cannot err" might now reflect on the merits of an approach to technology that draws on paradox as the essential complement to linear logic. The present work makes a good start.

Reference

Funtowicz, SO and Ravetz, JR, 1990. *Uncertainty and Quality in Science for Policy*, Kluwer Academic, Dordrecht.

Reviewed by J. R. Ravetz, Council for Science & Society, 3/4 St Andrews Hill, London EC4V 5BY

Knowledge acquisition for expert systems (2nd Ed.) by A Hart, Kogan Page, London 1989, pp 196, £14.95. ISBN 1 85091 830 9.

Among the advances in the expert systems world since this book first appeared in 1986 has been a growing understanding of why it is that successful deployment of a working system is a rare achievement. We are now more likely to appreciate the complexity and subtlety of human expertise. Systems are more commonly labelled "knowledge-based systems"; not least because few programs can be said to perform at the level of an expert, who displays intuition and judgement on sub-optimal or inadequate evidence, and furthermore learns from experience. Clearly-defined boundaries, within which the knowledge base can be confidently assumed to be complete and correct, are one factor critical to success. No longer cloaked in mystery, KBS technology is increasingly viewed as another branch of information systems. "Integration" and "embedding" have been buzz-words in recent years, indicating that commercial or practical benefits are significantly improved by linking KBSs with other technologies (such as relational databases). This insight has been accompanied by recognition of the essential importance of validation and of maintenance. Another advance, as expert systems came into the fold, was increased attention given to identifying and involving the user in (and even throughout) the project. A certain level of relevant knowledge must always be assumed on the part of the user—so the earlier and more precisely we can define that knowledge, the better. Computers lack common sense. Human interpretation is usually necessary to some degree between the external world and the inputs to the system, and also between advice or decisions conveyed by the system and actions performed by the human user as a result. The context can never be completely defined for the computer.

One perception which has remained largely unchanged, on the other hand, is that knowledge acquisition is difficult: we still hear references made to the "Feigenbaum bottleneck". It certainly seems usually true that where the knowledge engineer is not already familiar with the task or domain to be modelled, the whole process of knowledge acquisition takes more time than implementation. Whether in fact it is also so difficult as to merit its image (among some) of a "black art", depends partly on the objective of the knowledge engineer (knowledgician?) and partly on the nature of the knowledge. If, at one extreme, the aim is to accurately and completely capture the private knowledge of the one person regarded as significantly knowledgeable about the problem, without access to any verified documents or other realiable sources, then knowledge acquisition may be a failure. On the other hand, some (successful) so-called expert systems have been designed as replacements for technical manuals: acquiring the knowledge from such a source may be tedious, but hardly difficult. Most KBS projects lie somewhere between these two paradigms, where the knowledge engineer following "guided curiosity", with an eye (and an ear) for detail and a degree of common sense, should succeed.

Knowledge acquisition must, however, be disciplined and structured: these days, at any rate, there will usually be a fixed budget and a project plan, incorporating deadlines. There is a considerable gap to be bridged between the mass of disparate sources drawn upon by people in solving a problem, and writing a computer program which generates (or advises on) a solution. Often, too, those who are expected to use the KBS have significantly less relevant knowledge and experience than the expert who acts as the primary source for the knowledge base. Ensuring that the system's user interface is comprehensible for the ultimate user is a part of the knowledge engineer's job, often

expected to be more straightforward than is justified. Textbook guidance in knowledge acquisition has been hard to come by, especially at the level of a broad overview of the subject.

Anna Hart's book was indeed the first in this field, so far as the reviewer is aware (excepting perhaps Margaret Welbank's excellent 1983 technical report, available from British Telecom). The publishers of that first edition were quite clear about their market: "Of particular interest to students of computer studies, this book is also suitable for anyone in industry considering the development or use of expert systems". It was a readable and helpful introductory overview, and was deservedly successful. Some academics or established knowledge engineers might have cavilled, about perceived superficiality for example, but the book was not written for them. Within 180 pages, ten chapters (and an informative preface) covered the nature of expertise, programs as experts, a comparison with systems analysis, the knowledge engineer, fact-finding by interviews, reasoning and probability theory, fuzziness in reasoning, machine induction, the repertory grid (a particular technique) and two case studies. A bibliography, a glossary and an index were appended.

In a fast-moving field it is hardly necessary to justify revising an overview after three years. We may assume that readers of the first edition will mostly have moved on, and that the market of this new edition remains the same: students and those in business who are new to the subject. This is an assumption, because the publishers are now less clear on whom they perceive the "users" of this book to be: their blurb refers merely to the first edition having been "especially welcomed by non-specialists concerned with the practicalities of designing and building expert systems for use in business and industry". The author states in her preface that while written "for those who were relatively new to the field", the original book in fact "has appealed to a different audience from that intended": viz., those who wanted to find out "What's so hard about knowledge acquisition".

The main question is, does this new edition adequately incorporate in line with the overall level and depth of presentation) the advances of the intervening three years, as well as improving the original material? Broadly speaking, the answer is yes. The strongest new theme is one of more attention to practical, human-centred isssues, giving the book a more mature and relistic feel. The first two chapters (titles remain unaltered) introduce the notion of implicit (i.e. tacit or cultural) knowledge in a section headed "true expertise"; expand the meaning of "explanation"; and emphasise that the user and the computer should be seen (in effect) as a synergistic combination. One recognized consequence of an alternative "culture of dependence" approach is deskilling. The third chapter compares systems analysis with "knowledge elicitation" (introduced on page 32 without explanation about its relationship to knowledge acquisition: one has to return to the preface of the first edition to discover that the author does "not make a clear distinction" between the two terms. This reviewer considers it helpful to do so, both preferring and treating as a broader term "knowledge acquisition": it includes the conceptual analysis and modelling which follows "elicitation" of the knowledge). The topic of identification of users and their involvement in the project receives new discussion here, although unfortunately no entry is made for pp 36–8 against the index heading of "users". There is a good two pages on prototyping, followed by half a page of KADS (all of which is new to this edition); paper models attract only a brief mention.

Chapter four ("The knowledge engineer") has gained discussion of the need to know who the users are (by name), as well as experts, and to seek their views and expectations of the proposed system. A new (third) reason why experts should not be their own knowledge engineers is adduced: "they will have very poor perceptions of the users' needs, and sometimes little empathy for them" (p. 44). Anna Hart notes here that a team approach may sometimes be preferable to expecting one knowledge engineer to have all the desirable talents.

The chapter on interviewing has been considerably revised. New topics include: preparation and cosideration of which methods to use, different types of knowledge, intermediate representations, use of a prototype, validation (briefly: no entry in the index), structured versus unstructured interviews, simulation techniques (described as "imposing constraints"), and the twenty questions, laddered grid, and card sorting techniques.

Conversely, very few changes have been made to the chapters on probability and fuzziness: a paragraph on the relevance and use of probability by knowledge engineers, a note on the difference

between "uncertainty" and "imprecision", and a more helpful (though at one page, still brief) discussion of analogical reasoning and pattern matching.

Chapter eight on machine induction has been significantly improved by the inclusion of more warnings and qualifications concerning the use of induction. (Indeed, it is generally true of the book as a whole that many statements have been rewritten in less categorical terms, which is to be welcomed.) An important correction on page 125 is the description of induced results as a possible representation of "domain concepts" (cf, "his [the expert's] rules" in the first edition!). New points stressed include the need for human skills in analysing induction results, with an illustration of how two apparently valid decision trees may be produced from a single training set; the fundamental importance of human judgement in selecting good examples for the training set and assigning appropriate attributes; and the use of an independent set of data for validation. Genetic algorithms receive a brief mention, followed by a page and a half on neural networks.

The only change (apart from replacing one citation by a more recent one) to the chapter on repertory grids, involves a brief nod towards those tools (such as Nextra) being developed which are based on this technique. The final chapter includes a new, third, case study, which really only illustrates the point made earlier about adopting the team approach.

In general, then, the book is still a good introductory overview of what knowledge acquisition might involve. Some of the original preface's cautionary material about the book's limited depth have been jettisoned, however: this is a demerit of an otherwise appropriate new preface. "Students" are (apparently) no longer part of the target "user group", as referenced earlier, which perhaps makes it the more important that the book provides adequate guidance on further reading. Noteworthy here is the annotated bibliography: now extended from 35 to 61 references, organised more helpfully by chapter this time, rather than by general subject area, it includes many post-1985 citations. Corrections have been made too, for example the introduction of the standard Ericsson and Simon (1984) text on protocol analysis. Still missing, on the other hand, is citation of the original Michalsky and Chilausky paper on induction: frequently cited as an induction success story, but one wonders how many commentators have actually read the paper's inadequate reporting of the questionable "manual" knowledge acquisition, whose results were unfavourably compared with those from induction! No references to Mildred Shaw's work on repertory grids subsequent to 1981 have been added, which is doubly unfortunate in view of her attention to incorporating views of more than one expert (not an uncommon practice, except where the aim is to specifically model one particular person's knowledge) and to validation and testing.

This leads to two other general points of criticism. First, Anna Hart does not provide sufficient guidance on either the relative merits or appropriateness of different techniques. Probably the single question most commonly asked by the likely readers of this book is "What technique do *you* use in knowledge acquisition?" (or "What's the best way?"). In response, after one has muttered something about horses for courses, and many factors being involved, some generally helpful remarks can—and should—be made. For example, protocol analysis (rather dismissively treated here in less half a page) is a good technique for accessing procedural knowledge—which is often more difficult to acquire to acquire than the declarative type— and the order in which information is sought and used in solving a problem; as also, it might be said, are "twenty questions" and simulation techniques (the reviewer speaks from experience). Conversely, techniques like the laddered grid, card sorting and repertory grid are more suited to finding out how the expert structures a domain. Repertory grids are also much better suited to decision-making in "soft" problem areas such as education and management, from where published illustrations are usually drawn, than in areas where conceptual distinctions are more clear-cut. In the latter, the technique can also appear contrived and unnecessarily complex to domain experts, whose time is generally at a premium.

The second point concerns important issues like validation, verification and the relationship of knowledge acquisition to the other stages in KBS development. Their treatment really is not adequate: in even a fairly short book (196 pages—170 of which are text—with no more than 400 or so words per page, at most), more than a brief passing mention is required today. That is to say that although this book is a considerable step in the right direction, it needs to go further. If—and this

must be doubtful—economics required deletion of some material to make way for these extensions, the author might reconsider whether it really is appropriate, at this level, to devote 38 pages to the general topic of Uncertainty. These two chapters are certainly the least relevant to the practice of knowledge acquisition, at the level readers of this book are likely to experience. The index is generally good and better than is usually provided at this level, although it should include entries like "lifecycle" and "validation". It would also be helpful if the reader's attention was drawn (at an early stage, such as the preface) to the glossary at the back: terms pop up in the text without explication, which in fact the glossary often provides.

The standard of production of the book is good, without typographical errors, although the publisher's decision to set the price as high as £14.95 for a paperback edition is rather disappointing. This seems to confirm the impression that students are not expected to buy the book, which is a pity because there is a growing potential readership here; certainly at, say, £9.95, this reviewer would have recommended it for purchase by students in his department (rather than forming a queue in the library).

One final point relates to the origin of this second edition. In her preface, Anna Hart says that she expected a succession of ever-more-detailed books on knowledge acquisition would have appeared since 1986, making hers "more or less redundant", but that this has not happened. She is wrong: commencing with the book edited by Alison Kidd published (after two years in gestation) in mid-1987, followed by Mike Greenwell's a year later (neither, unfortunately, cited by Anna Hart), a stream is now in flow, particularly since the beginning of 1989 when two very relevant books edited by Dan Diaper and by Guida and Tasso appeared. (And besides this journal, the 1989 newcomer *Knowledge Acquisition* should be mentioned). On the other hand, neither do any of these yet make Anna Hart's book redundant: none match hers as a very clear, accessible introduction to the subject. Experienced knowledge engineers may not learn much themselves—except about writing—from this book, but many will be grateful to be able to point the enquirer towards it. A third edition should be considered very soon.

References

Diaper, D (Ed.), 1989. *Knowledge Elicitation: Principles, Techniques and Applications*, Chichester: Ellis Horwood.

Ericsson, KA and Simone, HA, 1984. *Protocol Analysis: Verbal Reports as Data*, Cambridge MA: The MIT Press.

Greenwell, M, 1988. *Knowledge Engineering for Expert Systems*, Chichester: Ellis Horwood.

Guida, G and Tasso, C. (Eds), 1989. *Topics in Expert System Design: Methodologies and Tools*, Amsterdam: Elsevier.

Kidd, AL (Ed.), 1987. *Knowledge Acquisition for Expert Systems: A Practical Handbook*, New York, NY: Plenum Press.

Knowledge Acquisition, published quarterly by Academic Press (started March 1989).

Michalski, RS and Chilausky, RL, 1980. "Knowledge acquisition by encoding expert rules versus computer induction from examples: a case study involving soybean pathology" *International Journal of Man-Machine Studies* **12**(1) 63–87.

Welbank, M, 1983. *A Review of Knowledge Acquisition Techniques for Expert Systems*, Ipswich: Martlesham Consultancy Services.

Reviewed by Ian Neale, ARIES at City, Department of Business Systems Analysis, The City University, London EC1V 0MB.