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

the use of a statistical significance test to prune hypotheses from the search. Well, the other four papers in this section continue in a similar vein ...

I turned avidly to the next section since its title "Meta-Knowledge and Explanation" loosely covers my research area, but was quite disappointed to find the first paper an airy, classification type with the provocative title: "Is AI a Sub-field of computer science—or is AI the science of explanations?" by Yves Kodratoff (the author suggests the latter of course). Fighting the urge to take the bait, I flicked past to more suitable reading. As I did so, Mitchell's infamous "safe-to-stack" EBL example caught my eye. Was it actually in Kodratoff's paper? I've seen it so many times in the last three years I could easily have imagined it.

The next paper looked promising, having an interesting title: "Learning control knowledge in an explanation-based framework" by R. Desimone. Techniques for learning, however, and details of how the author's system fitted into the EBL framework, seemed to be in short supply. Most of the paper was devoted to re-explaining Silver's earlier work on equation solving called "Precondition Analysis", but rather than learning control knowledge, the chief task of this sytem seems to be one of acquiring pre- and post-conditions in terms of predicates on equation states. The extensions to this earlier work, outlined towards the end of the paper, were akin to STRIPS with MACROPS with some operator non-linearity thrown in, but the significance of these improvements to the performance system was hardly explained.

The next section "Studies of Special Mechanisms for Learning" (!?) started with the paper "Inductive generalization: a logic framework" by N. Helft. Another logical framework for induction! Another clear semantic definition of generalization! I was obviously mistaken in thinking that all the theory had been wrapped up in the seventies and early eighties. Having spent enough time reading about the tedious (but noisy) Inductive Learning in Part 1, I decided to imitate Prolog's cut and failed the call to the whiole section.

The final part is subtitled "Learning in Complex Domains" and consists of three papers, two of which I read (I was put off the final one by its very long title). Both were interesting contributions: "Learning of qualitative models" by I. Mozetic, deals with the construction, refinement and operationalization of a deep domain model for an expert system. The second was another view of S. Muggletons "Duce", a program which restructures expert systems by introducing new concepts.

My general conclusion about *Progress in Machine Learning* is that it contains some competent research papers, but the progress made by them is generally minor and not very exciting. In other words, none of them seem likely to be considered milestones in future years. At best, the book would only serve the purposes of Machine Learning researchers; those specifically interested in Learning from Examples in Noisy Domains might find it useful, since this topic formed its largest and most coherent section. Finally, another drawback is that the book's papers are not particularly representative of the whole Machine Learning field.

Professional judgement: A reader in clinical decision making by J. Dowie and A. Elstein (eds.) 1988, Cambridge University Press, pp 565, £13.50 (paperback).

Through such pivotal early expert systems as MYCIN and INTERNIST, the domain of clinical medicine, more than any other, has helped shape the development of knowledge engineering as a discipline. It is therefore interesting to read a book which aims to provide a comprehensive introduction to prescriptive studies of professional judgement in this influential domain.

The book consists of 30 recent papers and chapters culled from the vast interdisciplinary literature on clinical decision-making, with contributors from medical specialists, psychologists, decision analysts, statisticians, economists, lawyers and others. These contributions are preceded by a useful 40-page linking commentary by the editors. The book is divided into four main parts: