the "focus is instead on the *logical* structure of the beliefs—what matters is how a belief coheres with the other beliefs that are accepted in the present state" (p. 8). Both approaches are given a full airing here, starting with a discussion of their relative merits by Jon Doyle himself.

With respect to the foundational approach, further distinctions may be made. One is to differentiate between *revision*, in which new information is obtained about a *static* world and *updating*, in which new information is obtained about a *changing* world. A paper by Katsuno and Mendelzon argues that there are important differences between the approach which one should take to updating as opposed to revision. Their argument that revision in the AGM sense is inadequate as a model of rational belief change caused by the second kind of information, is continued in a paper on "Planning from first principles" by Michael Morreau.

A second distinction is between those who consider revision should be applied to belief *sets*, and those who consider belief *bases* as the focal point. The AGM approach considers belief revision from the perspective of the axioms explicitly represented in a knowledge-base *and* the logical consequences thereof; a belief set. Papers by Hansson and Nebel in this volume consider that the postulates for revision operations should focus on the *bases* for belief sets (the explicit axioms) rather than on the sets themselves. This model is based "on the intuition that some of our beliefs have no independent standing but arise only as inferences from our more basic beliefs" (p. 89).

Making connections between apparently disparate branches of a discipline can often provide deep insights. This is exemplified here with a paper by Dubois and Prade in which possibility theory is linked to the epistemic entrenchment orderings of AGM. This enables them to explore the links between rational updating and the handling of inconsistencies in their possibilistic logic. They also demonstrate how their model of a belief state can be used to describe updating with uncertain pieces of information.

Further papers cover belief revision and communication (Julia Galliers), conditionals and knowledge-base update (Cross and Richmond) and a detailed analysis of the logic of theory change (Hans Rott). In all, the book provides a comprehensive and thoroughly up to date survey of the state of the art in formal models of belief revision. It is well and thoughtfully produced, and can be highly recommended both for those wishing to gain an initial understanding of the field and for those already familiar with it who wish to broaden their outlook.

Reviewed by Paul J Krause, Advanced Computation Laboratory, Imperial Cancer Research Fund, London, UK

Textbooks for artificial intelligence

Essentials of artificial intelligence by Matt Ginsburg, Morgan Kaufmann, 1993, £26.50, pp 430, ISBN 1-55860-221-6

Formal concepts in artificial intelligence—fundamentals by Rajjan Shinghal, Chapman & Hall, 1992, £19.95, pp 666, ISBN 0-412-40790-6

Artificial intelligence—structures and strategies for complex problem solving by George F. Luger and William A. Stubblefield, Benjamin/Cummings, 1993, pp 740, ISBN 0-8053-4780-1

Knowledge engineering, Volume 1, Fundamentals edited by Hojjat Adeli, McGraw Hill, 1990, £21.95, pp 345, ISBN 0-07-000355-6

Knowledge engineering, Volume 2, Applications edited by Hojjat Adeli, McGraw Hill, 1990, £31.95, pp 352, ISBN 0-07-000355-7

To assess any book that professes to be aimed at students of artificial intelligence one must first consider what is the best way to go about teaching the subject. This is no easy matter since it requires some opinion about the best way to go about "doing" artificial intelligence—by which I mean the best way of progressing towards the aims of the discipline. The question as to what constitutes these aims is somewhat controversial but is much less controversial than that of the best way in which one should pursue these aims. On the latter question the world of artificial intelligence research is divided into a number of camps. The first division seems to be between those who believe that logic is central to the pursuit of intelligence, and those who believe it is at best a useful tool and at worst a distraction. Contributions to this debate can be seen at their most vociferous in Nilsson (1991) and Birnbaum (1991), but can be found throughout the literature. The second division is between those who believe in formal approaches as a means of expressing ideas with perfect clarify and precision and those who prefer informal approaches on the grounds that the aims of research into artificial intelligence are beyond the bounds of applied mathematics.

Now I have no wish to stir up these debates, but since my views on the way that artificial intelligence should be taught are bound to be coloured by my beliefs about the way the discipline should be pursued, it is best if I admit which camp I stand in. Well, I don't really stand in any one camp—being a natural eclectic I sit on the fences between them. I think that logic has a good deal to offer artificial intelligence, but that making truely intelligent systems will require the incorporation of ideas that are drawn from non-logical approaches, for instance to deal with uncertainty and imprecision. I also believe in the use of formal descriptions to express ideas clearly, but accept that it is easy to "over-formalize" so that the essential idea behind a system is lost in a morass of mathematics when there is no parallel explanation in natural language. So, overall, I tend to the formal and logical, but with the proviso that logic is not everything and formalism can be incomprehensible.

