

First generation expert systems: a review of knowledge acquisition methodologies

IAN M. NEALE

ARIES at City, Room CM355, Centre for Insurance and Investment Studies, The City University, Northampton Square, London EC1V 0HB

Abstract

This paper reviews a wide range of knowledge acquisition techniques in the context of attempts to achieve a systematic methodology. These have been poorly documented by expert system builders, who are often inclined to overvalue textbooks and the ways experts themselves claim they solve problems. No one method has a universal advantage; each has some value. Techniques should be selected to suit the domain, the task, the expert and the knowledge engineer. Knowledge acquisition involves creating a conceptual model of expert knowledge and reasoning, from analysis of data elicited by these techniques. A survey of the literature indicates increasing emphasis on tools for knowledge acquisition, used directly by experts. Several projects currently directed towards providing a proper epistemological foundation for knowledge acquisition are discussed and compared. None has yet produced a complete epistemologically sound methodology; however, recognition of the need to create a conceptual model at the knowledge level (rather than the symbol level) is an important advance. An extensive bibliography is appended.

1 Introduction

1.1 Orientation

The approach taken by some enthusiastic promoters of expert systems suggests that all you need to make a program intelligent is to provide some general rules and lots of very specific knowledge. This has created a mechanistic conception implicit in the very name “knowledge engineering” coined to describe the activity of building expert systems by Feigenbaum; see Buchanan and Duda (1983) for an early description of the task. For example, in a book widely read in lay circles, Feigenbaum and McCorduck (1984) describe the job of knowledge engineers as to “mine those jewels of knowledge out of [experts’] heads one by one” (p. 107). Stamper (1988) suggests that many AI researchers treat knowledge as if it were oil to be extracted, suitably refined by knowledge engineers and then sold as a profitable commodity. Systems Designers Scientific (1986) speak of “brain scouring”, and Currier (1986) reportedly describes the use of an expert system development package as “knowledge crunching”. Describing the same package (“Super Expert”), Michie (1986a) is quoted reaching new heights of hyperbole in asserting that “it’s joined the rest of the human race”.

Whether or not it is entirely fair to construe these remarks as evidence of an over simplistic analysis of expertise, it is as well to note that “major benefits accrue to the cautious users of the best human expertise: significant risks are run by those who put much faith in a mechanistic diagnosis of a complex problem in nature” (Hawkins, 1983, p. 1). Knowledge is not a substance to be quarried (Young, 1988); rather it is something to be probed by knowledge engineers in an essentially social process (Goodall, 1985).

Fewer psychological and practical difficulties will be encountered if knowledge acquisition is considered as social science research, aiming to *model* the expert’s knowledge rather than merely *extracting data* (Breuker et al., 1987; Johnson and Johnson, 1987b; Kidd, 1987). With Evans (1987),

it is argued that the concepts and methods of cognitive psychology have often been unjustifiably ignored in the elicitation of human knowledge.

1.b Definition

It has been frequently stated that the problem of knowledge acquisition is "the critical bottleneck" (e.g. Feigenbaum and McCorduck, 1984 p. 107) in expert system development. The frustration which this can engender in knowledge engineers has contributed to the misguided belief, of some, that there is "one magical technique" (Kidd, 1987). It is more helpful to develop a "toolkit" of techniques, from which to select according to the nature of the domain and type of knowledge under investigation (Gammack and Young, 1985; Berry and Broadbent, 1986; Banathy, 1988); remembering that "nothing is so detrimental to the acquisition of insight than an underestimation of the complexity of the knowledge and reasoning in a particular domain" (Breuker and Wielinga, 1983 p. 6).

Knowledge acquisition involves elicitation of data from the expert (usually by some verbal technique); interpretation of the data to infer the underlying knowledge or reasoning process, and, guided by this interpretation, creation of a model of the expert's domain knowledge and performance (Kidd, 1987). These three phases have been similarly identified elsewhere in the literature as domain definition, fundamental knowledge formulation, and basal knowledge consolidation (Grover, 1983), or orientation, problem identification and problem analysis (Breuker and Wielinga, 1987a). The common thrust is towards the development of a methodology of knowledge acquisition. With particular reference to the first phase the terms "knowledge elicitation" and "knowledge acquisition" are sometimes loosely regarded as synonyms in the literature, or alternatively the former as a subset of the latter (Breuker and Wielinga, 1983). As Jackson (1985) observes, 'elicitation' carries overtones of extraction and coercion; whereas 'acquisition' suggests a piecemeal process and connotes intrinsic value, and is therefore to be preferred. The original usage of cited authors has been retained, however.

2 Overview

The paper discusses the variety of techniques and tools employed by practitioners of differing methodologies. I have deliberately eschewed detailed consideration of issues of knowledge representation, the nature of expertise and theories of human information processing (see Slatter (1987) for a relevant discussion) and machine learning apart from induction. A comprehensive treatment of knowledge engineering would certainly require inclusion of these topics, but the result would be a book and not a review article. Dangerous superficiality can result from attempting too much in a paper (e.g. see Garg-Janardan and Salvendy, 1988).

A summary review of the literature is followed by an outline of the essential elements of knowledge acquisition (as defined above). First, the questions which should be considered during the initial domain orientation phase of an expert system project are discussed. The review of psychological techniques of knowledge acquisition which follows covers a range of interview strategies, observational methods, some multidimensional techniques (card sorting, MDS, repertory grid and proximity analysis), and verbal protocol analysis. The next section discusses machine-aided knowledge acquisition and in particular reviews the literature on some of the tools which have been developed. The bibliography covers the period up to July 1988.

3 Literature on knowledge acquisition

Knowledge acquisition has been poorly described (Hart, 1986). Few papers have discussed what methods work for which domains or types of knowledge (Welbank, 1987). One typically vague report quoted by Wright and Ayton (1987, p. 14) is from Pauker et al. (1976), who elicited medical expertise by "introspection and through direct observation of the clinician's problem-solving behaviour. The insights gained in this way were represented in a computer program" (p. 983). As the

use of expert system shells becomes more commonplace it seems likely, regrettably, that even less attention will be paid by researchers to recording their methodology. For example, Perkin, Pitt and Price (1987) describe their strategy for knowledge acquisition for industrial real-time expert systems simply as "to acquire expertise from the body of academic research" (p. 233).

The principal overviews of the literature have for several years been the work of Welbank (1983) and of Breuker and Wielinga (1983, 1984; see also Wielinga and Breuker, 1985). A concise up-to-date account of techniques and tools, with many references, has been presented recently by Boose (1988). A general introduction to the subject, particularly strong on the topics of uncertainty and induction, is provided by Hart (1986); see also briefer overviews by the same author (Hart, 1985c; 1988). Kidd (1987) has edited a selection of more detailed and rigorous accounts of methodologies, preceded by her own introductory framework. A variety of approaches to knowledge elicitation are discussed in Diaper (1989), and another selection is reviewed in context by Slatter (1987); a shorter collection of papers appears in Pavelin and Wilson (1987).

For the past two years the state of the art has been formulated at a series of annual Workshops on both sides of the Atlantic. Reading University (UK) have published the Proceedings of the first European Workshop on Knowledge Acquisition for Knowledge-Based Systems (Department of Computer Science, September 1987). Papers presented a year earlier at the first North American Workshop have appeared in five subsequent special issues of the *International Journal of Man Machine Studies*, over volumes 26 and 27 (see especially Gaines (1987b)).

Proceedings of the second North American (AAAI sponsored) Workshop (Banff, Canada, October 1987) are likely to appear in the same journal in late 1988 and/or 1989. During the final stages of revision of this paper the author received the proceedings of the second European Workshop (Bonn, June 19–23, 1988): Boose, Gaines and Linster (1988); although a thorough consideration was precluded, a number of citations from this have been included. The Third AAAI sponsored workshop on Knowledge Acquisition for Knowledge-Based Systems will again take place in Banff, from 7–11 November, 1988; the proceedings are expected to appear in the *International Journal of Man Machine Studies* and perhaps also in an Academic Press book.

Burton and Shadbolt (1987a) offer a short practical introduction to some techniques of knowledge acquisition—reviews of similar scale may be found in Davies (1986), Wright and Ayton (1987) and Hoffman (1987). Olson and Rueter (1987) outline twelve methods of knowledge acquisition in more detail. Several psychological methods are discussed by Evans (1988) in the context of an extensive review of research in cognitive psychology. It is considered in a sketchier and maxim-orientated way by Buchanan et al. (1983), Waterman (1986) and Prerau (1987).

4 Knowledge acquisition methodologies

The first section covers the basic points for a knowledge engineer to consider in the *domain orientation* phase. These are summarized in Figures 1–4. There follows a discussion of techniques which have been employed to discover what knowledge the expert uses and the methods by which it is used to solve problems within the particular domain. The analysis and structuring of this knowledge should lead to a model of the domain concepts, relations and inference strategies, putatively corresponding to the expert's view of how such problems are solved. The construction of this conceptual model is the least mature stage in manual knowledge engineering methodology. For first generation expert systems it has often involved little more than the extraction of rules from the data, although the use of paper modelling techniques, drawing upon established systems analysis practice in many cases, is gaining acceptance.

4.a Initial considerations

Prior to commencing work with the expert the knowledge engineer must be oriented with respect to the domain. This involves listing and investigating sources of information, establishing familiarity




Figure 1 Necessary requirements for expert system development (Waterman, 1986 p. 129).




Figure 2 Justification for expert system development (Waterman, 1986 p. 130).




Figure 3 Characteristics which make the use of expert systems appropriate (Waterman, 1986 p. 132).

with the basic vocabulary, identifying authoritative experts, and assessing the feasibility of the project. The knowledge engineer should as a result be clear about the following points at least.

- 1 *Is expert system development possible?* (Fig. 1)
- 2 *If it is possible, is it justified in the particular circumstances?* (Fig. 2)
- 3 *Does the problem show characteristics which make expert system development appropriate?* (Fig. 3)
- 4 *Who are the prospective users of the expert system?* User definition is frequently overlooked (or considered in the final stages) in developing interactive systems; it should also include an analysis of the purposes for which the system will be used, and what the users already know about the domain and the way they use their knowledge (see also Collins, Green and Draper (1985) and Berry and

Broadbent (1987)). Particular care must be taken where users are not just lower down the same ladder of expertise, but in fact are trained for a different role; an example is a nurse who uses the doctor's expertise in the ACORN system to assess the significance of chest pains (Wyatt and Emerson, 1988). Josefson (1987, p. 175) quotes a similar instance. Mittal and Dym (1985) stress the paramount importance of matching the needs of the intended users: systems like MYCIN and MDX have never been deployed in clinics because they only perform tasks which doctors perform routinely anyway, and probably more slowly. Failure to define user requirements and involve users in system design from the start has been a significant cause of project failures (Candy and Lunn, 1988). Sparck Jones (1985) notes the ambiguity in the term "user" where a human intermediary is involved in an interactive system, highlighting the importance of user definition. Diaper (1987, 1989a) advocates extending user modelling to include those indirectly affected by the system, by developing a model of the client organization.

A point worthy of note in designing an expert system whose projected role is decision support is the possibility of using the technology to complement (rather than copy) human reasoning ability (Evans, 1987). Certain inherent limitations of human information processing have been noted, for example by Tversky and Kahneman (1974), Barr (1983), Jacob, Gaultney and Salvendy (1986) and Lehner and Adelman (1987); perhaps we should conceive of the human and the expert system as a joint cognitive system (Woods, 1986). The work reported by Fenn et al. (1986) on the development of a cooperative expert system is related to this point, and is supported by Zimolong et al. (1987).

5 Who are the experts to work with? In some complex domains where an individual's knowledge is restricted to a small subset it may be advantageous to collaborate with several experts (Mittal and Dym, 1985). On the other hand, it may be difficult to achieve consensus if their areas of expertise overlap. 'Getting a second opinion' could ultimately reduce the quality of the knowledge base if a compromise has to be made (Adelman, 1987).

It is of course, of critical importance to secure and maintain the willing cooperation of an articulate expert. This is no mean feat, as Welbank (1983) observes. If successful, various ways of maintaining the expert's interest have been suggested by Welbank and others, including appealing to his or her sense of intellectual challenge; avoiding the imposition of excessive structure and criticism, or any suggestion that the expert is failing to inform appropriately; insulating the expert from technical matters such as the intricacies of programming languages; getting something working quickly (although popular, "rapid prototyping", as it is called, has been severely criticized; e.g. by Breuker and Wielinga (1983, 1987a)) or at least frequently referring the emerging system to the expert for comments; above all, recognizing that experts' time is valuable and that they should be treated with respect. Tact, diplomacy, empathy and patience should be the watchwords of the knowledge engineer (Buchanan et al., 1983; Welbank, 1983; Hart, 1986). The obligations the expert feels towards the knowledge engineer and the project, and the extent of any benefits he or she perceives for himself or herself and/or his or her employers, will significantly influence the viability and success of the whole exercise.

6 What are the boundaries of the domain? These should be constrained as narrowly as possible to make the domain manageable, whilst encompassing a sufficient breadth of expertise to be useful (Hartley, 1981; Waterman, 1986).

7 What is the language of the domain? Many authors observe that a fundamental prerequisite for effective knowledge acquisition is a degree of mastery over the vocabulary of the domain (e.g. Breuker and Wielinga, 1984; Kahn, Nowlan and McDermott, 1985b), at about the level of the target user of the expert system (Berry and Broadbent, 1986; Jackson, 1986a). The expert and the user may not share the same vocabulary, in which case the knowledge engineer must be prepared to learn both. To be unaware of basic concepts and terminology is to risk inhibiting the expert, or even alienating him or her by requiring him/her to spend valuable time on explanations (Welbank, 1983). In interacting with the expert, the knowledge engineer should use the same terminology as the expert to avoid errors introduced by the expert translating questions into more familiar terms (Bainbridge, 1979); this may not always coincide with textbook terminology. One should not, however, fall into

1. *Is expert system development possible?*
2. *If it is possible, is it justified in the particular circumstances?*
3. *Does the problem show characteristics which make expert system development appropriate?*
4. *Who are the prospective users of the expert system?*
5. *Who are the experts to work with?*
6. *What are the boundaries of the domain?*
7. *What is the language of the domain?*
8. *Are there any relevant textbooks?*
9. *What hardware and software are available for use?*
10. *Has any similar research been reported in the literature?*

Figure 4 Ten questions to begin with in knowledge acquisition.

the trap of deluding oneself into feeling expert in a particular domain: that takes years of training and experience (Buchanan et al., 1983). Some workers assert that it is not even necessary for a knowledge engineer to become a "semi-expert" in advance of developing an expert system (Jackson, pers. comm.; Burton and Shadbolt, 1987a). A short period of apprenticeship is however strongly advocated by Collins, Green and Draper (1985): see 4.b.2. *Participant observation* below.

The knowledge engineer should also discover whether the language of the domain is formal (as in mathematics), or at least stable and well agreed (as in medicine). If it is neither, as in classroom teaching (Wood, 1987), knowledge acquisition will involve the creation of a "new" language to support reasoning, making it a more difficult task.

8 *Are there any relevant textbooks?* Reading them is often the quickest route to a general understanding of the domain (e.g. Roycroft and Loucopoulos, 1985), although formal textbook models of expertise can vary considerably from the way experts practise. The problem-solving and knowledge-structuring skills essential to an expert are developed from this idealized knowledge of textbooks in the process of learning to become an expert (Bobrow, Mittal and Stefik, 1986).

9 *What hardware and software are available for use in conjunction with knowledge acquisition?* It may be helpful to have access to an expert system shell, although many authors warn against preselection of a particular implementation as it may impose an inappropriate constraint. On the other hand, it may be that the computer already exists and perhaps software too, in a "fixed option" form, in which case it is as well to know from the start! (Such a situation was presented to Coopers and Lybrand by the UK Department of Health and Social Security when they began the project reported by Evans and others (British Computer Society, Expert Systems '86 Conference).

10 *Has any similar research been reported in the literature?* In addition to a retrospective literature search, consideration should be given to contacting anyone who might be currently interested in the field.

