little to do with logic. The section on "Expert Systems" (page 77) is similarly flawed. The system XCON/R1 was used to configure systems in the manufacturing plant, not help salesmen (that was XSEL). It was Edward Shortliffe who developed MYCIN, not Edward Shortcliff, and it is the Dempster–Shafer theory, not the Dempsey–Shafer theory. Further errors are obvious in the definition of "Genetic algorithm" (page 94), which is almost entirely misleading (and genetic algorithms certainly don't belong in the category "Theory/Philosophy") and that of "Monotonic Reasoning" (page 138) which completely misses the point about non-monotonic reasoning. This is not that contradictions are tolerated (they aren't in most systems), but that contradictions are prevented by withdrawing old incorrect conclusions when new information is added to the database (it is the fact that conclusions cannot be withdrawn that makes classical logic monotonic). Even when the basic ideas are correct, the authors are extremely slapdash in putting the entries together, as the entry for "Meta Rule" (page 134) clearly shows—the example simply does not make sense.

To be fair to the authors, it is extremely difficult to provide anything like a comprehensive overview of artificial intelligence in a single book, especially a slim volume such as this. Even some AI textbooks, written by people with long involvement in the field, occasionally contain factual errors in the description of subfields with which the authors are unfamiliar, and it must be almost impossible for people coming from outside the field to grasp any part of it quickly. Thus it is completely understandable that this book should include as many facts as it does that are just plain wrong. However, just because it is very difficult to write a book such as this well, it is not excusable to write it badly, or to fill it full of incorrect information.

So, if you want to learn about AI, go out and buy one of the many introductory textbooks, several of which are just as clear and informal (for the most part) as *Artificial Intelligence from A to Z*, and which have the definite advantage of being correct.

Reviewed by Simon Parsons, Advanced Computation Laboratory, Imperial Cancer Research Fund, PO Box 123, Lincoln's Inn Fields WC2A 3PX, UK

Intelligent decision support—handbook of applications and advances of the rough sets theory edited by Roman Słowiński, Kluwer Academic, The Netherlands, 1992, pp 471, £88.00, ISBN 0-792-31923-0.

The first thing to note about this book is that the title is partly misleading since the book says little about intelligent decision support. However, as the second part of the title suggests, it has a lot to say about rough sets, being a collection of papers on the subject. Rough sets, as the book seems to mention every other page, were introduced by Pawlak in 1984 as a means of extracting "decision rules" of the form "if a car is small and has a large engine then it is fast" from a mass of data about the size and speed of a car and the size of its engine. They have been quite widely adopted, forming the basis of a number of systems of machine learning and inexact reasoning, and any new work on them is automatically of interest. Unfortunately, this collection of papers, many of which are interesting in their own right, give the impression of having been thrown together rather randomly (perhaps because they have no more connection than having been presented at the same conference), and so, despite the editor's best efforts, sit rather uneasily together. As a result, it is difficult to see who is going to shell out £88.00 for the priviledge of having a copy of this book on their bookshelf. Even I, a confirmed advocate of the use of rough sets, would think twice about it. This is a shame, because in a more attractive (and cheaper) format the book might well go further, and give the idea of rough sets some well deserved publicity.

The book opens with a section of 13 papers on applications of rough sets to a wide range of problems from the control of a rotary clinker kiln to the prediction of earthquakes in Belgium, taking in on the way, the analysis of voting in the 1988 US Presidential election, and that favourite chestnut of the rough set community, the prediction of the prognosis for patients having undergone a highly selective vagotomy. All of these are fine descriptions of applications, except that none of them actually explain what a rough set is, so one has to read to page 205 to find a definition—the

publishers are clearly hoping to encourage the baffled reader to buy Pawlak's monograph on the subject. In each one the problem is sketched out in enough detail, the relevant mass of data presented in a table, and the results of applying Pawlak's method of analysis described. It would be churlish to complain about details such as the occasional lack of discussion in some of the papers (so that my reaction tended to be "so what?" when I got to the end), when the papers are so nicely presented, but there is a slightly more serious problem. That can be summarized by the question "Who is the intended audience for these papers?" The lack of description of the method is likely to make the papers impenetrable for anyone new to rough sets, but anyone who knows a lot about them is unlikely to be very interested in descriptions of applications. The only people likely to be impressed by such work are people who have a problem that they can't solve and might be inclined to adopt rough sets when they see someone solve a similar problem using them, and such people, I feel, will be put off by the lack of explanation. Anyhow, this is not too much of a problem, and as a whole the section is fine.

The second section contains six papers that relate rough sets to other techniques. These are all very interesting and would make a good point of departure for anyone considering developing the formalism. The papers include a typically worthwhile contribution by Dubois and Prade arguing for an eclectic combination of rough and fuzzy sets, as well as an interesting contribution on rough sets and topology, and two comparisons of the rough set method with numerical techniques for data analysis. It is unsurprising to find that one of these uses the data from the highly selective vagotomy patients.

The final section of eight papers, entitled "Further developments" is more of a mish-mash of different things. There is a description of some software for automating the application of Pawlak's method (in fact, this software was used by a number of the authors of the application papers), and a paper which gives an algorithm that seems likely to make this software redundant since it promises much faster solution in many cases. There is a paper describing the application of rough methods to the pipelining of processes on a multiprocessor machine, and a paper extending the use of rough sets to unify knowledge representation and classification. And there is, of course, a paper analysing the highly selective vagotomy data again, this time with the aim of determining which attributes are important in obtaining the "decision rules".

So, overall, a good book but not a great book, but one that deserves more than it is likely to get, which is to be buried in libraries rather than to be bought by individuals. That this will happen is a great shame, in my opinion, but is only too likely given the price of the book. It is possible to photocopy it, at commercial rates, for around one third of its retail price. Now to me this just does not add up. The publishers incur no typesetting costs since all the papers were provided in camera-ready form, so all they have done is to photocopy it and wrap it in a hardcover. This minimal service doesn't seem worth nearly £60, and of course it isn't. Kluwer, in their wisdom, have decided that they only want to sell it to libraries and have priced it accordingly, thus putting it out of reach of people who are genuinely interested in the subject but don't have £88 to spend—yet another example of commercial considerations coming before scientific ones. I wonder what the authors, and particularly the editor who seems to have made a considerable effort to tie the book together, think of the publisher cynically limiting the audience of their work . . . .

Reviewed by Simon Parsons, Advanced Computation Laboratory, Imperial Cancer Research Fund, PO Box 123, Lincoln's Inn Fields WC2A 3PX, UK

Text and context—document processing and storage by Susan Jones, Springer-Verlag, Berlin, 1992, pp 298, DM54, ISBN 0-387-19604-8.

This book was written to support an option course dealing with various aspects of document processing, and is intended to help readers see links between a collection of topics which, although currently fragmented, might in the future present many possibilities for integration. As one might imagine, this is quite a tall order to satisfy since the topics that fall under the general rubric of