind this work is to develop a logic for planning and for action, drawing heavily from Kenny's 'logic of satisfactoriness'. The latter seems to be a work aimed at the philosophical logic community. With the example

> I want to be in San Francisco by noon
> If I take the 11.30 train I will be in San Francisco by noon
> So I'll take the 11.30 train

Flannagan introduces the inference schema of 'co-modus ponens'. That is:

$$\frac{G,\ C \rightarrow G}{C}$$

This rather looks like abduction to me, in which case the properties of defeasibility do not seem at all surprising. Dov Gabbay among others has on occasion talked enthusiastically about the properties of abduction, and I suspect that what Tim Flannagan is trying to express could be expressed more clearly in these terms. But that is perhaps a statement of my prejudice against heavy philosophy.

Dov Gabbay's mischievously entertaining presentation style can belie the seriousness of what he has to say. The aforementioned paper by Gabbay included in this book is comprehensive and hard work. It is the second part of an extended study of modal and temporal logic programming, and does demand a much higher degree of mathematical sophistication to read than any of the other papers in the book. I think this is a pity, and suspect that the paper will go largely unread by most purchasers. A little more guidance through the paper for the uninitiated would have been useful. It is highly interesting and original work as you would expect. But in deference to the hard of understanding like myself, I rather think that this is the wrong place for this paper to have appeared in its raw form.

So what have we got left? The handling of incomplete information in a logical database is discussed by Kong and Williams; 'Virtual logical neurons' implemented in a concurrent logic language by Kozato and Ringwood; Kelpie, by Hamish Taylor, addresses the development of a concurrent logic programming system for large scale database applications. Referring back to my opening comments, the interfacing of logical languages to large databases is a particularly fast moving field of research at the moment, as is the development of concurrent logic languages themselves. One wonders a bit just how much these three papers reflect current technology?

This book does contain an interesting selection of papers, and they are all certainly worth reading. The main reason for a slight feeling of disappointment with it is its very limited coverage of the material which was actually presented in the workshop of February 1989. I think that in an ideal world I would like to have seen a slightly better balanced presentation of material, and a much broader coverage of the state of the art at the time. But it is worth a read all the same.

Uncertainty and vagueness in knowledge based systems by R Kruse, E Schwecke and J Heinsohn, Springer-Verlag, Berlin, £35.00. ISBN 3-540-54165-9.

When ECAI was held at Stockholm a couple of years ago, Mark Wallace of ECRC began his presentation with words to the effect: 'There seem to be two general approaches to giving short presentations at conferences like this. The first is that of the established figure who just gives one or two pointers and refers the audience to the literature for further information. The second is that of the enthusiastic young researcher who wants to tell the audience everything he or she has learnt.

Well, I'm afraid I am going to do the latter'. As I think do Rudolf Kruse and his colleagues in this book.

This is a very comprehensive review of the use of numerical techniques in uncertainty handling written in the context of a presentation of the authors' own research. The plus side is that there is a great deal of valuable information in this book, including much new work. The downside is that enthusiasm to include so much information has been at the expense of clarity of presentation, to a certain extent. And the sheer volume of work involved in producing the book to quite tight deadlines by very busy people has meant that there are quite a number of production errors in it. I will come back to these points later as I think they are quite important. But more important is the actual intellectual content of the book, which is significant and original.

This is very much an applied mathematician's book. In fact it is more constrained than that; with the exception of a very brief final chapter, it concentrates exclusively on *numerical* methods of uncertain reasoning. It is the author's contention that '. . . vagueness and uncertainty can be handled adequately by using measure theoretic methods. The presentation of applicable knowledge representation formalisms and reasoning algorithms (in this book) shows that efficiency requirements do not necessarily require renunciation of an uncompromising mathematical modelling approach . . .'. The term modelling tends to mean all things to all people, but here the building blocks used are measure theory and set theory, and it may help any intending reader if they have a reasonable familiarity with measure theory. I do wish to stress that this is but one approach to modelling uncertain reasoning in knowledge based systems, and this book does not address any of the alternative symbolic approaches to any serious extent. However, it is not my wish to use this review to extend the extensive discussion of the relative merits of each approach; the book should be judged on its own terms of reference.