The goal of domain definition should be the production of something like a Domain Definition Handbook (Grover, 1983), which might contain:

- A general description of the problem;
- Bibliography of principal references;
- Glossary of terminology;
- Identification of experts;
- Characterization of users;
- Definition of suitable measures of performance;
- Description of example reasoning scenarios.

Where possible, the knowledge engineer should aim to eventually formulate a “knowledge-level” description (Newell, 1982) of the selected problem solving task; i.e. the types of knowledge necessary to solve that problem. This will include the goals of the task, the actions of which the problem solver is capable, and the knowledge of which actions lead to which goals. (In Newell’s description these are all connected by a *principle of rationality* i.e. that “if an agent has knowledge that one of its actions will lead to one of its goals, then the agent will select that action” (p. 102)). The purpose is to provide a clear specification of the minimal requirements for the expert system, preferably in the form of a conceptual model of expert problem solving, independently of the facilities of any expert system tool (see also Jackson, 1986b) and therefore distinct from the symbol level at which the knowledge is encoded. The function is to bridge the gap between verbal data on expertise and implementation, whilst avoiding the constraints implicit in the choice of any particular implementation formalism. In first generation expert systems this gap is commonly ignored and the behaviour of the intelligent problem solver displayed at the symbol level in the form of rules, frames, procedures, etc. Modelling the behaviour instead at the knowledge level offers flexibility and predictability at a level above the confines of representation.

In practice, however, an expert is rarely carrying out a single problem solving task in isolation. The situation is often more complex: a human expert may fulfil several functional roles (communication tasks or “modalities”), executing different problem solving tasks within each (see Basden (1984) for a discussion of roles). An example is a hospital doctor advising patients in his or her role as consultant and giving explanation to students in his/her role as tutor. The distinction between modalities and problem solving tasks—which Breuker and Wielinga (1987a) note has hardly been acknowledged—is important, because expert systems sometimes cannot be switched from one modality to another without completely reconstructing the knowledge base. This is illustrated by Clancey’s attempts to convert MYCIN from a purely diagnostic system into a tutoring system (Clancey and Letsinger, 1981; Clancey, 1983; see also discussion in Coombs and Alty, 1984). Therefore, the modality for the proposed expert system should be decided at an early stage; then the set of problem solving tasks required by the system can be identified.

In approaching knowledge acquisition through any of the various techniques discussed below, the knowledge engineer should give a high priority to documentation of data and interpretations. The process should be foreseen as essentially incremental, operating in a breadth-first and cyclical manner to favour successive refinement and inclusion of the results of analysis of related topics (Breuker and Wielinga, 1987a). Acquisition of new data should alternate with analysis of data previously obtained. Documentation in the form of a “paper model” in plain English is a useful form of feedback for expert review (Trimble and Cooper, 1987; Welbank, 1987; Inder, Wai and Fraser, 1987). Diagrammatic or graphical representations, such as decision trees, are often more comprehensible to the expert than textual statements or a body of rules (see Alvey and Greaves (1987) for an example of the use of decision trees in building an expert system).

4.b Psychological techniques

4.b.1 Interviews

Interviewing is the most commonly used way of acquiring basic concepts and terminology as used by the expert (particularly if printed sources have been unavailable), being more efficient than obser-

vational methods. Interviews can be structured in various ways; the important thing to remember with all interviewing strategies is that there may be discrepancies between the way in which the expert presents his or her knowledge and the way it is actually used by the expert in practice (e.g. Welbank, 1983; see also the review in Gaines, 1987b). His/her judgement may be biased: detecting, avoiding and correcting errors of judgement is an important responsibility of the interviewer which is often overlooked (Cleaves, 1987). On the other hand, the interviewer may introduce bias and error through the way questions are phrased (Cooke and McDonald, 1986; Jackson, 1986c; Kessel, 1986; Silva and Regan, 1986; LaFrance, 1987). The field of subjective assessment of probabilities (and numerical representations of uncertainty in general) is a particular minefield, considered by Wright and Ayton (1987) and many others including Tversky and Kahneman (1974), Fox (1986) and Saffioti (1987); a detailed review of this problem is outside the scope of this paper, however. Breuker and Wielinga (1983) define the function of interviewing as "to obtain information that is private . . . not commonly shared . . . and that remains in general tacit"; they also warn that "not all private knowledge is inspectable, which means . . . that there is an unavoidable source of underreporting" (p. 11). Some practical advice on interviewing is offered by Davies and Hakiel (1988).

Tutorial interview

The simplest and least "interventionist" format: the expert is asked to prepare an introductory talk, outlining the main themes and ideas of the domain, and lasting about an hour (Gammack and Young, 1985; Gammack, 1987c). This is recorded, and from a transcript concepts are subsequently extracted. The method should be considered if, for example, a lack of textbooks or other reference works has prevented the knowledge engineer from achieving satisfactory domain orientation. It has the additional merit that a dictionary definition may fail to adequately represent the words used to describe a concept: as Gammack (1987c) notes, much human expertise is likely to consist in the personal and semantic associations (connotative meaning) that an expert brings to domain concepts and may result in the invention or appropriation of personalized terms to describe esoteric or subtle domain phenomena. A variant of this technique is the "compiled list" method (Fenn, pers. comm.) in which the expert is asked to write out specific areas of knowledge as a flat list of information; this is best suited to well-defined areas such as disorders/stereotypes, where complete coverage of the domain may reasonably be anticipated.

Focused interview

The interviewer prepares the topics in advance, although not the precise questions: this technique is most like normal conversation and therefore expert cooperation is most easily obtained. The structure, outlined by Breuker and Wielinga (1984), consists of three parts.

First, an introduction, to explain the goals, orient the expert to the role of interviewee, and agree the sequence of topics; thus the expert is motivated and enabled to plan strategies for responding. Second, topics are addressed sequentially from a list, often in an associative manner determined by the interviewer, and various probing interventions employed to elicit further information about each topic. Finally, the interview is closed with a summary and evaluation of the information obtained. The interview should be tape recorded. A focused interview should provide factual knowledge, the types of problems the expert solves, and the functions the expertise performs.

Burton and Shadbolt (1987a) outline a similar method centring on the use of probes, such as "why/how/when would you do that?", "what alternatives are there?", "can you tell me more about that?", etc. Their "formal interview", however, is more demanding of the expert than the approach of Breuker and Wielinga, expecting him/her in the course of a ten minute outline of the target task to provide a list of major rules which connect variables to possible solutions of the problem.

It should be remembered that some knowledge is not suitable for representation in the form of rules (Gammack and Young, 1985) and that experts do not normally speak in rules (Cooke and McDonald, 1986; Wielinga and Breuker, 1986). A few variant strategies deriving from the focused

interview technique, and not requiring the expert to be explicit about the rules he or she thinks he/she might use, have been reported in the literature.

Distinction of goals

The expert is presented with a specific goal (e.g. a diagnosis) and asked what evidence is necessary and sufficient to distinguish this goal from the alternatives. By doing this first for the final goal and then for each intermediate goal, a model of the expert's knowledge can be constructed. This technique is appropriate for medical and other diagnostic domains (Welbank, 1983; Hart, 1986).

Reclassification

The expert is presented with the evidence (facts), and asked which goals are supported by which facts. Each goal is then reclassified into the set of facts supporting it. These facts are then reclassified into subgoals, in a 'backward chaining' manner, and so it continues until the evidence has been broken down into facts directly observable by the expert (Grover, 1983; Hart, 1986). The medical consultation described by Keravnou and Johnson (1986, pp. 36–38) seems to incorporate this approach.

Dividing the domain

This is the reverse of reclassification: the expert starts with a set of facts (e.g. symptoms) and "forward chains" through successive subgoals (e.g. diagnoses) until the final goal is reached (Hart, 1986).

Systematic symptom-to-fault links

Lists of all possible faults in a system and all possible symptoms are presented to the expert, who is asked which faults would lead to which symptoms. Welbank (1983) reports that one electronics fault diagnosis knowledge base was produced in this way, and notes that it may be easier for experts to produce such lists of objects than sets of rules. It may be too artificial for other than simple restricted domains, however (Breuker and Wielinga, 1984); not least because it disregards the order in which symptoms are encountered in a real situation (Hart, 1986).

Intermediate reasoning steps

A more detailed representation of expert diagnostic reasoning may be obtained by asking the expert for the evidence at each stage which points to the next step in the symptom-to-fault link. Intermediate steps could involve reasoning "from initial symptoms, to affected part, to type of infection, to diagnosis, to treatment" (Welbank, 1983, p. 20).

Structured interview

There is no sharp distinction between some of the above strategies and the structured interview; the principal difference lies in the amount of questioning performed by the knowledge engineer. The structured interview involves the detailed, depth-first sequencing of topics in an attempt to elicit all knowledge related to a particular concept or model. The interviewer requires a good knowledge of the terminology of the domain and the ability to adapt his/her strategy of questioning according to his/her evaluation of an answer. Care must be taken to present the questions so that they access the expert's knowledge in the most natural way (Bainbridge, 1979). Although this technique may reveal the detailed structure of a concept better than other methods, there is a danger that the knowledge elicited is an artefact created under the pressure of questioning (Breuker and Wielinga, 1984). The

structured interview is a relatively poor way to elicit problem solving strategies, if-then rules, and certainty factors (Welbank, 1987).

Systems Designers Scientific (1986) used this technique in the early stages of their project to acquire knowledge about a chemical reaction. It was also the chosen method for work in the insurance domain under the British ALVEY research program by the ARIES Club (Butler and Chamberlin, 1988).

Twenty questions

This is described by Breuker and Wielinga (1984) as a particular strategy of structured interviewing in which questions are designed to be answered “yes” or “no”. The aim is to collect in a fast way the major characteristics of concepts in a domain. Burton and Shadbolt (1987a), however, appear to interpret the technique differently (and more appropriately): in their account, the knowledge engineer has a set of examples of solved problems in the chosen domain and the *expert* asks the questions. The preparation of examples and answers to all anticipated questions is very time consuming (Schweickert et al., 1987). A similar procedure is described by Spradley (1979), by Welbank (1983) and also by Wright and Ayton (1987) under the heading *Context focusing: short-cut protocol analysis*. From this, one can discover what information the expert needs to solve the problem, and to some extent also the structure of the problem solving, from the order of questioning. Gammack (1987a,b) reports using “twenty questions” successfully to validate the structure of a knowledge base derived from other techniques. Sorensen and Nordhus (1988) used a similar technique. Welbank (1983) counsels *asking* the expert to explain why he asks for the information he does, rather than attempting to work it out from transcripts.

Laddered grid

This technique is suitable where the knowledge engineer strongly suspects there is a hierarchical structure to the domain, and it quickly produces a “map” of the domain on paper. The interviewer starts the expert with a “seed” item from the domain: for example, in zoology, “cat”. Lower levels in the hierarchy are established by asking “can you give any examples of ‘cat?’” (domestic cat, lion, tiger, etc.) and then asking for criteria which distinguish each of them. To go up the hierarchy, ask “what is ‘cat’ an example of?” (mammal).

Burton and Shadbolt (1987a), from whom this account derives, also describe a ‘real-world’ case from the domain of VDU fault diagnosis. They report that while experts are unimpressed and find the technique hard, the laddered grid does provide good quality rules, and does so more easily than an interview transcript. A similar technique is discussed in relation to personal construct psychology (Kelly, 1955), by Boose (1985) and also by Shaw and Gaines (1986); see also 4.b.3 *Repertory grid* below.

Teachback interview

Teachback interviewing is based on the Conversation Theory of Pask (1975), in which the notions of concepts and understanding are treated “as entities made public or objective by an interaction between participants” (Johnson and Johnson, 1987a,b; see also Gregory, 1986). The technique is distinguished from other interviewing methods by being a participant activity rather than one between an interrogator and a respondent (although the principle is identifiable in the stress placed by Spradley (1979) on “restating what the informant says, in his own words”, in ethnographic interviewing (see 4.b.1 *Ethnographic interview*). It is very much “expert driven”, in the sense that any topic to be discussed must originate from the expert, who has to specify at least one link with another topic.

The expert describes a procedure to the interviewer, who then “teaches” it back in the expert’s terms and to the expert’s satisfaction. When they agree, it is said that they share the same concept:

this is Pask's Level O. At Level 1 of the interaction, the interviewer asks the expert how the concept was reconstructed, and the teachback procedure continues until the expert is satisfied with the interviewer's version. Then the interviewer has understood the expert. Level 1 is a matter of explanation of explanations; or perhaps, knowledge about knowledge (metaknowledge). It may correspond to the "epistemological level" identified by Breuker and Wielinga (1987a) as "the missing level in developing expert systems"; at this level of analysis concepts are integrated into a framework of types of entities.

The achievement of understanding (by whatever means) and its expression in a *publicly* examinable form is central to the knowledge acquisition process. Words are less explicit and less reliable as a representation of the mental model of either participant than a graph, diagram or other pictorial representation (Regoczei and Plantinga, 1987). Both the content and the meaning of the knowledge conveyed by the expert must be shared by the knowledge engineer. The conceptual model which results is in effect a formal definition of the language and semantic constraints (meaning) which the expert and knowledge engineer agree to share (Twine, 1988).

Johnson and Johnson (1987b) discuss two very different case studies: children's arithmetic knowledge (treating the 12 year-old child as an expert), and VLSI Design. Teachback methodology is central to both, but it is neither a complete elicitation technique nor an all-purpose one (perceptual skills are excluded, for example). Strengths of the technique include the relative absence of conflation of any preconceptions about the domain which the knowledge engineer might have, with the expert's conception; its non-judgemental nature; and its success in gaining and retaining the expert's interest. Furthermore, there is no doubt about the expert authenticity of the data. On the other hand, it is extremely tiring for the interviewer and as it is not a strongly structured technique, it is particularly important for the interviewer to be trained in interviewing. There is also a very onerous mass of interviewer transcription involved in the method. All of this should encourage short interviews (less than one hour!).

The representational tool used and advocated by Johnson and Johnson for analyzing the qualitative interview data is the Systemic Grammar Network (SGN). An SGN is a graphical structure of nodes suggested by the data, joined by links or relations. It is a more transparent mediating device between verbal data and implementation formalisms than interview transcripts or rule statements (see Johnson (1987) for further explanation).

Ethnographic interview

Spradley (1979) has described a methodology for representing "the acquired knowledge that people use to interpret experience and generate social behaviour" (p. 5); in other words, their culture or shared knowledge. A large part of this is tacit and assumed to be outside the awareness of ethnographers for whom Spradley writes. Ethnographic interviews are a means of accessing this knowledge through "friendly conversations" which nevertheless have an explicit purpose, include frequent explanations of what is going on by the interviewer, and incorporate carefully formulated questions. Spradley identifies about thirty kinds of ethnographic questions, broadly categorized as descriptive, structural (domain organization) and contrast questions (to find meanings which distinguish objects and events). These are discussed in considerable detail, with many examples, for ethnography is a branch of anthropology in which the researcher and the interviewee would normally have a great deal less in common than is usual between knowledge engineers and their experts. Spradley's question classification is worthy of attention from the knowledge engineer who feels "foreign" to the expert and his domain, or in situations where the informant is not used to regarding himself or herself as an expert or is defensive (Welbank, 1987).

One example of structural questioning is card sorting (see 4.b.3 *Card sorting* below). Spradley's ideas about the principle of contrast are based on the work of Kelly (1955): it essentially states that "the meaning of a symbol can be discovered by finding out how it is different from other symbols" (Spradley, 1979, p. 157). Contrast questions include the "twenty questions game" (4.b.1 *Twenty*

questions above), and “triadic contrast questions” (described as repertory grid analysis in 4.b.3 *Repertory grid* below).

A similar taxonomy of questions is presented by LaFrance (1987) as a way of tackling some of the more fundamental problems of knowledge acquisition. These include inadequate understanding of the complex nature of expertise, poorly refined comprehension of the psychology of interpersonal communication in interviewing, and a limited repertoire of questioning strategies (LaFrance, 1987; see also acknowledgement by Linster, 1988). The Knowledge Acquisition Grid, described by LaFrance, has been developed by Digital Equipment Corporation to assist in the manual transfer of expertise. It involves a matrix of five forms of knowledge (“layouts”, “stories”, “scripts”, “metaphors” and “rules of thumb”) and six question types (“grand tour questions”, “cataloguing the categories”, “ascertaining the attributes”, “determining the interconnections”, “seeking advice”, and “cross-checking questions”). This approach offers a way of systematizing knowledge acquisition in a more rigorous and efficient manner and is applicable to any domain. It also underlines the potential contribution of the social sciences to knowledge acquisition (see, for example, the discussion by Gaines, 1987b)). Although knowledge acquisition is fundamentally an anthropological activity (Regoczei and Plantinga, 1987), awareness of this has yet to percolate widely amongst expert system developers.

Introspection

This essentially involves asking the expert to imagine how he or she would (or did) solve a problem or class of problems, typically by an instruction like “how would you solve this (kind of) problem(s)?”. The expert needs to imagine a situation that is concrete or familiar. He/she should be encouraged to “think aloud” in a monologue mode, with the knowledge engineer’s interventions confined to “how”—probing for more information, rather than diversionary “why”—or “what”—probing (Breuker and Wielinga, 1984). Asking “why” may in fact create a source of inaccuracy through the expert inferring an answer without reference to a specific memory record (Ericsson and Simon, 1984). Three techniques have been described in the literature.

Retrospective case description

The expert is asked to describe how he/she handled one or more typical cases (preferably in the recent past). The problem for the knowledge engineer lies in assessing how representative the case is: atypical cases tend to be more readily remembered (Breuker and Wielinga, 1984). The fallibility of retrieval of information from long-term memory has been noted (Ericsson and Simon, 1984) and there is also the problem of assessing the extent to which the expert’s justification of his or her actions is coloured by hindsight (Fenn et al., 1986).

Critical incident

Here the expert is asked to describe his or her experiences of remarkable or difficult cases. Being inherently more memorable (Hart, 1986), they produce more detailed and accurate recall than retrospective case description and can create a stimulating start to knowledge acquisition. Asking how difficult problems were solved may also provide useful information on bottlenecks in expert functioning (Breuker and Wielinga, 1984). On the other hand, problems may arise with anecdotal digression (Welbank, 1983) and from mistaking interesting cases for typical ones (Hart, 1986).

Forward scenario simulation

The expert describes in detail how he/she would solve a hypothetical case, chosen either by the expert (Grover, 1983) or, as Burton and Shadbolt (1987a) emphasize instead, by the knowledge engineer prior to the interview. Trimble and Cooper (1987) used the latter approach in order to elicit

certainty values (Bayes factors). The interview takes place “off the job”, in a laboratory setting. There are several problems. Unless the hypothetical case is chosen with great care (and perhaps even then), the validity with respect to real-world instances may be poor (Breuker and Wielinga, 1984). The expert may be unable to describe his or her reasoning. Worse, his/her account may contain post-hoc justification instead, leading to the construction of spurious decision rules which do not correspond to how he or she actually reasons (Welbank, 1983; Burton and Shadbolt, 1987a). The technique was used for deriving planning strategies in a financial domain by Gleeson and West (1988).

4.b.2 Observation

Watching the expert in action as he or she solves real problems is an obvious alternative to asking him/her what he/she does. It provides some insight into the complexity of the problems, and when coupled with interviewing has been a major source of experiential knowledge (Systems Designers Scientific, 1986). Without any verbal input, however, the technique is limited in its usefulness and is really only for the experienced observer, in a situation where the expert's time is especially valuable (Welbank, 1987). From a study involving experts in copier subsystem design, Dym (1987) recommends asking the experts to perform their task (in this case, working out the design solution at a blackboard) rather than describe it. Several variant methods of observation have been reported in the literature: most involve use of video or at least audio recorders.

Dialogues (non-participative)

Noting the difficulty experts have in verbalizing, particularly their strategic knowledge, Cookson, Holman and Thompson (1985) suggest that a better way of acquiring medical expertise may be to study the clinician's performance while working with a patient. This is the way the experts themselves acquired their diagnostic expertise and the context in which they practise it (Wilkins, Buchanan and Clancey, 1984). However, the problem of assessing how typical the interaction is, noted in 4.b.1 *Retrospective case description* above, arises here too, and as with teachback interviewing the transcription and analysis of the dialogues is very labour intensive.

Expert knowledge can also be observed in the context of an expert-client dialogue where the “client” is a prospective user of an expert system. This has obvious advantages for inferral of the requirements of the man-machine interface (Breuker and Wielinga, 1984), the aim of a study reported by Diaper (1986, 1987, 1989a) in which the so-called “Wizard of Oz” technique was used (a person simulates the expert system, unknown to the user). This technique was also found helpful by Sandberg, Winkels and Breuker (1988) in developing an intelligent tutoring system. Fenn et al. (1986) describe an experimental method of dialogue simulation in which medical experts take the roles of simulated patient and simulated expert system. The potential user and the simulated patient are in one room and the simulated expert system in another. The user carries out “examinations” and can seek advice from the “expert system” by communicating through a voice intercom manned by a human intermediary in a third room (heightening the illusion of interacting with a computer). The researchers found this technique a satisfactory means of assessing the user requirements of an expert assistant and of recording information about an expert's test planning and performance, although it is slow and labour intensive (but reportedly enjoyable); see also the discussion by Berry and Broadbent (1987). Welbank reported at the British Computer Society Expert Systems '86 Conference that Coopers and Lybrand have used a similar technique, with user and expert separated by a curtain to exclude visual or tactile input; she suggested this indirect acquisition technique is good when dealing with defensive experts. Coopers and Lybrand also find observing two experts and recording their dialogue is a very effective technique (Ellman, 1987).

Participant observation

This is appropriate particularly for skill domains, such as medical diagnosis (Kassirer, Kuipers and Gorry, 1982), in which Collins, Green and Draper (1985) argue that participant observation is

necessary to acquire the essential tacit knowledge (that which cannot be articulated). Without spending a period of apprenticeship on the task, it would not have been possible for the knowledge engineer in the crystal-growing project reported by Collins, Green and Draper to distinguish when the expert was asserting real experiential knowledge and when the schematic outline of idealized practice—"textbook knowledge"—took over. The importance of the thesis presented by Collins, Green and Draper has yet to receive general recognition in expert systems research. It is the cultural context within which expertise is practised that gives rules their meaning (Bloomfield, 1988).

4.b.3 Multidimensional techniques

The purpose of these is to elicit structural criteria which the expert uses to organize his or her concepts of the problem domain and thus to form an intuitive map of his/her representation: this structure is commonly not accessible to verbal consciousness (Gammack, 1987b). A well-known example of such a multidimensional map quoted by Burton and Shadbolt (1987a) from the domain of chemistry is the Periodic Table of the elements, which shows their interrelationships.

Card sorting

The knowledge engineer prepares a set of cards, each bearing the name of one concept. These are randomly distributed on a table and the expert is asked to group the cards into piles according to any criterion he or she chooses: this reveals one "dimension" along which the expert classifies the elements (concepts). The task is iterated for all other dimensions along which the expert perceives the elements to vary. Burton and Shadbolt (1987a) offer two illustrations: the first involving animals, a naive analysis which bears little relation to any zoologist's conceptual structure; the second from the domain of VDU fault analysis, where the expert sorted twenty cards along five dimensions into a total of thirteen piles: this was sufficient to distinguish most elements from each other. Schweickert et al. (1987) used this technique with an expert on lighting for industrial inspection tasks; a problem was the expert's tendency to sort on dimensions for which potential system users would not know the values.

A more sophisticated methodology is presented by Gammack (1987c), examining the technical domain of central heating systems (see also Gammack and Young, 1985). First the expert was told to group 75 elements into as many small piles as possible, until no more subdivision was feasible according to the dimension chosen: 25 groups resulted. A "think aloud protocol" (see 4.b.4) was recorded during this process, which the researchers found helpful to the knowledge engineer in understanding the domain and therefore a good idea, providing the expert is agreeable. The expert was then asked to label each group, and next to merge them into slightly larger groups, again while thinking aloud, and to relabel these. This was repeated until no further amalgamation according to similarity was possible for the expert, and resulted in a six-level hierarchy of concepts. Wright and Ayton (1987) suggest the opposite approach, namely to start by asking the expert to sort the cards into just two groups initially, and then into successively smaller groupings. Whilst, arguably, this is intuitively a more likely model of the way many experts would approach such a sorting task, class duplication within the elicited hierarchy is also more likely to result. If the same name recurs in the hierarchy, the risk of confusion is increased.

Gammack and Young (1985) recommend this technique when there is a large set of concepts covering the whole domain, and in need of structuring to become manageable. However, where the domain in question lacks a natural hierarchical structure (evidenced by finding some concepts belonging to more than one group, for example), the result of a card sort will be more a tangled net and therefore of limited usefulness. In its favour, Gammack (1987c) notes that "sorting is a task people find natural and easy, and not just concepts, but pictures, sentences or domain problems may be used as stimuli". On the other hand, if it does not seem to "work" to the expert, it may irritate him; Burton and Shadbolt (1987a) reports that experts are unenthusiastic about the task.

Multidimensional scaling (MDS)

This is a class of procedures for divining structure amongst a set of objects (e.g. domain concepts) whose underlying dimensions of classification are not well-known, and is particularly suitable where the objects can be presumed to vary along only a few dimensions.

Typically, each object is compared with all others and some estimate of their similarity is given. Burton and Shadbolt (1987a) suggest these techniques are complicated, hard to interpret without statistical expertise, and of dubious efficacy in knowledge acquisition; Welbank (unpublished) regards them as special purpose, last resort tools which are also rather strenuous. However, while Gammack (1987c) agrees that it is very demanding of the expert—probably too much so when used on concepts seen as varying along too many dimensions, too few of which are common to all—he describes how once the “right” basic structure has been arrived at by MDS, it can be a useful “conversation piece” for eliciting more detail. This assessment is supported by Olson and Rueter (1987). Scaling techniques reduce dependence on introspection and verbal reporting by experts, compared with traditional interviewing and protocol analysis techniques (Cooke and McDonald, 1987).

Repertory grid

The repertory grid method of knowledge acquisition arises from the theory of personal constructs advanced by Kelly (1955) and summarized by Aleksander (1984), who notes that “the word ‘personal’ in Kelly’s theory draws attention to the fact that different people may classify elements in different ways” (p. 17). The repertory grid is a way of representing the expert’s perception of a problem, in a matrix of constructs and elements. The elements are the domain concepts or other objects whose relationships are to be investigated. Traditionally (but not necessarily), these elements are presented in threes to the expert, who is asked to state a distinction which separates one from the other two; this dimension of difference is a construct, and is necessarily bipolar (e.g. heavy/light). Every other element from the prepared list (usually containing between six and twelve) is then rated along this dimension, often on a numerical scale such as 1–5. All other possible triads are presented in the same way, until all elements are distinct from one another in multidimensional space and the expert is satisfied that the resultant grid accurately represents his or her view of the problem. The grid can then be analyzed by techniques such as cluster analysis and entailment analysis to exhibit the underlying patterns, or to derive rules for input into an expert system shell (Shaw and Gaines, 1986). There is a risk, however, if common sense is subordinated to an excess of mathematical zeal in grid analysis, of constructing an impressive but opaque edifice which is unjustified (Eden and Jones, 1984). Some arguments in favour of the theory of personal constructs as a foundation for a conceptual framework for knowledge elicitation are presented by Garg-Janardan and Salvendy (1987).

Several differences between repertory grid and MDS techniques are discussed by Gammack (1987c), who concludes that “MDS allows averaging and statistical comparisons across a number of experts, while repertory grid is more suitable for modelling an individual” (this is rather unduly restrictive, as will be noted below). Identification of all the appropriate dimensions or constructs is a common problem, however (Cooke and McDonald, 1987). Both techniques are particularly appropriate when there exist closely related concepts, not well-differentiated by novices, but susceptible to discrimination by experts—who are often unable to articulate their decision criteria. The repertory grid technique can elicit finer grain criteria than interviews (Gammack and Young, 1985) and is particularly suitable for accessing implicit knowledge (Gammack, 1987b).

Credit for the development of repertory grid and related techniques based on personal construct psychology is due particularly to Shaw, who in cooperation with Gaines has developed a suite of computer programs which interact with experts to enable them to express the constructs underlying their knowledge (e.g. Shaw and Gaines, 1987). An example arising from their recognition that “much expertise only resides within the social context of co-operating individuals and requires elicitation across the group” (p. 277), is the SOCIO analysis program, which facilitates comparison of the construct systems of several users. The major part of this state-of-the-art report is discussed in

more detail by Shaw and Gaines (1986). A less current overview, but nevertheless a useful introduction to the subject, is offered by Easterby-Smith (1981). The principal handbook is still the text by Fransella and Bannister (1977). For a shorter account, concentrating on the analysis of a case study, the reader is referred to chapter nine in the book by Hart (1986).

Apart from relieving the tedium of manual analysis of the repertory grid, another advantage of computerization may be to permit the inclusion of more elements from the domain: Welbank (1983) has noted that (under manual operation) the procedure “becomes unmanageable with more than about 10 objects, and most domains have many more” (pp. 19–20). Although it might seem natural to suppose a great multiplicity of elements would be required to create an accurate representation of the expert’s view of a problem, this is apparently not so: Easterby-Smith (1981), for example, suggests that no more than twelve elements should be necessary to provide adequate coverage.

The principal criticism of the repertory grid technique has been the subjectivity of the results (Hart, 1986). However, the comfortable and slightly old-fashioned notion of expert “objectivity” has been powerfully attacked by Partridge (1985) and by Agnew, Brown and Lynch (1986). In any case, as Hart (1986) notes “the expert does not simply identify a cause which explains an effect: he sees the problem in a particular way” (p. 134), which distinguishes him from a novice. Whether an expert should be expected to see it in terms of bipolar constructs is another question. It is rather easier to envisage this approach succeeding in the “soft sciences”—Shaw (1981) includes applications in clinical psychology (Kelly’s field), education and management—but evidence of knowledge acquisition using the repertory grid in the hard sciences or manufacturing is difficult to find in the literature. The technique may be best suited to domains where there is a low level of consensus among experts and a correspondingly high degree of perceived subjectivity.

Proximity analysis

This is a method developed by Gammack (1987a,b) to collect information relevant to creating a network representation of domain concepts. The expert is presented with a set of concepts and asked to sort them by referring each one in turn to one concept selected as a reference. Three piles result: “a relationship exists between the two”; “no relationship exists”; “there may be some relationship”. The relatable concepts (from the first pile) are then arrayed on a scale from 0 to 100 and moved around until the expert is satisfied with the ordering and relative distance between them; scale values are then entered in a data matrix. The procedure is repeated until all concepts have been used as references. Gammack analyzes the proximity matrix using a computer program called PATHFINDER (Schvaneveldt, Durso and Dearholt, 1985) to produce a network. One advantage of this method is that it shows which links in the network of domain concepts are likely to be meaningful, with associated strengths, since the links represent conceptual proximity. Proximity analysis complements MDS (4.b.3 *Multidimensional scaling (MDS)*) by retaining information on local structuring. Researchers at New Mexico State University have independently developed similar techniques (e.g. Cooke and McDonald, 1987; McDonald and Schvaneveldt, 1987).

Matrix technique

A further development of proximity analysis, not yet fully developed and tested (Gammack, 1987b) is to use the network to extract propositional information about the relationship among its concepts (i.e. the labels for the links). Gammack took each pair of concepts identified by the expert as relatable, in turn, and asked the expert to record a one-predicate sentence describing the relationship between them. After weeding out symmetrical propositions, around 124 propositional relationships were obtained among the 24 core concepts Gammack analyzed; in the space of just two hours. The procedure is easy to automate. Cooke and McDonald (1986, 1987) are also researching a methodology of identifying link labels.

Gammack points out an advantage of the technique over protocol analysis (4.b.4), using as an example (from his chosen domain) the phrase “. . . the burning gas heats the water in the boiler . . .”.

Unless the knowledge engineer is sufficiently knowledgeable about the domain, such invalid propositions as “gas is in boiler” may be created in analyzing the protocol: all sorts of errors, of omission and commission, may be made (Gammack, 1987b).

4.b.4 Protocol analysis

The question of accuracy of the verbal self-report has been extensively discussed in the literature (see Nisbett and Wilson, 1977; Ericsson and Simon, 1980, 1984; Wielinga and Breuker, 1985) leading to the conclusion that it provides “a nearly complete record of the sequence of information that is heeded during task performance” (Ericsson and Simon, 1984, p. 62). The classical technique of protocol analysis involves asking the expert to “think aloud” all the time while solving a problem; this is tape recorded, transcribed in each detail (along with notes of actions) into a “protocol”, and then analyzed for meaningful associations (often with the aim of deriving production rules). The knowledge engineer intervenes during the task performance only with non-directive reminders to keep thinking aloud. The effects of verbalization upon task performance have been examined by Ericsson and Simon (1984), who found that, provided the subject was not asked to explain his/her thoughts, instructions to think aloud do not significantly alter the sequence of cognitive processes.

The principal merit of protocol analysis is that in deriving from a much more true-to-life task situation than any interview method, the unconstrained verbalizations of the expert may reveal the knowledge (particularly the heuristics) that experts use in problem solving, which they are unable to articulate in an interview. This knowledge usually has to be inferred from the data. Welbank (1983) suggests that protocol analysis is most often used to verify rule sets derived from theory, or for comparison with the reasoning methods which experts say they use.

The results are less satisfactory in the analysis of perceptual-motor tasks than of other kinds, because subjects tend to verbalize only higher-level evaluations of solution attempts (e.g. “yes”, “looks okay”, “oh, no”, etc.) and indicators of attention to control processes (e.g. “let’s see now”, “what’s it doing”, etc.): the subject is apparently unaware of the lower levels of his thought processes (Ericsson and Simon, 1984). Welbank (1983) and Bainbridge (1986) also note this problem of lack of conscious awareness in the context of motor skills such as skating, swimming and riding a bicycle. However, it should be remembered that such tasks, though often verbally employed as illustrations of the difficulty of explicating expertise, are unsuitable for expert system development anyway: “the task must require cognitive, not physical skills” (Waterman, 1986, p. 128).

The whole process of protocol analysis is very time consuming, putting a high premium on making a representative selection of problems in the beginning. Asking the expert to do this sorting may also provide information on the way he or she classifies the domain (Breuker and Wielinga, 1984). A certain amount of private practice in thinking aloud may be necessary or even desirable, to release the inhibitions in some experts. Breuker and Wielinga (1984) recommend also that protocol analysis should follow the use of interview techniques because the latter are likely to be perceived as more “helpful” by experts. Knowledge engineers in any case need to have acquired a good understanding of the domain beforehand in order to understand the expert’s task, and as Bainbridge (1979) notes, interviews may be the best way of acquiring this knowledge. Waldron (1985), on the other hand, takes the opposite view, arguing that “process tracing [through protocol analysis] will enable interviewing time to be more efficient by familiarizing the knowledge engineer with the knowledge acquisition process” (p. 664). Wilson (1987) similarly argues cogently for using the verbal protocol as a set of cues for subsequent interviews. Waldron discusses many practical problems associated with interviews, particularly sources of inefficiency, and stresses also the importance of the structure of the knowledge base reflecting the expert’s natural approach to the problem.

The methodology of protocol analysis has been discussed with reference to industrial process operators by Bainbridge (1979) and to causal reasoning in medicine by Kuipers and Kassirer (1984), Patel and Groen (1986) and Kuipers, Moskowitz and Kassirer (1988). An interesting—and cautionary—point related to the general problem identified by Bainbridge (1979) that the expert may simply not verbalize what is “obvious” to him, is reported by Kuipers and Kassirer (1984):

The expert physician, with many years of experience, has so “compiled” his knowledge that a long chain of inference is likely to be reduced to a single association. This feature can make it difficult for an expert to verbalize information that he actually uses in solving problems. Faced with a difficult problem, the apprentice fails to solve it at all, the journeyman solves it after long effort, and the master sees the answer immediately. Clearly, although the master has the knowledge we want to study, the journeyman will be much easier to study by our methods.

Wright and Ayton (1987) make a similar point in their discussion of protocol analysis. A general review of problems associated with protocol analysis as a method of knowledge acquisition may be found in Berry (1987).

Protocol analysis is used within the KRITON system for computer-aided knowledge acquisition (see 5.b.2 *KRITON* below) to observe “in action” the knowledge ideally elicited in some previous interview or text analysis process (Diederich, Ruhman and May, 1987a). Johnson, Zuolkernan and Garber (1987) describe a representation of the kind of knowledge used to perform some design and diagnostic tasks, based on analysis of verbal protocols.

Dillard and Mutchler (1987) present a study of expertise in assessing solvency problems in accounting, using protocol analysis techniques. An empirical analysis of the knowledge that knowledge engineers themselves use in their job, based on videotaped protocols, is reported by Littman (1987). Lewis (1982) and Jerrams-Smith (1987) have both used protocol analysis in developing an expert system within a man-machine interface. Barnard, Wilson and Maclean (1986) used the technique to elicit a user’s knowledge of an office software package. Protocol analysis was chosen as a means of understanding media planning at an advertising agency by Mitchell (1987). Fox et al. (1987) discuss its use and limited success in developing an expert system for test interpretation in leukaemia diagnosis. Other applications are mentioned by Welbank (1983) and by Bainbridge (1979, 1986).

Shadbolt et al. (1987) counsel strongly against the use of protocol analysis alone in knowledge acquisition, claiming it is a much less effective technique than the ladder grid (see 4.b.1 *Ladder grid*) and card sort (4.b.3 *Card sorting*) methods, which primarily access declarative knowledge; but this is exactly what one would expect in a classification task such as Shadbolt et al. examine (the identification of igneous rocks; see also Burton and Shadbolt, 1987b). More detailed discussion of protocol analysis is to be found in Wielinga and Breuker (1985), as part of their development of a systematic knowledge acquisition methodology (see 5.b.2 *Knowledge acquisition and document structuring (KADS)*). A contribution towards standardization of methods for encoding and analyzing protocols, called Structural Analysis, is presented by Bailey and Kay (1986). The fundamental text references are Newell and Simon (1972) and Ericsson and Simon (1984).

Eidetic reduction

This is a refinement of introspective and verbal self-report techniques, introduced by Killin and Hickman (1986a,b) to access the subconscious activity of the expert’s mind when formulating a problem definition. Some real world problems such as mathematical modelling are sufficiently complex to resist a classificatory approach and a problem definition has to be constructed (especially through the use of analogy). In this process the expert relies in part upon a set of implicit assumptions of which he or she is unaware.

The technique of eidetic reduction, or epoché, involves a shift in consciousness from the unreflected “doing” of a task to awareness of that activity and one’s role in it. Killin and Hickman (1986b) illustrate this change of consciousness with the situation which arises if, while walking, one is confronted by another pedestrian whose attempts to avoid a collision exactly mirror your own. Eidetic reduction enables the expert to mentally “stand aside” from the actual doing of a task in order to critically examine what is happening and to describe the associations he or she is making, to the elicitor. “The benefit of this technique lies in its near immediacy, and the associated absence of

reliance upon memory, which thus inhibits the scope for post rationalised justification” (Killin and Hickman, 1986b, p. 3). Whereas for protocol analysis the expert generates a verbal self-report by “thinking aloud” as he/she solves the problem, eidetic reduction seeks to access that which he/she is not even aware he/she is doing.

Video recording is important, to maximise explicit capture of data and avoid any later need to infer; in this way the elicitor gets as close as possible to the expert. The technique is only suitable for use with a case study new to the expert; it depends also upon the expert grasping exactly what is required of him or her, and the dialogue being strictly controlled (Killin, pers. comm.).

4.c Major conclusions and issues for section 4

The principal techniques of knowledge acquisition at the level of human interaction have been reviewed and distinguished in terms of ease of use, expert appeal, problems and drawbacks, suitability, accuracy and reported actual use in practice. If comparatively little attention has been paid until recently to documenting knowledge acquisition methodology, authors have typically offered even fewer clues to the modelling of knowledge. Implicitly or explicitly, there seems in many cases to have been a move directly from the data to implementation by encoding the knowledge in IF . . . THEN rules.

This leaves out the crucial intermediate stage of formulating a conceptual model of the expert’s knowledge. This is important not least for reasons of expert validation: there must—if the system is intended to be a genuine attempt at replication of the expert’s problem solving expertise—be agreement between expert and knowledge engineer that the latter has properly and adequately acquired a knowledge of the concepts involved and their interrelationships, as used by the expert. Rather than presenting a collection of production rules, which is generally an alien representation to most experts, the achievement of the proper level of understanding is more likely to be gained by examining together a map of the knowledge on paper first. This paper model might incorporate network or hierarchic forms of entity relationships, including inheritance features or weighted links; other graphical forms such as decision trees, and assertions and conditional statements.

It should be recognized that the expert may within the domain be performing several problem-solving tasks and on top of that, be communicating in more than one functional role or modality. Identification of the intended modality for the system (e.g. advising, diagnosing, tutoring) will facilitate definition of the problem solving tasks within that modality for which conceptual models are required.

5 Machine-aided knowledge acquisition

At the British Computer Society Expert Systems ’86 Conference, Gaines presented a state-of-the-art report from the AAAI Workshop on Knowledge Acquisition held in November 1986. In this he noted that the trend seems to be towards eliminating the knowledge engineer and getting the expert working directly with a shell and predicted that knowledge acquisition tools will be a big development area in the next few years (Gaines, unpublished). Recent market surveys suggest that Europe is about 12–18 months behind the USA (more in finance, insurance and defence); this gap is particularly noticeable in the use of advanced tools (Hewett, 1987). From this the future of knowledge acquisition may seem to lie with machine aids. Shaw and Gaines (1987) identify several reasons to doubt that human labour is the appropriate solution for “the knowledge engineering problem”. These seem to reduce largely to economics (human labour as a proportion of expert system development costs rising more rapidly than other costs; a demand-pull market dictates more rapid production), although Shaw and Gaines also suggest that using a human intermediary may be less effective: “knowledge may be lost through the intermediary and the expert’s lack of knowledge of the technology may be less of a detriment than the knowledge engineer’s lack of domain knowledge” (p. 111). Many would argue against this—Waterman (1986, p. 154), for example, advises succinctly “don’t be your own expert!”—but the point here is that expert systems are not immune from universal trends

towards packaging, pre-processing and convenience and away from handcrafting useful things from raw materials. An important benefit of research in tool development has been the attention focused on the issue of modelling expertise.

Expert system tools may be classified into three broad groupings (Waterman, 1986); programming languages (particularly the symbol manipulation languages LISP (and its many dialects) and PROLOG); knowledge engineering languages; and system-building aids. From the many research programs developing system-building aids (including knowledge acquisition tools) examples of six principal approaches are discussed. Proponents of induction from examples as an alternative to interviewing for knowledge acquisition were among the first in the field: a typical result is the EXPERT-EASE tool. A second class of tools function essentially as knowledge base editors and tend to require a good knowledge of expert systems at the implementation level: an early example is TEIRESIAS. More recently, tools based upon problem solving methods have been developed; these vary according to the specificity and explicitness of the model and constitute the third and fourth classes: examples are ETS and MOLE. Some other system building aids are specific to particular tasks, such as KNACK. Finally, several research programs are following a language based approach, incorporating facilities for building expert systems across a wide spectrum of expertise: an example is the KADS project. In the discussion of knowledge acquisition tools the author is particularly indebted to the insights of Breuker et al. (1987), McDermott (1986 and in press) and Musen (1988).

Programming languages are rather peripheral to this enquiry. System building aids (including induction-based tools) are particularly relevant, but the much greater activity (up to now) in the market place of knowledge engineering languages merits some consideration of the latter first. The author wishes to stress that this section of the paper is based upon the literature rather than personal knowledge of the tools discussed.

5.a Knowledge engineering languages

Waterman (1986) categorizes these as either skeletal systems (more commonly known as “shells”) or general purpose systems (also variously termed “programming environments”, “development environments”, “hybrid toolkits” or “toolkits”). Shells are simpler and therefore will be examined first.

5.a.1 Shells

According to the journal of one leading vendor, Texas Instruments, “about 85% of all expert systems are now being produced with shells, as opposed to being developed in a [programming] language” (Boothroyd, 1986). The principal merits and demerits of shells (here taken to mean tools, often designed for use on PCs and generally very restricted in the facilities offered) for building expert systems may be summarized thus: they offer a faster, cheaper development route but they constrain the designer in the limited formalisms they support, beyond what is acceptable if the knowledge to be represented is of any real complexity (Hayward, 1985). One (admittedly partisan) view is that “often, they turn out to be Procrustean beds, where only some form of disabled expertise will fit” (Breuker and Wielinga, 1987a). Almost all shells, according to one authority, are restricted to one kind of reasoning, usually backward chaining, and many are based on the production rule formalism (Goodall, 1985); examples are the archetypal EMYCIN, Teknowledge’s M.I., Expertech’s Xi Plus and Intelligent Environments’ CRYSTAL. The predominance of backward-chaining inference is no accident; it is the easiest form of inferencing to implement and also the most suitable form of reasoning for diagnosis, which is one of the largest application domains for expert systems developers so far (Milne, 1987). The situation is rapidly developing, however: many shells such as Xi Plus, I.S.I’s SAVOIR and Texas Instruments’ PERSONAL CONSULTANT PLUS can perform forward chaining also, and other shells such as Creative Logic’s LEONARDO, which is an object structured system, offer additional facilities such as procedures, screen editors and mechanisms for

representing uncertainty. The line between the more advanced shells and hybrid toolkits is becoming increasingly blurred and meaningless.

The most rapidly developing sector of the market (in Britain at least) is in shells which can run on microcomputers (Bramer, 1987a): both Xi Plus and CRYSTAL, two developments of 1986, will run on an IBM-PC and compatibles, and can hold several thousand rules in memory. These and other shells for micros are chiefly aimed at business organizations interested in assessing the worth of an expert system for a particular application, or just finding out what expert systems can do (Gold, 1986). Gold asserts however that “most are beyond the technical proficiency of the average business user” (p. 77) and notes that “most vendors recommend that users have at least some experience in programming” (p. 79); Forsyth’s review of Xi Plus concludes similarly that “the learning curve is rather steep” (Forsyth, 1987, p. 50). The same author describes PERSONAL CONSULTANT PLUS as “unsuitable for domain experts who want to build their own expert systems without undue effort” (Forsyth, 1986, p. 246). An important class of exceptions to Gold’s contentions (which he acknowledges) are shells based on examples, which operate by inductive reasoning (such as EXPERT-EASE). Consideration of induction is given in section 5.b.1. The strong demand for increased flexibility and “user-friendliness” will undoubtedly be a major influence in improving the ease with which knowledge can be acquired from computer-naïve users by shells, but as Bramer (1986) noted, “the technology is simply not at a stage where expert system building can be turned into a production line activity”.

5.a.2 Hybrid toolkits

These offer the user a choice of knowledge representations and inference methods: examples are KEE, ART and KNOWLEDGE CRAFT. They vary a great deal in features, flexibility and cost (see the four-page survey, covering eighteen products, in the December 1986 issue of *Expert Systems User* or the review by Gevarter (1987)). As suggested earlier, recent technical developments have led to a rapprochement between toolkits and shells and use of the latter term in the US extends to KEE, ART, etc. Reichgelt and van Harmelen (1986) contend that these high level programming environments are not the solution to the limitations of shells as described in the previous section, particularly for non-AI programmers, because “they provide the knowledge engineer with a bewildering array of possibilities and little, if any, guidance under what circumstances which of these possibilities should be used” (p. 2). This view is supported by Breuker and Wielinga (1987a) and by Mettrey (1987).

A general conclusion which may be drawn is that knowledge engineering languages offer a ready-made framework for building an expert system and in this way facilitate development, particularly of early prototypes. Against this advantage must be set the varying degrees of constraint they impose, and the possible resulting inaccuracy of the knowledge acquisition process in representing expertise. Indeed, as Cooke and McDonald (1986, p. 1424) observe, to choose the knowledge representation scheme before acquiring the knowledge is “putting the cart before the horse”. Nevertheless, where there is a proven development already—such as PUFF, a system for diagnosis of lung disorders implemented in EMYCIN (Waterman, 1986)—it may on balance be beneficial to use a knowledge engineering language for a closely-related activity: Fox et al. (1987) chose to use EMYCIN for a study of leukaemia diagnosis. Buchanan (1986), who was involved in the EMYCIN project, observed that “its utility is limited to a class of problems like Mycin’s” (p. 32).

5.b System-building aids

5.b.1 Induction

An important class of system-building aids relies upon the principle of induction of rules from examples supplied by experts. Developers of these tools, such as EXPERT-EASE, have attempted

to eliminate both the knowledge engineer and the need for previous structuring or textual analysis of the domain knowledge. Induction is one form of machine learning.

Interview methods of knowledge acquisition have many acknowledged inadequacies. Winston (1982) considers them too slow, too expensive, and unreliable. The difficulty experts often have in articulating even part of what they know, or what is typical and representative of the domain, is compounded by their lack of access to the tacit element of their expertise (Collins, Green and Draper, 1985) and indeed by the ways in which they solve problems generally (Hart, 1985b; Bloomfield, 1986). Experts may not mention what is "obvious" to them (Fox, 1984). Even where the expert is articulate and willing to cooperate, there may be a great difference between the way the expert describes his/her reasoning and the method he/she actually uses (Bell, 1985): he/she may simply rationalize from theory (Fox, 1984). The very act of asking experts about their methods can change their view of them (Hart, 1985b).

Whilst expert systems have often been based upon rules, there is evidence that experts do not consciously think in terms of rules, but instead base their reasoning on cases or analogies (see e.g. Kolodner, 1984; Carbonell, 1985). However, Sweller, Mawer and Howe (1982) observe that this "history-cued strategy may facilitate rule induction" (p. 455), and it is certainly characteristic of experts that they can readily supply good cases or examples of their decisions (Michie, 1986b). Indeed, for experts unaccustomed to formalising their expertise, there may be no other natural notation (Gruber and Cohen, 1987). Experts do not eternally reinvent wheels: they pass on their knowledge to apprentices, through their ability to identify key concepts of the domain and present these through an appropriately sequenced set of tutorial examples (Michie, 1986b). This is the rationale behind machine induction.

The principle of induction is that the expert supplies a set of domain examples of different types of decisions, called a training set, together with the attributes which describe the examples, and values he/she assigns to those attributes. From the training set, a computer program using an inductive algorithm induces a set of rules, which are often constructed in the form of a decision tree. Domains where rules are not appropriate are unsuitable for induction: these include some real-time problems where quick reaction to an unpredictable sequence of observations is required (Hart, 1987b).

A famous example of success with machine induction is the work reported by Michalski and Chilausky (1980). They used a set of 630 examples of 15 different soybean diseases and 35 attributes such as "leaf malformation" and "seed size" as well as environmental descriptors like "temperature". The inductively-derived decision rules were found, to the authors' expressed surprise, to perform better than rules derived from "discussions" with plant pathologists lasting about 20 hours. They acknowledged the possible deficiencies in their "manual" acquisition methodology (which is not described in any detail), but this has been generally ignored in favour of their conclusion that induction is "a viable knowledge acquisition method if the problem domain is sufficiently simple and well-defined" (p. 79). This qualification is significant and important.

The training set has to cover all possible cases—an incomplete or inadequate set is likely to result in poor rules (Hart, 1986)—including unusual or "difficult" cases (White, 1987): but since an expert cannot account for all he or she knows, how can he/she (or a knowledge engineer) be sure that the training set is adequate? Furthermore, how can one be sure the attributes supplied by the expert constitute a sufficient set for the construction of a valid decision tree? Quinlan (1982), the originator of the well-known ID3 algorithm (a program which induces a decision tree of IF . . . THEN rules from a training set of examples; a generalization of it is the basis of the system-building aid EXPERT-EASE), reports that "finding small but adequate sets of attributes for the [chess end game] king-rook king-knight problems was a considerable task" (p. 201). And even if the examples and attributes are adequate, can the expert unambiguously assign values for the attributes? These and many other non-trivial problems (Hart, 1985a) associated with induction are discussed in the critical review by Bloomfield (1986), who also observes that one cannot be certain that an inductive inference is true unless all of the objects in the domain are known, because induction is really a form of conjecture. Some guidelines for inductive knowledge generation are provided by Hart (1987a).

The ID3 algorithm has been particularly discussed by Bramer (1987b) and by Hart (1985a,b; 1986; 1987b), who notes that whilst it assumes (in its original form) that the examples are complete and correct, this is often not so. The domain of chess endgames (for which ID3 was developed) is probably exceptional in containing no noise or uncertainty. Several extensions of ID3 have been developed to cope with this problem (e.g. Quinlan, 1986; Mingers, 1987; Niblett and Bratko, 1987; White, 1987) and also the INSTIL project which is investigating noise in the knowledge acquisition phase as well as in rule induction (Blythe, 1987; Blythe, Needham and Corsi, 1987). However, it is interesting to note Mingers' conclusion about the potential of induction for knowledge acquisition for expert systems: whilst the degree of success is encouraging,

“... it will always be limited by the quality and amount of data available; it still requires a good deal of expert time to supply the attributes and examples, and to validate the rule tree; and it will only really be suitable for fairly straight-forward, well-defined domains” (Mingers, 1987, p. 46).

The last point is supported by Trimble and Cooper (1987), who report other problems with induction:

“One of these is that the natural sequence of questioning that is inherent in a domain is not respected. For example a pair of questions might be:

*Is the pipe a drain?

*Have you performed a drain-test?

If the sequence of these questions is reversed as it may be in rule induction the confidence of the user will quickly evaporate.” (p. 6).

An alternative improvement on programs based on ID3 is described by van Terheyden and Chalcraft (1987): the knowledge base generated, from examples, by their program (either as a set of rules or as an inference network) can be readily modified by the expert as he sees fit. This proposal for unifying the inductive and deductive approaches may overcome some of the limitations of induction mentioned above. A similar interactive approach is the basis of the Berlin Learning by Induction Project (BLIP), which allows the expert to reverse layouts and rules at all times and thus supports both top-down and bottom-up acquisition (Morik, 1987a,b).

5.b.2 Knowledge acquisition tools

TEIRESIAS

Named after the blind seer in the ancient Greek tragedy “Oedipus the King”, this program like the prophet has a form of “higher order” knowledge, or meta-knowledge (Davis, 1979). This is knowledge about knowledge: in TEIRESIAS, knowledge about how MYCIN's knowledge is represented and used (such as about the structure of rules) (Brachman et al., 1983). In effect, “the system has a model of the knowledge it is likely to be acquiring from the expert” (Davis, 1979, p. 138). The incorporation of TEIRESIAS and EMYCIN provides better explanatory abilities, because the meta-knowledge allows understanding of a rule to be offered at various levels of detail instead of giving just a bald restatement of the rule (Brachman et al., 1983).

The function of TEIRESIAS is to help a domain expert add new knowledge to MYCIN. This process of interactive transfer of expertise “requires the strongest degree of cooperation from the expert” because “TEIRESIAS does not attempt to derive knowledge on its own” (Davis, 1979, p. 124). Thus TEIRESIAS is particularly applicable at the later stages of knowledge acquisition, involving refinement and expansion of the knowledge base: it needs a context upon which to work and does not aid the initial knowledge acquisition (Cooke and McDonald, 1986). In operating in this way at the symbol level it offers no support for analysis at the knowledge level. Most system building aids are knowledge base editors like TEIRESIAS (Musen, 1988). Welbank (1983, p. 34)

offers an overview of TEIRESIAS, which was developed at Stanford University as a research system (Waterman, 1986).

Expertise transfer system (ETS)

Based on personal construct psychology (see 4.b.3 *Repertory grid*), ETS is a program which can be used to supply the initial knowledge base to TEIRESIAS: unlike the incremental approach of TEIRESIAS to system building, ETS is designed to facilitate rapid prototyping “from scratch” (Boose, 1985, 1986a). In a more recent statement, Boose (1986b) describes ETS as a complete methodology for knowledge acquisition:

“Currently ETS interviews experts, analyses the elicited knowledge, builds knowledge bases to test the knowledge, uses feedback to help refine the knowledge bases and helps the expert combine his knowledge with that of other experts.” (pp. 192–3).

Thus unlike TEIRESIAS the interaction does not centre on rules, but instead upon a process of knowledge acquisition driven by an implicit model of problem solving. Because this model is shielded from the user, little or no knowledge of the computer system is required. The program does the job of fitting the data into its internal classification model. ETS expresses the expert’s knowledge in the form of a rating (repertory) grid, and generates rules from this accompanied by certainty factors based on the relative strength of the rating and the relative importance of the construct in each case. Boose (1986b) suggests that this process need occupy no more than an hour or two, and elsewhere (Boose, 1985) that “an initial 15 min. explanation is usually all that is needed” (p. 496) to train an expert in the use of ETS. Thus in a sense the technique becomes “expert-driven” and avoids many of the problems of human interviewing—initially.

The limitations of ETS for knowledge acquisition are frankly acknowledged by Boose (1986b)—although attempts are being made to overcome these (Shaw and Gaines, 1987), in particular through the development of AQUINAS (Boose and Bradshaw, 1987; Kitto and Boose, 1987a; Boose, Shema and Bradshaw, 1988). Grid methodology is only well-suited to analytical tasks such as diagnosis and classification, in which the *selection* of solutions involves less complex reasoning structures and variations than in other areas where solution is a process of *construction* by configuration of the elements (Breuker et al., 1987). Experts may not perceive meaning in the statistical certainty factors attached to rules derived from the grid. ETS does not easily elicit “deep” causal knowledge and additional human interviewing is still necessary to produce an expert system: ETS elicits only simple classifications from the expert, does not exploit distinctions between hierarchical, definition and heuristic relations, and has no provision for data abstraction (Clancey, 1985, p. 343).

ETS is now used (Boose, 1988) for building small-scale prototypes in the early stages of expert system projects. Its successor, AQUINAS, is a larger system which integrates several tool sets into a workbench and makes the problem solving method more explicit to the expert using it; on the other hand, it is less easy to use than ETS because it is much more complex. AQUINAS is useful for a variety of analysis tasks.

ROGET

The ROGET system (Bennett, 1985) is an attempt to generalize from MYCIN, based upon a survey and analysis of diagnostic expert systems, and was the first knowledge acquisition tool to be based on a problem solving method. ROGET uses its knowledge base of “conceptual structures” derived from this analysis to help the expert identify the vocabulary of a problem representation and its organization for a diagnostic expert system. A large part of ROGET’s knowledge base consists of a thesaurus of phrases and synonyms, which enables a simple semantic-grammar parser to recognize which diagnostic problem solving task is involved. The system then proceeds to acquire the conceptual structure for the new expert system by asking questions and presenting descriptions of the structure it holds. This is elaborated and pruned (ROGET can advise the user on the feasibility of

systems) and finally the conceptual structure is made operational by translating it into a shell (currently only a substantially modified version of EMYCIN). Unlike ETS therefore, ROGET makes its intrinsic problem solving method explicit in conducting the knowledge level analysis.

Bennett recognizes that much work is required before ROGET can be accepted as an expert system design tool, particularly on acquisition of inference methods and strategies. Although it may be extensible to other problem solving tasks—such as planning or design—ROGET can at present recognize only domains very similar to MYCIN, and it is difficult to see how well the conceptual structure of MYCIN would apply to other tasks. Although the original philosophy was “multi-model”, Breuker et al. (1987) conclude that in practice ROGET is single-model based.

MORE

MORE (Kahn, Nowlan and McDermott, 1985a,b) is a diagnosis knowledge acquisition system similar in some ways to TEIRESIAS and ETS in providing a mechanism for interviewing domain experts. It differs in its understanding of what kinds of knowledge are likely to be significant in diagnosis, and is thus able to recognize gaps and errors in an existing knowledge base. This drives the interaction process with the expert and leads to stronger diagnostic conclusions. Like ROGET, it is based on an explicit problem solving method, a form of heuristic classification (Clancey, 1985). MORE has been used to build fault diagnosis systems in domains such as computer networks (Kahn, Nowlan and McDermott, 1985a).

MOLE

MOLE (McDermott, 1986; Eshelman et al., 1987) is similar to TEIRESIAS, ETS and MORE in that it builds expert systems that perform diagnosis and classification and like MORE is based on a variant of heuristic classification: the problem-solving or “role-limiting” method presupposed by MOLE is termed “cover and differentiate” (McDermott, in press). MOLE initially asks the expert to list all the commonly relevant events (hypotheses and symptoms) in his/her domain, and then requests knowledge which may explain or cover the symptoms and knowledge which helps differentiate which hypothesis is the most likely explanation of a symptom. MOLE then uses this knowledge to build a network of associations.

MOLE essentially uses the interviewing methods described in 4.b.1 *Systematic symptom-to-fault links* and *Intermediate reasoning steps*, although it recognizes that the expert may not initially be able to provide a complete specification of events and associations. A second phase, dynamic analysis, involves comparison of the diagnosis of the performance system within MOLE with the expert's diagnosis. If they differ, MOLE-the-knowledge-acquisition-tool assumes some knowledge is missing, reasons what type of knowledge is required, and asks the expert for it.

MOLE has been used to build systems for diagnosing problems in car engines and in a steel rolling mill, and diagnosing inefficiencies in a coal burning power plant (McDermott, in press). Further information about MOLE will be provided in the forthcoming book by Marcus (in press).

SALT

SALT (Marcus, 1987; see also McDermott, 1986, in press), is, like MORE and MOLE, a “single model” tool; in this case for building expert systems which use a propose-and-revise method for synthesis tasks such as construction, plan creation and design. Like MOLE, SALT is intended for use by domain experts who it is claimed do not need prior knowledge of programming (although Musen et al. (1987) suggest the conceptualization process required of the expert may seem unnatural). It assumes that a complete set of solutions can be pre-specified for each possible problem which may arise during completion of the task. SALT is not, therefore, suitable for building systems whose tasks are less tractable, such as XCON's (McDermott, in press; see also Soloway, Bachant and Jensen (1987) on the RIME methodology developed to rewrite XCON).

Like ROGET, MORE and MOLE, SALT conducts the interaction with the expert at the knowledge level by asking for a task description in terms of the types and relationships of the model within the tool. SALT has been used to build VT, which configures elevator systems (Marcus, Stout and McDermott, 1988). For further information about SALT, see the relevant chapter in Marcus (in press).

KNACK

Another product (like MORE, MOLE and SALT) of research at Carnegie-Mellon University (CMU), KNACK builds expert systems called WRINGERS which use an acquire-and-present method to help non-programmer domain experts produce reports on plans or designs (Klinker et al., 1987; McDermott (in press); see also specific chapter in Marcus (in press)). Whereas MORE, MOLE and SALT are each suitable for a family of related tasks, KNACK appears to be based on the more specific task of report generation. The expert is required to provide a skeletal report and the domain concepts and vocabulary; KNACK then helps the expert integrate these elements into a WRINGER. The expert is then asked to define sample strategies that the WRINGER is to use to acquire values instantiating the concepts in the generalized sample report (McDermott, in press). A WRINGER may, for example, be created to provide information on flaws and possible remedies in designs.

OPAL

OPAL (Musen et al., 1987; Musen, 1988) is a graphical tool which allows cancer experts to enter and review treatment plans for use by the expert system ONCOCIN. Like KNACK, the conceptual model adopted by OPAL is task-based and applies a form of skeletal-plan refinement as the problem solving method. The knowledge-level analysis (of cancer treatment planning) has already been done by the tool developers; the user is required only to enter into the template the content knowledge for instantiating the concepts for the particular therapy programme to be prescribed.

Knowledge acquisition and document structuring (KADS)

KADS (standing both for a methodology and the set of tools that support this methodology) is the work of Breuker and Wielinga at the University of Amsterdam (e.g. Breuker and Wielinga, 1984, 1987a,b) and co-workers (e.g. Anjewierden, 1987; Breuker et al., 1978; Hayward, Wielinga and Breuker, 1987; Schreiber et al., 1988). Currently the best short summary of the methodology is provided by Hayward (1987). The idea of KADS is very similar to that of ROGET, but whereas the latter is modelled on the analysis of existing expert systems (and MYCIN in particular), KADS focuses more broadly on the analysis of human problem solving and expert behaviour.

The need for a modelling approach to the analysis of data acquired from domain experts is almost a matter of general agreement (e.g. Clancey, 1985; Cooke, 1985; Wielinga and Breuker, 1986; Winston, 1986; Gaines, 1987a). This analysis should produce a high level functional specification, or conceptual model, for the expert system, according to the KADS developers, who believe that:

“... conceptual models of expertise in different domains—but representing similar tasks—may have the same higher level structure. Thus a conceptual model can be abstracted from its domain specific content and function as a skeletal model for the conceptual model of expertise in another domain, but with a similar task.

Such an abstracted model is called [an] interpretation model ... An interpretation model is an intermediate between ‘what to look for in the data’ and the conceptual model that is to be constructed. In the KADS methodology, a conceptual model is a product; an interpretation model is a tool ...” (Breuker et al., 1987, p. 7)

“Interpretation models are generic models of the problem solving processes in classes of tasks and types of domains . . . An interpretation model is a kind of catalogue of types of ingredients the knowledge engineer can look for in the data, and thus functions as an organiser that provides coherence to these data.” (Wielinga and Breuker, 1986, p. 22)

The idea of the interpretation model is central to the KADS methodology and expresses the most important development of the classical procedure for building expert systems, as outlined by Buchanan et al. (1983).

As noted earlier, the creation of the conceptual model occurs at Newell’s (1982) “knowledge level” or the epistemological level identified by various authorities as the missing level in developing expert systems (see Breuker et al., 1987, p. 9). Interpretation of verbal data at this level is important because the gap to the implementation level is too wide to be bridged in one step, without risking at least an unstructured and opaque system as a result.

Breuker et al. (1987) distinguish four layers of description of expert knowledge: static knowledge describing concepts and relations, knowledge of different types of inference that can be made, knowledge representing tasks and goals, and strategic knowledge. Their use of the last three levels as a vehicle to generalize from one particular task domain to classes of tasks is the basis of KADS. A summary of the Breuker et al. model of task analysis may be found in Wilson (1987).

Following the process of defining the functions of the prospective expert system and the role of its users, the second stage of knowledge acquisition involves specification of the tasks the system is to perform and construction of a conceptual model. From data acquired by the various techniques outlined in 4.b above, both static-knowledge and inference-knowledge (which includes tasks and strategies) may be deduced—the latter particularly from “expertise-in-action”, especially protocols—and an interpretation model selected or constructed and then filled out with data (Breuker and Wielinga, 1984).

Viewing the process of knowledge base creation “as a mapping from an understanding of behaviour in the real world to the description of the form of an artifact (i.e. as a category of engineering)”, Breuker et al. (1987, p. 4) assert that tools such as KEE assist in “mapping backwards” i.e. in the composition of elements of the artificial world to create the desired behaviour. The alternative involves mapping the real world onto the artificial via a model, or more closely, constructing a description in the analysis space (which includes all potential real world models), mapping it onto a description in the design space (which includes all system models) and from there to an implementation. Breuker et al. (1987) see the principal aim of knowledge acquisition research, from an engineering perspective, as to understand how these spaces are structured and points within them defined. Breuker et al. categorize three approaches to this.

The first, exemplified by provision of techniques such as repertory grid analysis, provides detailed information about individuals’ construct systems and particular domains, but according to Breuker et al. (1987) fails in the wider aim. Boose (1985, 1986b) recognizes that a single repertory grid cannot account for the links between multiple domains, although Shaw and Gaines (1986) see no intrinsic limitation to the range of application of techniques based on personal construct psychology. This approach may ultimately not be incompatible with the more explicitly model-based philosophy of Breuker et al. (although they do not share quite the same concept of abstraction as discussed by Darden (1987)); particularly when the work on combination of expertise from different sources (e.g. Boose, 1986b; Boose, Shema and Bradshaw, 1988) is further developed. The application of personal construct techniques to facilitate rapid prototyping may have biased the opinion of Breuker et al. against this approach, in view of their well-expressed criticisms of rapid prototyping (e.g. that the knowledge represented is conceptually and epistemologically unsound (Breuker and Wielinga, 1983, 1984; see also Johnson, Zualkernan and Garber, 1987)).

The second strategy is the one exemplified by single-model systems such as ROGET; it involves selecting a particular implemented system (i.e. a point within the design space), building a knowledge acquisition tool to define the mapping from the analysis space and then experimenting with the tool to build systems in what are believed to be related domains. In fact, although ROGET was

based on MYCIN there is no evidence that the more recent systems like SALT were founded upon a particular implemented system, so this characterization of Breuker et al. may be out of date now.

The third approach—which Breuker et al. view as complementary to the second—involves attempting to characterize “the languages adequate to describe a significant number of points within these spaces”; they do not underestimate the difficulties in this approach, and admit that “at present we are far from having a coherent view” (Breuker et al., 1987, p. 13).

Breuker et al. regard KADS as combining the generality of this language-based approach with the strength of a single-model system. This strength is conferred by allowing some constraint on design at the upper layers of their four-layer model, in the form of fixed task structures such as goal-trees, with corresponding task decomposition.

Although a degree of analogy may have occurred earlier to the reader between the notion of an interpretation model and an expert system shell, Breuker et al. (1987, p. 22) assert that the degree of this design constraint “is obviously of a different order of magnitude compared with conceptualising on the basis of an implementation formalism”.

Applications of the KADS methodology have been reported for commercial loans assessment (Wielinga and Breuker, 1986); export credit guarantee underwriting (Breuker et al., 1987); printed circuit board manufacturing (Hayball and Barlow, 1988); a natural language access system (Jansen-Winkel, 1988); text analysis (Moller, 1988); and incorporation in an expert system tool (Karbach and Tong, 1988).

KEATS (The knowledge engineer's assistant)

The developers of KEATS (Motta et al., 1986; 1988) aim to provide support for the knowledge engineer with a tightly integrated set of tools operating in an essentially bottom-up way, particularly in the early stages. A problem with the original KADS methodology has been the limited assistance provided in structuring the data, although it is a large and continuing project and this problem is being addressed (e.g. Anjewierden, 1987; Hayward, Wielinga and Breuker, 1987).

KEATS incorporates four subsystems: CREF, a Cross-Referencing Editing Facility that helps the knowledge engineer to structure interview protocols; KDL, a Knowledge Description Language which provides the user with a frame-like representation; GIS, an intelligent Graphical Interface System providing a graphical front-end (a “sketchpad”) to the KDL knowledge base; and COPS, a Context Oriented Production System rule interpreter which interacts with the KDL knowledge base, provides forward and backward chaining and allows rules to be factored into contexts to make integration of different rule bases easier and to speed up rule selection. The KEATS system assists the understanding of what is often initially a large unstructured collection of data. It allows the knowledge engineer to build an abstract conceptual model of the domain without forcing early decisions about the concrete implementation formalism. The input-driven user interface provided by GIS is a particularly useful model-building tool, helping the knowledge engineer to visualize the domain and communicate easily with the expert. One handicap to the wider use of KEATS may prove to be the Symbolics machine to which it is presently restricted, compared with the KADS system for example which runs on any hardware supporting UNIX.

Future development of KEATS will focus particularly on model-based support, including the use of coding sheets (questionnaires based on interpretation models), and the development of a library of task models similar to KADS interpretation models (Motta et al., 1988). Thus there is some convergence of approach between the KEATS and KADS teams; reporting on the design process in KADS, which converts the output of the analysis phase into a design model, Schreiber et al. (1987) tentatively suggest it is a combination of top-down and bottom-up i.e. “a middle out” or bidirectional approach. Breuker and Wielinga (1987b) and Anjewierden (1987) also make explicit the inclusion of bottom-up tools within the second version of the KADS system. These KADS Power Tools (KPT) are currently under revision (Hickman, pers. comm.). The present (June 1988) prototype implementation, according to KADS participants SD-Scicon, contains facilities for protocol

analysis (similar to the CREF and GIS subsystems of KEATS), glossary building, building taxonomies of domain concepts and choosing and editing interpretation models.

KRITON

KRITON (Diederich, 1987; Diederich, Ruhmann and May, 1987a; Diederich et al., 1987b; Linster, 1988) is a general purpose system (named after a disciple of Socrates) with a similar approach to KADS and more especially KEATS (see Breuker et al. (1987) and Motta et al. (1988) for comparisons). Both Breuker et al. and Linster (1988) consider KRITON could be incorporated into the KADS methodology. KRITON supports automated interviewing, combining repertory grid, forward scenario simulation and laddering techniques; protocol analysis; semantic text analysis; and machine learning via a learning-by-example system called COLE. The output is translated into an intermediate knowledge representation system involving a description language in the form of a semantic net, and a set of operator-argument structures derived from the protocol analysis. However, the process of translation is by no means clear, which is a criticism that can also be levelled at the KADS system: the evolution of their "interpretation models" is also inadequately described. Related to this criticism also is the construct introduced by Diederich et al. (1987a,b) which they call an AKB (Acquisition Knowledge Base), the purpose of which is "for better guidance of the ongoing elicitation process". It is unclear whether these AKBs are in fact intended as "predefined declarative deep models of a domain" or as domain-independent KADS Interpretation Models (see Diederich, Ruhman and May, 1987a).

Thus it is not yet sufficiently clear how precisely the intermediate representation or interpretation model in either the KRITON or KADS system is constructed. KEATS has also been included in this criticism, in that inadequate support is offered for a conceptual analysis stage driven by the data. Johnson (1987) (see also chapter in Diaper (1989)) contends that, instead, some fundamental choices or presuppositions are encouraged (e.g. task models). Johnson argues that a Mediating Representation such as a Systemic Grammar Network is essential to support the process of deriving data driven concepts (qualitative data analysis) and to generate a description of the knowledge (a "grammar") which can be tested against the original data (see also Johnson, 1985). A picture such as the GIS within KEATS produces, no matter how detailed according to Johnson (1987), cannot be a grammar or as adequate as a grammar.

All these Intermediate or Mediating Representations can be viewed as lying on a spectrum from humanistic to computer-like (Young and Gammack, 1987), and no dogmatic assertion that one is universally superior seems appropriate at present. It is felt by this author that it is too early to say whether the still-evolving KEATS project will produce competent expert systems or if it will develop task structures and representations founded upon valid models of expertise. Presently KEATS lacks an explicit underlying methodology, such as KADS possesses. The possibilities represented by integration of machine learning into knowledge acquisition (e.g. Blythe, 1987; Diederich et al., 1987b), or as in the balanced cooperative modelling advocated by Morik (1987a,b), deserve further attention.

5.c Major conclusions and issues for section 5

From the research literature, at least, there is no doubt that a trend exists towards the use of tools for knowledge acquisition. This review has concentrated on tools which offer the benefits of increased productivity in expert system building, largely through aiming to dispense with the knowledge engineer as intermediary and allowing the expert to interact directly with the computer; whilst also being the fruits of intellectually vigorous research programs. Even so, only representative examples have been chosen from the range recently reported.

There is a clear trend towards inclusion of a graphical interface to ease the task of experts with limited knowledge of computers in encoding their domain knowledge. Additional instances include

the Honeywell-proprietary domain-specific acquisition tool KLAMShell (Cochran, 1988), and the implementation of Sowa's conceptual graphs (Sowa, 1984) by Eklund and Kuczora (1988) as intended to replace production rules. This latter and associated work at Brighton Polytechnic (UK) by Esfahani and Teskey (1987, 1988) and by Kellett and Esfahani (1988) shows strong parallels with the KEATS project. Other parallels among the KADS, KEATS and KRITON language-based systems have already been noted.

It may be that for some experts and some domains or tasks it is difficult even with graphical representations to design a way of displaying the knowledge which is comprehensible to the expert. For such expertise, tools based upon induction from examples or cases may be appropriate (Gruber and Cohen, 1987).

A historical development path has been traced in the development from TEIRESIAS through ROGET and MORE to MOLE and SALT of systems based upon particular conceptual models. There are two aspects which both these tools and the language-based systems such as KADS share, to the extent that a process of convergent evolution might be already discernable. The first is a developing consensus that analysis of a domain should take place at the knowledge level rather than the symbol level, although how tasks should be described at the knowledge level is a more open question. The second interesting possibility lies in the degree of similarity apparent between KADS Interpretation Models (and intended similar developments within KEATS and KRITON) and McDermott's role-limiting methods (McDermott, *in press*). "An interpretation model is a kind of catalogue of types of ingredients the knowledge engineer can look for in the data, and thus functions as an organiser that provides coherence to these data" (Wielinga and Breuker, 1986, p. 22); "... role-limiting methods ... both have a quite broad scope of applicability and provide substantial help in specifying what knowledge needs to be collected to perform a particular task and how that knowledge can be appropriately encoded" (McDermott, *in press*). As the two programs are fleshed out it will be interesting to compare KADS' intended "library of interpretation models" with McDermott and co-workers' "taxonomy of problem-solving methods".

The distinction often made between domain-specific tools and domain-independent tools is now beginning to seem less helpful. While knowledge acquisition tools such as OPAL are tied fairly closely to a particular domain task, others like SALT are suitable for a family of closely-related tasks but not necessarily in similar domains. Tools having the broader overall scope of the language-based tools like KADS are intended for a variety of tasks and domains. It would seem more appropriate to characterize tools by their degree of applicability to tasks rather than domains, in the first instance, or alternatively by the problem solving method they employ (less useful at present), and to regard projects like KADS as workbench systems capable of producing the sort of tools of which MOLE and SALT are individual examples.

A different approach to workbench systems is represented by AQUINAS (see especially Kitto and Boose (1987b)), although at present it is confined to analysis problems and the heuristic classification problem solving method (which is the most completely described (Clancey, 1985) method and also the most widely used in knowledge acquisition tools (Boose, 1988)). The limitation on the scope of tasks AQUINAS can include may arise from its grounding in the repertory grid technique. The accent appears to be on the speed of knowledge base construction which is possible with AQUINAS through rapid prototyping techniques.

The generic task approach of Chandrasekaran's group at Ohio State University is based on the anti-traditional view that "knowledge representation and use cannot be separated" (Chandrasekaran, 1987, p. 1184) and seemingly rejects the KADS idea that problem analysis and solution design are distinct activities. A generic task is "an elementary generic combination of a problem, representation, and inference strategy about concepts" (Bylander and Chandrasekaran, 1987, p. 235). Diagnosis was initially seen (Chandrasekaran, 1983) as such a task, but later (Chandrasekaran, 1986, 1987) as a compound task which can be decomposed into more "atomic" building blocks. One of these is termed "hierarchical classification", but it seems that this in turn involves several semi-autonomous problem-solving agents (Chandrasekaran and Goel, 1988): how far should decomposition proceed? It is thus unclear whether generic tasks are more analogous to KADS inter-

pretation models or to single knowledge sources within them; some generic tasks are apparently more complex than others. A consensus has yet to develop about what in fact is the “right” level of abstraction for knowledge acquisition.

The real need in facilitating veridical knowledge acquisition is for support in epistemology (Smart, 1987): we are still some way from a well-defined, epistemologically-sound methodology. This is why further development of programmes such as the KADS project is particularly important: such a methodology is the aim, but the inadequate support presently given to the user to select an interpretation model illustrates the need. Here “methodology” should not be construed as something rigid which (like the Law of the Medes and the Persians) “altereth not”, neither according to the domain nor the humans inquiring of it. Rather, following Banathy (1988), it is contended that within an epistemologically-based structured methodology, involving successive levels of abstraction and stages of refinement, techniques and methodological tools or approaches should be selected to suit the context and participants of the enquiry.

Finally, it may be necessary to recognize limitations upon the machine-aided approach to knowledge acquisition. The goal of genuinely accurate replication of the expert’s problem-solving ability in a computer program is regarded by some as unreachable: Winograd and Flores (1986), for example, suggest that “it is fruitless to search [in the expert’s mind] for a full formalisation of the pre-understanding that underlies all thought and action . . . human cognition includes the use of representations, but is not based on representation” (p. 99). (This is also of course a criticism of expectations of expert systems in general). Another critic (Kornell, 1987) has examined three representative tools—RULEMASTER (Michie et al., 1984), MOLE (Eshelman et al., 1987) and AQUINAS (Boose and Bradshaw, 1987)—in the light of human thought processes. Kornell distinguishes two ways of thinking: formal thought, exemplified by predicate calculus, and narrative thought, exemplified by metaphors, analogies and gestalts. Formal thought deals with closed, well-defined worlds; narrative thought seeks to construct meaning in open, dynamic systems, i.e. in the real world. Expert systems which aim to model the problem solving expertise of human experts must therefore model narrative thinking. This involves a combination of facts and heuristics with *patterns of reasoning*, represented for example by scripts, and *kinds of reasoning*, such as reasoning by analogy. Kornell argues that whilst knowledge acquisition tools can represent patterns of reasoning, they offer no help in identifying and acquiring patterns of reasoning; the existence of different kinds of reasoning is not even acknowledged by such systems. As Morik (1987b) puts it, “the model to be represented in the expert system is presupposed. The system supports the encoding of a given model, but not the building of the model itself.”

6 Overall summary and conclusions

Very little recognition has been accorded until recently to the need for a systematic approach to knowledge acquisition. The literature of systems analysis and design deserves much closer attention than has been apparent in this field. This is particularly true now that the desirability of following a structured and properly documented approach to expert system building is gradually gaining recognition over the early “rapid prototyping” paradigm. An example demonstrating the evolution of a structured methodology from fast prototyping is the paper by Thomas (1988). The NIAM methodology (Twine, 1988) for conceptual schema design in database research is another instance of a development outside the expert systems world which could assist knowledge analysis.

Unjustified faith has often been placed in textbook knowledge, and in what experts say they do during informal interviews. Psychological techniques have been described which, chosen with care to suit the domain and, when identifiable, the task, as well as the expert and the knowledge engineer, have been found to be helpful in knowledge acquisition. From analysis of the data acquired, the aim should be to create a conceptual model of the expert’s knowledge and reasoning. For conducting expert review and validation of this, the value of using mediating or intermediate representations (which may include paper models) is increasingly recognized.

It is also noted that some psychological techniques are also very time consuming, particularly

those like verbal protocol analysis which seek to elicit the expert's conceptual structure rather than just straightforward facts about the domain. This is one reason why a strong trend exists towards eliminating the knowledge engineer and getting the expert to work directly with the computer. Particularly if a simple shell based on the production rule formalism and incorporating limited inferencing facilities is used, there is a trade-off involved between speed of expert system production and accuracy of representation. The more ambitious the intention to match the expert's own problem-solving expertise accurately, the greater the need for care in selecting a tool based on the correct problem-solving method: more support is needed at this point. Several research programs and tools founded on particular theories of knowledge acquisition have been described and some commonalities identified, in particular the need to formulate a conceptual model of the expertise at the knowledge level.

This paper stresses the importance of a thorough domain orientation process, including an initial feasibility study, which makes minimal demands upon the expert. Defining user and organizational requirements at this stage also affects knowledge acquisition strategy. Knowledge acquisition is a social process, which will be both better understood and more successful when the contribution of social scientists to an understanding of knowledge and how people communicate is more widely recognized.

Acknowledgements

This paper is based on part of a dissertation submitted to Loughborough University in 1987. Anne Morris provided guidance and useful advice for which I am grateful. I would also like to particularly thank John Gammack for his assistance and several helpful discussions. The paper has been improved by comments and suggestions from three referees and the editor, whose help I very much appreciate. The inaccuracies and inadequacies which remain evident to the reader are no doubt a product of my present ignorance: I shall be grateful for enlightenment.

References

- Adelman, L (1987). "Toward considering psychological measurement issues when developing expert systems" *Proceedings of the 1987 IEEE International Conference on Systems, Man and Cybernetics* (Cat. No. 87CH2503-1), Alexandria, VA, USA, 20-23 October 1987 (Piscataway, NJ: IEEE), 2 pp. 1044-1048.
- Agnew, N McK, Brown, J L and Lynch, J G (1986). "Extending the reach of knowledge engineering" *Future Computing Systems* 1 (2) 115-141.
- Aleksander, I (1984). *Designing intelligent systems: an introduction*, London: Kogan Page.
- Alvey, P L and Greaves, M F (1987). "Observations on the development of a high performance system for leukaemia diagnosis" In: *Research and development in expert systems III*, Bramer M A (Ed.). Cambridge: Cambridge University Press, pp. 99-110.
- Anjewierden, A (1987). "Knowledge acquisition tools" *AI Communications* 0 (1) 29-38 (August 1987). Also in: *Proceedings of the First European Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Reading, UK, 2-3 September 1987, Reading: Reading University, section E2, pp. 1-12.
- Bailey, W A and Kay, E J (1986). "Toward the standard analysis of verbal data" *Proceedings of the 1986 IEEE International Conference on Systems, Man and Cybernetics* (Cat. No. 86CH2364-8), Atlanta, GA, USA, 14-17 October 1986, New York: IEEE, 1 pp. 582-587.
- Bainbridge, L (1979). "Verbal reports as evidence of the process operator's knowledge" *International Journal of Man-Machine Studies* 11 411-436.
- Bainbridge, L (1986). "Asking questions and accessing knowledge" *Future Computing Systems* 1 (2) 143-149.
- Banathy, B H (1988). "Matching design methods to system type" *Systems Research* 5 (1) 27-34.
- Barnard, P, Wilson, M and Maclean, A (1986). "The elicitation of system knowledge by picture probes" In: *Human factors in computing systems*. Mantei, M and Orbeton, P (Eds.). Proceedings of CHI '86 Conference, Boston, MA, USA, 13-17 April 1986, New York, NY: ACM, pp. 235-240.
- Barr, A (1983). "AI: cognition as computation" In: *The study of information: interdisciplinary messages*. Machlup, F and Mansfield, U (Eds.). New York, NY: John Wiley, pp. 237-262.
- Basden, A (1984). "On the application of expert systems" In: *Developments in expert systems*, Coombs, M J (Ed.). London: Academic Press, pp. 59-75.
- Bell, M Z (1985). "Why expert systems fail" *Journal of the Operational Research Society* 36 (7) 613-619.

- Bennett, J S (1985). "ROGET: a knowledge-based system for acquiring the conceptual structure of a diagnostic expert system" *Journal of Automated Reasoning* 1 49–74.
- Berry, D C (1987). "The problem of implicit knowledge" *Expert Systems* 4 (3) 144–151.
- Berry, D C and Broadbent, D E (1986). "Expert systems and the man-machine interface" *Expert Systems* 3 (4) 228–231.
- Berry, D C and Broadbent, D E (1987). "Expert systems and the man-machine interface. Part two: the user interface" *Expert Systems* 4 (1) 18–28.
- Bloomfield, B P (1986). "Capturing expertise by rule induction" *Knowledge Engineering Review* 1 (4) 30–36.
- Bloomfield, B P (1988). "Expert systems and human knowledge: a view from the sociology of science" *AI & Society* 2 (1) 17–29.
- Blythe, J (1987). "Machine learning as an aid to knowledge acquisition" Unpublished paper presented at FringES '87, British Computer Society Expert Systems '87 Conference, Brighton, UK, 14–17 December 1987.
- Blythe, J, Needham, D and Corsi, P (1987). "An experimental protocol for the acquisition of examples for learning" *Proceedings of the first European Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Reading, UK, 2–3 September 1987, Reading: Reading University, section F1.
- Bobrow, D G, Mittal, S and Stefik, M J (1986). "Expert systems: perils and promises" *Communications of the ACM* 29 (9) 880–894.
- Boose, J H (1985). "A knowledge acquisition program for expert systems based on personal construct psychology" *International Journal of Man-Machine Studies* 23 495–525.
- Boose, J H (1986a). *Expertise transfer for expert system design*, Amsterdam: Elsevier.
- Boose, J H (1986b). "Rapid acquisition and combination of knowledge from multiple experts in the same domain" *Future Computing Systems* 1 (2) 191–216.
- Boose, J H (1988). "A research framework for knowledge acquisition techniques and tools" In: *Proceedings of the European Knowledge Acquisition Workshop (EKAW '88)*, Boose, J H, Gaines, B R and Linster, M (Eds.). Sankt Augustin, West Germany: GMD, pp. 10.1–10.13.
- Boose, J H and Bradshaw J M (1987). "Expertise transfer and complex problems: using AQUINAS as a knowledge-acquisition workbench for knowledge based systems" *International Journal of Man-Machine Studies* 26 3–28.
- Boose, J H, Gaines, B R and Linster, M (Eds.) (1988). *Proceedings of the European Knowledge Acquisition Workshop (EKAW '88)*, Bonn, West Germany, 19–23 June 1988, Sankt Augustin, West Germany: GMD.
- Boose, J H, Shema, D B and Bradshaw, J M (1988). "Recent progress in AQUINAS: a knowledge acquisition workbench. In: *Proceedings of the European Knowledge Acquisition Workshop (EKAW '88)*, Boose, J H, Gaines B R and Linster, M (Eds.). Sankt Augustin, West Germany: GMD, pp. 2.1–2.15.
- Boothroyd, D (1986). "Shells the key to success" *intelligence* 3 2.
- Brachman R J et al. (1983). "What are expert systems?" In: *Building expert systems*, Hayes-Roth, F, Waterman, D A and Lenat, D B (Eds.). Reading, MA: Addison-Wesley, pp. 31–57.
- Bramer, M A (Ed.) (1985). "Research and development in expert systems" *Proceedings of the Fourth Technical Conference of the British Computer Society Specialist Group on Expert Systems*, University of Warwick, UK, 18–20 December 1984, Cambridge: Cambridge University Press.
- Bramer, M A (1986). "Codifying rules of thumb" *The Times Higher Education Supplement*, 14 March 1986.
- Bramer, M A (Ed.) (1987). "Research and development in expert systems III" *Proceedings of Expert Systems '86, the Sixth Annual Technical Conference of the British Computer Society Specialist Group on Expert Systems*, Brighton, UK, 15–18 December 1986, Cambridge: Cambridge University Press.
- Bramer, M A (1987a). "Expert systems in business: a British perspective" *First International Symposium on Artificial Intelligence and Expert Systems*, Berlin, Germany, May 1987.
- Bramer, M A (1987b). "Automatic induction of rules from examples: a critical analysis of the ID3 family of rule induction systems" *Proceedings of the first European Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Reading UK, 2–3 September 1987, Reading: Reading University, section F2, pp. 1–31.
- Breuker, J A and Wielinga, B J (1983). "Analysis techniques for knowledge based systems. Part 2: methods for knowledge acquisition." Report 1.2, Esprit Project 12 (Memorandum 13 of the Research Project "The Acquisition of Expertise"), December 1983, Amsterdam: University of Amsterdam.
- Breuker, J A and Wielinga, B J (1984). "Techniques for knowledge elicitation and analysis" Report 1.5, Esprit Project 12 (Memorandum 28 of the Research Project "The Acquisition of Expertise"), July 1984, Amsterdam: University of Amsterdam.
- Breuker, J A and Wielinga, B J (1987a). "Use of models in the interpretation of verbal data" In: *Knowledge acquisition for expert systems: a practical handbook*, Kidd, A L (Ed.), New York, NY: Plenum, pp. 17–44.
- Breuker, J A and Wielinga, B J (1987b). "Knowledge acquisition as modeling expertise: the KADS methodology" *Proceedings of the first European Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Reading, UK, 2–3 September 1987, Reading: Reading University, section B1.

- Breuker, J A et al. (1987). "Model-driven knowledge acquisition: interpretation models" Deliverable task A1, Esprit Project 1098, Amsterdam: University of Amsterdam.
- Buchanan, B G (1986). "Expert systems: working systems and the research literature" *Expert Systems* 3 (1) 32–51.
- Buchanan, B G and Duda, R O (1983). "Principles of rule-based expert systems" *Advances in Computers* 22 163–216.
- Buchanan, B G et al. (1983). "Constructing an expert system" In: *Building expert systems*, Hayes-Roth, F, Waterman, D A and Lenat, D B (Eds.). Reading, MA: Addison-Wesley, pp. 127–167.
- Burton, A M and Shadbolt, N R (1987a). "Knowledge engineering" Technical report 87-2-1, Artificial Intelligence Group, Department of Psychology, University of Nottingham.
- Burton, A M and Shadbolt, N R (1987b). "A formal evaluation of knowledge elicitation techniques for expert systems: domain 1" In: *Knowledge acquisition for engineering applications*, Pavelin, C J and Wilson, M D (Eds.). pp. 20–28.
- Butler, A R and Chamberlin, G (1988). "The ARIES Club—experience of expert systems in insurance and investment" In: *Research and development in expert systems IV*, Moralee, D S (Ed.). Cambridge: Cambridge University Press, pp. 246–257.
- Bylander, T and Chandrasekaran, B (1987). "Generic tasks for knowledge-based reasoning: the 'right' level of abstraction for knowledge acquisition" *International Journal of Man-Machine Studies* 26 231–243.
- Candy, L and Lunn, S (1988). "Design strategies for expert systems: a case study" *Conference on Human and Organisational Issues of Expert Systems*, Stratford-on-Avon, UK, 4–6 May 1988 (sponsored by ICL and The Ergonomics Society).
- Carbonell, J G (1985). "Derivational analogy: a theory of reconstructive problem solving and expertise acquisition" Report CMU-CS-85-115, Department of Computer Science, Carnegie-Mellon University, Pittsburgh, PA, USA.
- Chandrasekaran, B (1983). "Towards a taxonomy of problem-solving types" *AI Magazine* 4 (1) 9–17.
- Chandrasekaran, B (1986). "Generic tasks in knowledge-based reasoning: high-level building blocks for expert system design" *IEEE Expert* 1 (3) 23–30.
- Chandrasekaran, B (1987). "Towards a functional architecture for intelligence based on generic information processing tasks" *Proceedings of the Tenth International Joint Conference on Artificial Intelligence (IJCAI-87)*, Milan, Italy, 23–28 August 1987. Los Altos, CA: Morgan Kaufman, 2 1183–1192.
- Chandrasekaran, B and Goel, A (1988). "From numbers to symbols to knowledge structures: artificial intelligence perspectives on the classification task" *IEEE Transactions on Systems, Man and Cybernetics* 18 (3) 415–424.
- Clancey, W J (1983). "The epistemology of a rule-based expert system—a framework for explanation" *Artificial Intelligence* 20 215–251.
- Clancey, W J (1985). "Heuristic classification" *Artificial Intelligence* 27 289–350.
- Clancey, W J and Letsinger, R (1981). "NEOMYCIN: reconfiguring a rule-based expert system for application to teaching" *Proceedings of the Seventh International Joint Conference on Artificial Intelligence (IJCAI-81)*, Vancouver, Canada, 24–28 August 1981, pp. 829–836.
- Cleaves, D A (1987). "Cognitive biases and corrective techniques: proposals for improving elicitation procedures for knowledge-based systems" *International Journal of Man-Machine Studies* 27 155–166.
- Cochran, E L (1988). "KLAMSHell: a domain-specific knowledge acquisition shell" *AAAI-88 Workshop on Integration of Knowledge Acquisition and Performance Systems*, St. Paul, MN, USA, 21 August 1988.
- Collins, H M, Green, R H and Draper, R C (1985). "Where's the expertise? Expert systems as a medium of knowledge transfer" In: *Expert Systems 85* Merry, M (Ed.). Cambridge: Cambridge University Press pp. 323–334.
- Cooke, N M (1985). "Modelling human expertise in expert systems" Memoranda in Computer and Cognitive Science MCCS-85-12, Computing Research Laboratory, New Mexico State University, Las Cruces, NM USA.
- Cooke, N M and McDonald, J E (1986) "A formal methodology for acquiring and representing expert knowledge" *Proceedings of the IEEE* 74 (10) 1422–1430.
- Cooke, N M and McDonald, J E (1987). "The application of psychological scaling techniques to knowledge elicitation for knowledge-based systems" *International Journal of Man-Machine Studies* 26 533–550.
- Cookson, M J, Holman, J G and Thompson, D G (1985). "Knowledge acquisition for medical expert systems: a system for eliciting diagnostic decision making histories" In: *Research and development in expert systems* Bramer, M A (Ed.). Cambridge: Cambridge University Press, pp. 113–116.
- Coombs, M J (Ed.) (1984). *Developments in expert systems*, London: Academic Press.
- Coombs, M J and Alty, J L (1984). "Expert systems: an alternative paradigm" In: *Developments in expert systems* Coombs, M J (Ed.). London: Academic Press, pp. 135–157.
- Currier, K (1986). "Softsync: Something old, something new" Report in *IEEE Expert* 1 (4) 83.
- Darden, L (1987). "Viewing the history of science as compiled hindsight" *AI Magazine* 8 (2) 33–41.