My viewpoint is also coloured by my experience of teaching, learning and "doing" artificial intelligence. When I first came across artificial intelligence, the aspects that intrigued me sufficiently to encourage me to start "doing" it were the fun aspects such as game playing, solving problems like "missionaries and cannibals", and constructing plans by which robots could carry out tasks. The small amount of logic that I came across seemed rather tedious and unnecessary. However, when I came to start "doing" artificial intelligence in earnest I wished I had learnt more logic, and indeed signed up for a course on it soon afterwards. When I taught part of a course on artificial intelligence I found that the students had much the same view that I had had. As the course evolved over several years, becoming more formal and involving more logic, it became known as a hard rather than a fun course, fewer students took it, and those that did take it tended to avoid questions on logic in the exams.

Taking all of these together it is reasonably clear what kind of a textbook I would most like to see. It would introduce logic in some form, preferably without being too dry, but would also talk about game-playing and state space searching, relating search to theorem proving, and giving an impression of how all the basic techniques are related to one another. It would also introduce methods of knowledge representation other than logic, and talk about the tasks that the basic techniques can be applied to such as planning, vision and expert systems. And, if possible, it should be easy to read.

Thus it was a great pleasure to come across Matt Ginsberg's *Essentials of artificial intelligence*. Ginsberg is a believer in the use of logic, but he is also a pragmatist. Thus the book includes four chapters on logic, but also includes four chapters on other methods of knowledge representation, and always stresses the kind of tasks that can be handled rather than the theorems that can be proved. Ginsberg is also definitely a formalist, but he knows the value of being informal. He does not stray into the trap of including pages and pages of symbols that might discourage faint hearts, and always includes a clear explanation of those symbols that he does use. Best of all, in my opinion, there is a strong sense of fun that runs through the whole book from the unprovoked picture of the Eiffel tower, through the anecdotal footnotes, to the picture of "Ginsburg when not working on this textbook" (upside down in an aeroplane). Apart from the momentary amusement of each little joke, this kind of irreverence lightens the whole book making the whole thing easier to read and understand without ever trivialising the content.

The content of *Essentials of artificial intelligence* is exactly what I require of a text book. After the usual overview of the subject, the first topic is search, covering basic methods, heuristics and game playing. This is followed by a comprehensive introduction to logic and then a survey of other methods of knowledge representation discussed in terms of their relation to logic—what they add to first order logic and why they are necessary. The other methods are fairly comprehensive including assumption-based truth maintenance, non-monotonic logic, probability theory (including the only account of the state-of-the-art approach of Bayesian networks that I have seen in a general textbook), frames and semantic networks. Having discussed these foundational techniques, the ways in which they can be applied to solving problems is dealt with. In turn the problems of planning, learning, vision, natural language understanding and processing, and expert systems are sketched and some approaches to their solution outlined. In my view this content steers a good line between broad coverage and detail, handling all the necessary basics and showing how they can be applied to the problems in which people are interested.

The only slight flaw in *Essentials of Artificial Intelligence* is that it is so very American. Thus the reader is invited to sneer at the quaint traditions of Oxford, where Ginsberg held a research post and didn't drink sherry, and to accept with a straight face the formulation in first-order logic of the statement that "Ginsburg loves his country". But this minor difficulty is well worth putting up with in order to profit from this otherwise excellent book.

Unfortunately I don't think that Rajjan Shinghal's *Formal concepts in artificial intelligence* is in the same league. To some extent that is because of the degree to which Ginsberg's book meets my requirements for a textbook, but it is also because I think *Formal concepts in artificial intelligence* makes a couple of serious errors. First of all the book places too much emphasis on formality in general and logic in particular. Thus it includes great tracts of indigestible symbols which I found rather off-putting, and this is not helped by the appalling typesetting which makes it easy to wander from the text into a lengthy figure caption without really noticing. The book also plunges straight into logic and automated reasoning (including a chapter on resolution refutation and its variants which is unnecessarily detailed) without saying anything about search. This then leads to some interesting convolutions when trying to elucidate backtracking, which only really makes sense in terms of graph search. Thus my first objection to the book is closely bound up with my second, which is that the order of topics is rather peculiar. Not only is logic introduced in great detail before anything else, and to the exclusion of most other methods of knowledge representation, but directly it is over and done with the book launches into two hundred pages of detailed discussion of natural language processing. To me this seems totally bizarre. Whilst it is doubtless an interesting and fruitful research area, natural language processing is hardly a "fundamental", and to my mind the prospect of having to wade through a hundred and thirty pages of grammar would be enough to put the most enthusiastic student off artificial intelligence for life. Especially since the book has some very strange ideas about what constitutes an English sentence (since the sentence "the girl injured began to scream" is not an interesting example of a postnominal adjective but bad English).

