problems. Knowledge engineers may also be directed to long established methods of knowledge representation used by information scientists which may be of value in their own work. The modest investment of £5.00 may repay both audiences with some useful insights.

On being a machine. Vol. 1: Formal aspects of artificial intelligence by A. Narayanan 1988, Ellis Horwood/John Wiley, pp 200, £25.00.

Why write and publish a book? We can imagine several reasons for doing so, let aside the more political ones: an author may wish to discuss some new ideas of his or her own; one may wish to discuss old ideas in a new, more coherent or more popular and yet correct way; in particular, when teaching a course one may feel the need to have some material covering its very content.

The purpose of Narayanan's book is to explore the thesis that AI "has not developed significantly as far as its theoretical and philosophical aspects are concerned, because from the very beginning AI has been miscategorized by theoreticians, philosophers, and even AI researchers" (p. 9). The author has split his exposition into two volumes, the first of which (reviewed here) contains some formal and philosophical background. In particular, Turing's Imitation Game is discussed at length (Ch. 1); and there are introductions to concepts of theoretical computer science such as computability, automata theory and formal languages (Ch. 2) and logic (Chs. 3 and 4). Besides being presented formally, these theoretical concepts and their implications are related to the more philosophical questions raised by Turing's paper. The book aims at "philosophers (who) may want an introduction to formal theory" and experts "to remind expert formal theoreticians ... that at the roots of their subjects are very many definitions and assumptions which they may have forgotten about ..." (p. 10).

Unfortunately, the introductory presentation is very often imprecise, misleading or even incorrect (not to mention the incredible number of misprints throughout the book). The beginner in one of the fields covered will, at best, be confused. Take as a deterring example Chapter 3, the introduction to predicate calculus. You will find there, without any explanation, two different definitions of the syntax of predicate calculus; a third definition, with an absolutely unnecessary and unexplained shift in terminology, is given in Chapter 4. On the other hand, the author never presents a complete list of inference rules, rather surprising for a more than 50-page introduction to predicate calculus. Formal systems are defined as usual but then mixed up with logical calculi having a model-theoretic semantics. Properties of specific logical systems, e.g. that a formula is derivable if and only if its negation is unsatisfiable, are presented as general properties of formal systems. Even such basic notions as satisfiability and entailment are mixed up at various places (pp. 106f, p. 110).

And how about the following statement:

> ... Gödel showed that it is not possible for a formal system (of a sufficiently powerful kind) to be both complete and consistent *at the same time*. That does not mean that a formal system is not both complete and consistent; only that if it is required that completeness and consistency be demonstrated simultaneously, this cannot be done. (p. 143)

We have no idea what the author wants to say with these remarks, and what they have to do with Gödel's theorems.

Just as some of the remarks in Narayanan's book are amazing so is what is *not* in it. What does he take AI to be? It is just classical symbolic AI or are subsymbolic issues included? There is no reference to the Physical Symbol Systems Hypothesis! Searle's Chinese Room argument is presented, but there is no reference to the lively discussion it stimulated. The author takes philosophical views (which might simply be described as "mechanism is wrong") but the arguments in favour of them are never discussed in detail; you may either believe it or not. As an example, take the remark "that ascription of mental predicates to computers is logically impossible, simply because the ascription of mental predicates to an entity implies, or presupposes, that the entity in question has consciousness" (p. 191). Is this a dogma or an axiom? We thought this should be one of the questions to be *discussed* in a book *On Being a Machine*!

And what, by the way, *is* a machine? Does the author take machine to be a synonym for Turing Machine? (Why, then, isn't the title of the book *On Being a Turing Machine*?) Or are humans allowed to be instances of machines? Is there a difference of kind or just of definition? The argument from the provability of the unsolvability of the halting problem for humans being "more" than a Turing Machine isn't convincing: didn't the author just choose the wrong alphabet for arguing about proofs? Aren't humans finite in every respect? Aren't they, from this viewpoint, even "*less*" than a Turing Machine? None of these questions is new; all of them seem obvious to us, but they are not discussed.

Maybe all this will be in the second volume which, in fact, is extensively forward referenced. But then the author has done a poor job of splitting the matter. Everything in this book is contained in, for example, *Gödel, Escher, Bach* (which is referenced) which is just more readable, more witty and more detailed. So why not start with the second volume (whatever it will contain) and reference other texts for the basics?

To sum up: whatever Narayanan's reasons may have been to write and publish this book, the first volume, at least, is not convincing. We simply do not see anything sufficiently new or particularly well presented in it.

Progress in machine learning—Proceedings of EWSL 87: Second European working session on learning by I. Bratko and N. Lavrac (eds.), Sigma Press, pp 256, £14.95.

The field of Machine Learning generates some of the most interesting questions in Artificial Intelligence and consequently is a fast-growing field. The European Working Session on Learning (EWSL) is an annual workshop on Machine Learning, linked to a European Initiative in AI and Pattern Recognition. This paperback records the proceedings of the 1987 conference held in Bled, Yugoslavia.

The book consists of 16 papers (split into four sections by subject area) written by Scottish, French and Yugoslavian authors. It is at something of a disadvantage since other more global and comprehensive conference proceedings on Machine Learning were published in the same year. The same areas are covered by Proceedings of the Fourth International Workshop on Machine Learning, and the Knowledge Acquisition section of IJCAI-87, for instance, each consisting of over 40 papers. Its successor, *Proceedings of EWSL-88*, has also already been published and, in my opinion, contains a much better selection of papers. I will not attempt to cover all or even most of the papers in this book but will rather pick out a "representative sample" (a dodge often made by inductive learning researchers in the face of real-world data), in the order in which I actually studied the text. As I do, I will refer to the two major areas in Machine Learning since the mid-eighties: Similarity-based Learning (SBL) and Explanation-based learning (EBL).

The book starts with a small preface but no general introduction, and contains only reproduced, unedited papers. The first section is firmly in the region of SBL, and is called "Learning from Examples in the Presence of Noise". As expected it reported "bump on a log" research containing an introverted mix of papers stemming from Quinlan's ID3, and Michalski's AQ projects. The not unimportant bump in question was the addition of facilities to cope with both noisy data and data posed in an incomplete language.

The first paper, by P. Clark and T. Niblett, is entitled "Induction in noisy domains" and presents an inductive learning algorithm called CN2 which works like Michalski's "AQ" algorithms. The paper starts by outlining approaches to dealing with the induction of classification rules from real-world data, and points out the approach they use—that of building rules which may not be totally consistent but give a good coverage, and are therefore quite predictive. The authors then present their rule description language, which is similar to Michalski's VL1. At this point one spots that their rule variables start with capitals and a top-down ordering of classification rules is used, and the scent of a Prolog implementation permeates from future pages. Their algorithm is interesting, however, in