The distinction is made between two aspects of ignorance which are of importance in knowledge based systems; uncertainty and vagueness. By vagueness is meant that the exact value ω_0 of some datum is not known, but is restricted to some subset which covers ω_0. Thus, whilst we may not know the precise location of a ship, we may be able to say it is 'at sea' (the set of all possible locations in the oceans). By uncertainty is meant that the validity of the datum has to be determined. That is, how much do we believe the proposition 'the ship is at sea'? If Ω is the frame of discernment (the set of all possible outcomes) then vagueness is modelled by allowing data to be expressed as subsets of the frame of discernment, while uncertainty induces a mass assignment m to subsets, where m: $2^{\Omega} \rightarrow [0,1]$.

The word 'imprecision' is perhaps a better word to use than 'vagueness' in the above, in that what is being modelled is imprecise data, and not the compatibility of data with vague concepts (as the authors make clear). In fact, for the straightforward expression of data as subsets of the frame of discernment the authors do prefer the use of the word 'imprecision'. They usually reserve the word 'vagueness' for the presentation of data in terms of layered sets (L-sets) rather than subsets. This is a refinement which is given a precise meaning in the book and which may become a little clearer in the following.

First, consider the case of precise data. Here we may consider a set of sensors or agents Θ as providing data as elements of the frame of discernment Ω; Γ: $\Theta \rightarrow \Omega$ where Γ is an 'observation mapping'. The sensors may be unreliable. If each sensor has an associated probability of giving the correct value, then the observation mapping Γ will map this to a probability distribution over Ω. Suppose now the sensors can only give imprecise data. That is, they can only return a subset which covers the true value of the datum of interest. In this case, the observation mapping maps sensors to elements of the power set of Ω; Γ: $\Theta \rightarrow 2^{\Omega}$. Then Γ maps the probability assigned to the sensors to (probability) mass assignments on subsets of the frame of discernment. L-sets arise as a mathematical representation for the case when the identification problem can be considered from a number of different points of view, or *contexts*. An imprecise description of the datum $\omega_0 \in \Omega$ of interest is obtained in each of these contexts. So the sensors return a class of subsets instead of a single subset; Γ: $\Theta \rightarrow \mathcal{F}_L(\Omega)$, where $\mathcal{F}_L(\Omega)$ is the set of layered sets of W. This last is what the authors mean by vague data. For this case of vague data the authors introduce the notion of a

generalised mass distribution in which the evidence mass may be assigned to elements of $\mathcal{F}_L(\Omega)$. Actually, what they say (on p 150) is 'In *spite* of those cases we introduce the notion of . . .', which I don't think is quite what they mean (my emphasis). A slight mistranslation perhaps?

The following completes a very crude outline of the authors' model. Let Θ be a set of sensors or agents, $\mathcal{P}$ a class of probabilities on Θ. Then if Γ is an observation mapping from Θ to Ω, 2^Ω, or $\mathcal{F}_L(\Omega)$, where Ω is the frame of discernment, the pair $(\mathcal{P}, \Gamma)$ is an *information source*. One needs now to consider the dynamic aspects of the model; methods of updating the mass assignments derived from an information source in the light of additional information. In the case where Γ: $\Theta \rightarrow 2^\Omega$, the combination of evidence from two independent information sources corresponds to Dempster's rule of combination. As well as *combining* information sources, methods of *updating* information sources are also considered in two basic forms; conditioning and data revision. The first is the case of integrating absolutely reliable but imprecise evidence. That is, the datum of interest is known with certainty to lie in a subset E of Ω, and the mass distribution derived from an information source must be revised accordingly. The second case corresponds to an actual revision of the observation mapping Γ.

This approach can be used to model the probability, Dempster-Shafer and Possibility theories. It is more general than all of them, and gives a clear account of the set theoretic and probabilistic natures of each of them. However, it does require some fairly painstaking analyses of the various mappings between different sample spaces, and of the probability and mass distributions they all induce. Consequently, this is not by any means an easy book to follow, and this is where we come back to the opening remarks.

A request to 'tell me everything you know' does not necessarily lead to a coherent and clearly laid out presentation. As well as including material which develops the above approach in a way which is suitable for a research monograph, the book also includes an extensive survey of existing numerical methods of uncertain reasoning in knowledge-based systems; various heuristic and probabilistic approaches, the Dempster–Shafer theory of evidence and fuzzy set theory. Some of these discussions are informed by the authors' own perspective, but many, particularly the discussions of the heuristic models of MYCIN, RUM and INFERNO, add little to the material which already abounds in the literature. That is not to say that it is not useful to have all this collected together in one place. It's just that there is so much in this book that it may have been better to have spent the time on clarifying the presentation of the important parts of the book, rather than to include some of the outline survey material.

I feel the presentation does need clarifying. For example, I found, indeed still find, the introduction of L-sets in Chapter 3 a little bit too succinct. Worked examples are used quite effectively throughout the rest of the text, and it would be useful to have a slightly more extended example here to get a clearer picture of what is meant by the viewing of data from different contexts. Throughout the text most concepts are defined clearly, but once or twice something slips in without introduction. For example on p 86, Definition 5.1 starts with 'Let $(\Theta, \mathcal{A}, P)$ be a probability space . . .'. I could not recall the idea of a probability space being introduced when I came to that point. The notion of a probability space is a well defined concept in measure theory, so the reader may be expected to know what is meant. But I think that the relationship of $\mathcal{A}$ to Θ (i.e. $\mathcal{A}$ is a set of events on Θ; strictly a *σ-field* of subsets of Θ) should be made explicit. A minor oversight, but it is symptomatic.

I mentioned what I take to be an unintended use of the phrase 'in spite of'. Two of the authors (Kruse and Schwecke) are at the Technical University of Braunschweig, whilst the third (Heinsohn) is at the German Centre for Artificial Intelligence at Saarbrücken. As the book is aimed at an international scientific audience, they have not written it in their native language but have chosen to write it directly in English. There are a few grammatical and other errors, but mostly the English is very good. Once or twice one gets a feeling that they would have been much happier explaining things in their native language (for which they have every sympathy). But I think that most places where I felt there was a need for further clarification, it was due to the succinctness of the presentation rather than their use of English.

In conclusion, this is a slightly unhappy juxtaposition of a research monograph and a survey of numerical methods in uncertain reasoning. Unhappy mainly because the work required to produce such a comprehensive volume has been to a certain extent at the expense of the clarity of presentation of the important novel aspects of the work reported in the book. Further signs that the burden slightly overloaded the authors are given by a number of errors present in the cross-referencing in the text, and omissions in the references (for example, many of the references in Chapter 4 in particular are not listed at the back of the book). But it is an important statement of current research on numerical methods in uncertain reasoning, and is a significant contribution to the literature for that.

A beginner's book of TEX by Raymond Seroul and Silvio Levy, Springer-Verlag, Berlin, 1991, DM 58, ISBN 0-387-97562-4.

TEX is a powerful document preparation program, developed by Donald E Kuth to produce high quality typesetting, especially for technical and mathematical texts (it also has enough features to make it useful in other fields). It is a markup system, where the document text is interspersed with typesetting instructions, and it needs to be run through the TEX program before the formatted document can be viewed. The TEX program produces a page description in a device-independent form that needs in turn to be fed into an appropriate driver for it to be displayed on screen or to be sent to the printer.

TEX is macro programmable. There are a number of macro packages available which provide very high level commands for creating most common styles of documents with a relatively small amount of formatting effort. Perhaps one of the most widely used of these macro packages is Latex.

TEX is public domain and is available for a number of hardware platforms; PC compatibles, Unix workstations, etc. TEX documents only contain printable characters from the ASCII character set, making their transfer by electronic mail particularly straightforward.

A beginners book of TEX, as the name implies, does not aim to be a replacement for *The TEXbook*, written by Knuth himself, which remains the definitive reference for TEX. This book aims to provide a more accessible alternative for users who do not need to know the more obscure features of TEX.

The book has a large combined dictionary index section which makes up roughly one third of the book. You may need to refer to this section to follow some of the examples in the book that contain TEX features before they are actually introduced. The descriptions for the entries in this part are good enough to make the book useful for reference.

The main part starts with an introductory chapter which, among other things, gives an overview of the different stages involved in processing a TEX document. This chapter provides useful explanations for some of the cryptic prompts and error messages that are displayed when TEX encounters errors while processing a document.

It then quickly covers the basics of TEX in the next seven short chapters, using examples to demonstrate the application of each feature. The remaining five chapters cover the typesetting of mathematics and the more advanced features of TEX programming.

As a TEX novice whose experience of TEX was limited to what is required to use Latex to format simple non-mathematical documents (which is very little!), I found it much easier to negotiate my way through this book than Knuth's *TEXbook*. This was especially so when looking up a particular TEX feature.

How useful this book will be to beginners I think depends very much on how they intend to use TEX. If one of the existing macro packages like Latex meets all your formatting needs, then this book is going to be of limited use to you. However, if your formatting requirements are more demanding, and on occasions you need to use raw TEX to achieve a desired layout, then this book should prove very useful.