However, that is just my point of view. Given that after the tract on natural language *Formal concepts in artificial intelligence* goes on to cover production rules and expert systems, state-space search and game playing, it is general enough to count as a textbook for fans of logic and natural language processing who have a healthy contempt for all areas of artificial intelligence outside their own interests. I certainly would not recommend it.

Much more to my taste is Luger and Stubblefield's *Artificial Intelligence—structures and strategies for complex problem solving*, which ressembles nothing so much as a hybrid of Ginsberg's *Essentials of artificial intelligence*, Bratko's *Prolog programming for artificial intelligence*, and Charnaik, Riesbeck and McDermott's *Artificial intelligence programming*. By this I mean that the book covers the bulk of the material that *Essentials of artificial intelligence* contains, and also has detailed discussions on how to create basic implementations of a large part of the theoretical material using LISP and PROLOG. In my opinion, this wealth of knowledge on artificial intelligence programming largely makes up for the lack of some of the material included in *Essentials of artificial intelligence*. In addition, through the humour of some of the quotations at the beginnings of chapters, which range from Newell and Simon to Gary Snyder, *Artificial Intelligence—structures and strategies for complex problem solving* manages to project the same easy-going feeling that made Ginsberg's book so appealing.

The theoretical material that is presented is structured in what I consider to be a reasonably sensible manner. It starts with some history and an overview, follows on with a little logic, state space search and game playing, and then introduces Lisp and Prolog and discusses how to implement search mechanisms using them. Having covered some basic principles the book then addresses other methods of knowledge representation, rule-based expert systems, natural language understanding, automated reasoning (including resolution and the relation between state space search and resolution techniques), and a number of methods of machine learning. Then, as before, some of the theory is turned into programs with the implementations including some Prolog meta-interpreters and a natural language parser, along with a Lisp expert system shell and a version of the ID3 learning algorithm. Object-oriented programming is also discussed.

The stress that is placed on implementation is clearly a very good thing if one is interested in teaching artificial intelligence in a "hands on" fashion. With the code that is available, it would be possible to choose a path through the book such that every concept covered was illustrated by a piece of code that students could play with and yet sufficient material was covered that they ended up with a good grasp of the subject. Given that the publishers will supply the code on request, along with a public domain PROLOG interpreter, using the book as a basis for an introductory course on artificial intelligence is an attractive idea.

The introduction to the volumes *Knowledge engineering* edited by Hojjat Adeli claim that they "are suitable as primary reference books in introductory courses on applied AI and knowledge engineering." In my opinion this is simply not true. What Adeli has done is to invite various luminaries in various sub-fields of artificial intelligence to write a chapter on a particular topic. These chapters are then gathered together in no particular order, and with no linking material, so that they are much closer to being the usual collection of papers than a book that every student should have. Of course, there is nothing wrong with collecting together a set of papers, and many of the papers are very fine, but they do not constitute any sort of coherent whole and they certainly don't make a textbook. Of the twenty papers that make up the two volumes I believe only a quarter could be considered introductory, and these include some of the worst—a couple giving the impression that they were thrown together in the course of an afternoon. However, having dealt with these the rest are considerably better. A number of the papers give authoritative surveys of particular areas that are of sufficient standard and depth to be of interest to postgraduate students just starting out in the field as well as being engaging for people "doing" work in other areas of artificial intelligence. However, the best papers are the kind that one would find in any collection of papers, namely detailed descriptions of the author's own work, and this fact in my opinion finishes off any claims that *Knowledge engineering* might have to being a textbook. However interesting descriptions of a particular piece of work are to a given writer and those interested in their work, they are hardly relevant to or useful for students who have just begun to study the subject. However, these same papers may well prove interesting to those with more experience of the field, and it is to them that the books are most likely to appeal.

References

Nilsson, NJ, 1991. "Logic and artificial intelligence". *Artificial Intelligence* **47** 31–56.Birnbaum, L, 1991. "Rigor mortis: a response to Nilsson's 'Logic and artificial intelligence'. *Artificial Intelligence* **47** 57–78.

Reviewed by Simon Parsons, Advanced Computation Laboratory, Imperial Cancer Research Fund, London, UK

Learning in embedded systems by Leslie Pack Kaelbling, Bradford Books. MIT Press, USA, 1993, pp 176, $29.95, ISBN 0-262-11174-8.

An *embedded system* is one which "has an ongoing interaction with a dynamic external world" (p. 1—all references are to *Learning in Embedded Systems*). It has perceptual inputs and produces effector outputs, and can thus partly influence its environment. It may have internal states which will allow its *action mapping*, the mapping from inputs to outputs, to depend on any, all, or none of