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:

Part 1 (three papers) introduces the major theme of uncertainty in clinical judgement. This is shown to pervade almost every aspect of clinical decision-making. The debate between those who believe that this uncertainty should be handled analytically and explicitly, as against those who claim that clinical expertise is essentially intuitive and implicit, is well brought in the papers by Schon and Hamm.

Part 2 (nine papers) focuses on modelling the clinician and the clinical task, and distinguishes two major approaches. "Policy capturing" and related statistical approaches seek to model the relationship between the inputs (cues, symptoms) available to the clinician and the judgement or decision that results. Here the concern is with modelling behaviour—rather than what is going on inside the clinician's head—in the form of a mathematical equation. In contrast, "process tracing" approaches use techniques such as protocol analysis and computer simulation to model the clinician's cognitive processes.

Part 3 (ten papers) introduces decision analytical approaches to enhancing clinical decisions. This is a rational approach to decision-making under uncertainty based on expected utility theory. Several papers are devoted to outlining and illustrating this established prescriptive approach. Others are in the tradition of behavioural decision theory, which is concerned with how the human handling of uncertainty corresponds to (and frequently diverges from) statistical (e.g. Bayesian) expectation.

Part 4 (eight papers) is about the economic, ethical and legal contexts in which clinical decisions are made. The professional context is also considered, with interesting discussions of how medical failures are managed and why doctors are reluctant to disclose their uncertainty to patients.

The book is intended for teachers, students and practitioners in many related fields. However, even allowing for the fact that it is not designed for a knowledge-engineering audience, the inclusion of only a single paper on knowledge-based systems (Fox, 1984), (compared with a dozen or more decision analytic and related formal methods) is a little disappointing. Indeed, the editors appear somewhat sceptical about the whole expert system enterprise:

> Despite the fact that human experts make mistakes, expert systems research, as we have seen, esteems clinical judgement so highly that it finds it worthwhile to design computer programs that will mimic these judgements in so far as it is possible. (p. 18)

This well-produced volume is mot likely to appeal to researchers and practitioners concerned with the handling of uncertainty or medical decision support systems, or to those simply seeking a deepen their understanding of the nature of human expertise.

References

Fox, J, 1984. "Formal and knowledge-based methods in decision technology", *Acta Psychologica* **56**, 303–331.