- Davies, M (1986). "A practical investigation of knowledge acquisition techniques" Report No. 72/86/R147/U, April 1986; Plessey Electronic Systems Research Limited, Romsey, Hants, UK.
- Davies, M and Hakiel, S (1988). "Knowledge harvesting: a practical guide to interviewing" *Expert Systems* 5 42–50.
- Davis, R (1979). "Interactive transfer of expertise: acquisition of new inference rules" *Artificial Intelligence* 12 121–157.
- Diaper, D (1986). "Identifying the knowledge requirements of an expert system's natural language processing interface" In: *People and computers: designing for usability*, Harrison, M D and Mark, A F (Eds.). Proceedings of the Second Conference of the British Computer Society Human Computer Interaction Specialist Group, University of York, UK, 23–26 December 1986. Cambridge: Cambridge University Press, pp. 263–280.
- Diaper, D (1987). "POMESS: a People Oriented Methodology for Expert System Specification" *Proceedings of the first European Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Reading, UK, 2–3 September 1987. Reading, UK: Reading University section D4.
- Diaper, D (Ed.) (1989). *Knowledge elicitation: principles, techniques and applications*, Chichester: Ellis Horwood, in press.
- Diaper, D (1989a). "Designing expert systems: from Dan to Beersheba" In: *Knowledge elicitation: principles, techniques and applications*, Diaper, D (Ed.). Chichester: Ellis Horwood, in press.
- Diederich, J (1987). "Knowledge-based elicitation" *Proceedings of the Tenth International Joint Conference on Artificial Intelligence (IJCAI-87)*, Milan, Italy, 23–28 August 1987. Los Altos, CA: Morgan Kaufmann, 1 201–204.
- Diederich, J, Ruhman, I and May, M (1987a). "KRITON: a knowledge-acquisition tool for expert systems" *International Journal of Man-Machine Studies* 26 29–40.
- Diederich, J, Linster, M, Ruhman, I and Uthman, T (1987b). "A methodology for integrating knowledge acquisition techniques" *Proceedings of the first European Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Reading, UK, 2–3 September 1987. Reading: Reading University, section E4.
- Dillard, J F and Mutchler, J F (1987). "Expertise in assessing solvency problems" *Expert Systems* 4 (3) 170–179.
- Dym, C L (1987). "Issues in the design and implementation of expert systems" (*AI EDAM*): *Artificial Intelligence for Engineering Design, Analysis and Manufacturing* 1 (1) 37–46.
- Easterby-Smith, M (1981). "The design, analysis and interpretation of repertory grids" In: *Recent advances in personal construct psychology*, Shaw, M L G (Ed.). London: Academic Press, pp. 9–30.
- Eden, C and Jones, S (1984). "Using repertory grids for problem construction" *Journal of the Operational Research Society* 35 (9) 779–790.
- Eklund, P and Kuczora, P (1988). "An implementation of conceptual graphs for knowledge based systems development" *Conference on Human and Organisational Issues of Expert Systems*, Stratford-on-Avon, UK, 4–6 May 1988 (sponsored by ICL and The Ergonomics Society).
- Ellman, J (1987). "Practical experience of knowledge acquisition for expert systems" In: *Knowledge acquisition for engineering applications*, Pavelin, C J and Wilson, M D (Eds.). pp. 39–43.
- Ericsson, K A and Simon, H A (1980). "Verbal reports as data" *Psychological Review* 87 (3) 215–251.
- Ericsson, K A and Simon, H A (1984). *Protocol analysis: verbal reports as data*, Cambridge, MA: The MIT Press.
- Esfahani, L and Teskey, F N (1987). "KET, a knowledge encoding tool" *Proceedings of the first European Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Reading, UK, 2–3 September 1987. Reading: Reading University, section E5.
- Esfahani, L and Teskey, F N (1988). "A self-modifying rule-eliciter" In: *Proceedings of the European Knowledge Acquisition Workshop (EKAW '88)*, Boose, J H, Gaines, B R and Linster, M (Eds.). Sankt Augustin, West Germany: GMD, pp. 16.1–16.16.
- Eshelman, L, Ehret, D, McDermott, J and Tan, M (1987). "MOLE: a tenacious knowledge-acquisition tool" *International Journal of Man-Machine Studies* 26 41–54.
- Evans, J St B T (1987) "Human biases and computer decision-making: a discussion of Jacob et al." *Behaviour and Information Technology* 6 (4) 483–487.
- Evans, J St B T (1988). "The knowledge elicitation problem: a psychological perspective" *Behaviour and Information Technology* 7 (2) 111–130.
- Feigenbaum, E A and McCorduck, P (1984). *The fifth generation*, London: Pan Books.
- Fenn, J A, Worden, R P., Foote, M H and Willison, R G (1986). "An expert assistant for electromyography" *Biomedical Measurement Information and Control* 1 (4) 210–214.
- Forsyth, R (1986). "Software review: Personal Consultant Plus" *Expert Systems* 3 (4) 244–247.
- Forsyth, R (1987). "Software review: Expertech Xi Plus" *Expert Systems* 4 (1) 48–51.
- Fox, J. (1984). "Doubts about induction" *SPL-Insight Bulletin* 2 (2) 32–36.

- Fox, J (1986). "Three arguments for extending the framework of probability" In: *Uncertainty in artificial intelligence*, Kanal, L N and Lemmer, J F (Eds.). Amsterdam: North-Holland, pp. 447–458.
- Fox, J, Myers, C D, Greaves, M F and Pegram, S (1987). "A systematic study of knowledge base refinement in the diagnosis of leukaemia" In: *Knowledge acquisition for expert systems: a practical handbook*, Kidd, A. L. (Ed.). New York, NY: Plenum, pp. 73–90.
- Fransella, F and Bannister, D (1977). *A manual for repertory grid technique*, London: Academic Press.
- Gaines, B R (1987a). "Foundations of knowledge engineering" In: *Research and development in expert systems III*, Bramer, M A (Ed.). pp. 13–24.
- Gaines, B R (1987b). "An overview of knowledge acquisition and transfer" *International Journal of Man-Machine Studies* 26 453–472.
- Gammack, J G (1987a). "Modelling expert knowledge using cognitively compatible structures" *Third International Expert Systems Conference*, London, UK, 2–4 June 1987. Oxford: Learned Information, pp. 191–200.
- Gammack, J G (1987b). "Formalising implicit domain structure" In: *Knowledge acquisition for engineering applications*, Pavelin, C J and Wilson, M D (Eds.). pp. 29–38.
- Gammack, J G (1987c). "Different techniques and different aspects on declarative knowledge" In: *Knowledge acquisition for expert systems: a practical handbook*, Kidd, A L (Ed.). New York, NY: Plenum, pp. 137–164.
- Gammack, J G and Young, R M (1985). "Psychological techniques for eliciting expert knowledge" In: *Research and development in expert systems*, Bramer, M A (Ed.). Cambridge: Cambridge University Press, pp. 105–112.
- Garg-Janardan, C and Salvendy, G (1987). "A conceptual framework for knowledge elicitation" *International Journal of Man-Machine Studies* 26 521–531.
- Garg-Janardan, C and Salvendy, G (1988). "The contributions of cognitive engineering to the design and use of expert systems" *Behaviour and Information Technology* 7 323–342.
- Gevarter, W B (1987). "The nature and evaluation of commercial expert system building tools" *Computer* 20 (5) 24–41.
- Gleeson, J F J and West, M L J (1988). "CLINTE: Coopers & Lybrand International Tax Expert System" In: *Research and development in expert systems IV*, Moralee, D S (Ed.). Cambridge: Cambridge University Press, pp. 18–31.
- Gold, J (1986). "Do-it-yourself expert systems" *Computer Decisions* 18 (2) 76–81.
- Goodall, A (1985). *The guide to expert systems*, Oxford: Learned Information.
- Gregory, D (1986). "Delimiting expert systems" *IEEE Transactions on Systems, Man and Cybernetics SMC-16* (6) 834–843.
- Grover, M D (1983). "A pragmatic knowledge acquisition methodology" *Proceedings of the Eighth International Joint Conference on Artificial Intelligence (IJCAI-83)*, Karlsruhe, West Germany, 8–12 August 1983, pp. 436–438.
- Gruber, T R and Cohen, P R (1987). "Design for acquisition: principles of knowledge-system design to facilitate knowledge acquisition" *International Journal of Man-Machine Studies* 26 143–159.
- Hart, A (1985a). "Experience in the use of an inductive system in knowledge engineering" In: *Research and development in expert systems*, Bramer, M A (Ed.). Cambridge: Cambridge University Press, pp. 117–126.
- Hart, A (1985b). "The role of induction in knowledge elicitation" *Expert Systems* 2(1) 24–28.
- Hart, A (1985c). "Knowledge elicitation: issues and methods" *Computer-Aided Design* 17 (9) 455–462.
- Hart, A (1986). "Knowledge acquisition for expert systems", London: Kogan Page.
- Hart, A (1987a). "Automatic knowledge generation: possibilities and restrictions" In: *Knowledge acquisition for engineering applications*, Pavelin, C J and Wilson, M D (Eds.). pp. 56–67.
- Hart, A. (1987b). "Role of induction in knowledge elicitation" In: *Knowledge acquisition for expert systems: a practical handbook*, Kidd, A L (Ed.). New York, NY: Plenum, pp. 165–190.
- Hart, A (1988). *Expert systems for managers*, London: Kogan Page.
- Hartley, R (1981). "How expert should an expert system be?" *Proceedings of the Seventh International Joint Conference on Artificial Intelligence (IJCAI-81)*, Vancouver, Canada, 24–28 August 1981, pp. 862–867.
- Hawkins, D (1983). "An analysis of expert thinking" *International Journal of Man-Machine Studies* 18 (1) 1–47.
- Hayball, C and Barlow, D (1988). "Skills support in the ICL Kidsgrove Bonding Shop—a case study in the application of the KADS methodology" *Conference on Human and Organisational Issues of Expert Systems*, Stratford-on-Avon, UK, 4–6 May 1988 (sponsored by ICL and The Ergonomics Society).
- Hayes-Roth, F, Waterman, D A and Lenat, D B (Eds.) (1983). *Building expert systems*, Reading, MA: Addison-Wesley.
- Hayward, S A (1985). "Is a decision tree an expert system?" In: *Research and development in expert systems*, Bramer, M A (Ed.). Cambridge: Cambridge University Press, pp. 185–192.
- Hayward, S A (1987). "Methodology: analysis and design for knowledge based systems". Esprit Project 1098 Report STC-Y-RR-001(3.0). Harlow: STC.

- Hayward, S A, Wielinga B J, and Breuker, J A (1987). "Structured analysis of knowledge" *International Journal of Man-Machine Studies* 26 487-498.
- Hewett, J (1987). "Commercial progress in Europe" *Third International Expert Systems Conference*, London, UK, 2-4 June 1987, Oxford: Learned Information, pp. 159-164.
- Hoffman, R R (1987). "The problem of extracting the knowledge of experts from the perspective of experimental psychology" *AI Magazine* 8 (2) 53-67.
- Inder, R, Wai, H C P and Fraser, J (1987). "Towards a methodology for incremental system development" In: *Knowledge acquisition for engineering applications*, Pavelin, C J and Wilson, M D (Eds.). pp. 44-45.
- Jackson, A (1985). "Teaching an American to play cricket" *Datalink*, 28 October 1985, p. 14.
- Jackson, A (1986a). "The cost of coconuts" *Datalink*, 6 January 1986, p. 9.
- Jackson, A (1986b). "A little learning is a dang'rous thing" *Datalink*, 3 February 1986, p. 10.
- Jackson, A (1986c). "Inadmissible evidence" *Datalink*, 7 April, 1986, p. 10.
- Jacob, V S, Gaultney, L D and Salvendy, G (1986). "Strategies and biases in human decision-making and their implications for expert systems" *Behaviour and Information Technology* 5 (2) 119-140.
- Jansen-Winkeln, R M (1988). "WASTL: an approach to knowledge acquisition in the natural language domain" In: *Proceedings of the European Knowledge Acquisition Workshop (EKAW '88)*, Boose, J H, Gaines, B R and Linster, M (Eds.). Sankt Augustin, West Germany: GMD, pp. 22.1-22.15.
- Jerrams-Smith, J (1987). "An expert system within a supportive interface for UNIX" *Behaviour and Information Technology* 6 (1) 37-41.
- Johnson, L and Johnson, N E (1987a). "Research and methods in building knowledge systems at Brunel University" In: *Knowledge acquisition for engineering applications*, Pavelin, C J and Wilson, M D (Eds.). pp. 84-95.
- Johnson, L and Johnson, N E (1987b). "Knowledge elicitation involving teachback interviewing" In: *Knowledge acquisition for expert systems: a practical handbook*, Kidd, A L (Ed.). New York, NY: Plenum, pp. 91-108.
- Johnson, N E (1985). "Varieties of representation in eliciting and representing knowledge for IKBS" *International Journal of Systems Research and Information Science* 1 69-90.
- Johnson, N E (1987). "Mediating representations in knowledge elicitation" *Proceedings of the first European Workshop on Knowledge Acquisition for Knowledge-Based Systems*. Reading, UK, 2-3 September 1987, Reading: Reading University, section A2.
- Johnson, P E, Zualkernan, I and Garber, S (1987). "Specification of expertise" *International Journal of Man-Machine Studies* 26 (2) 161-181.
- Josefson, I (1987). "Knowledge and experience" *Applied Artificial Intelligence* 1 (2) 173-180.
- Kahn, G, Nowlan, S and McDermott, J (1985a). "MORE: An intelligent knowledge acquisition tool" *Proceedings of the Ninth International Joint Conference on Artificial Intelligence (IJCAI-85)*, Los Angeles, CA, USA, 18-23 August 1985, pp. 581-584.
- Kahn, G, Nowlan, S and McDermott, J (1985b). "Strategies for knowledge acquisition" *IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-7* (5) 511-522.
- Karbach, W and Tong, X (1988). "Filling in the knowledge acquisition gap: via KADs' models of expertise to ZDEST-2's expert system" In: *Proceedings of the European Knowledge Acquisition Workshop (EKAW '88)*, Boose, J H, Gaines, B R and Linster, M (Eds.). Sankt Augustin, West Germany: GMD, pp. 31.1-31.17.
- Kassirer, J P, Kuipers, B J and Gorry, G A (1982). "Towards a theory of clinical expertise" *The American Journal of Medicine* 73 251-259.
- Kellett, J and Esfahani, L (1988). "Vegan and KET, an integrated graphical approach to knowledge representation and acquisition" *Conference on Human and Organisational Issues of Expert Systems*, Stratford-on-Avon, UK, 4-6 May 1988 (sponsored by ICL and The Ergonomics Society). See also a similar paper to appear in *Knowledge Based Systems*, September 1988.
- Kelly, G A (1955). *The psychology of personal constructs* (two volumes), New York, NY: W. W. Norton.
- Keravnou, E T and Johnson, L (1986). *Competent expert systems. A case study in fault diagnosis*, London: Kogan Page.
- Kessel, K L (1986). "Methodological tools for knowledge acquisition" *Proceedings of the 1986 IEEE International Conference on Systems, Man and Cybernetics* (Cat. No. 86CH2364-8), Atlanta, GA, USA, 14-17 October 1986, New York, NY: IEEE 1986, 1 pp. 541-546.
- Kidd, A L (1987). "Knowledge acquisition—an introductory framework" In: *Knowledge acquisition for expert systems: a practical handbook*, Kidd, A L (Ed.). New York, NY: Plenum, pp. 1-16.
- Kidd, A L (Ed.) (1987). *Knowledge acquisition for expert systems: a practical handbook*, New York, NY: Plenum Press.
- Killin, J L and Hickman, F R (1986a). "Model formulation as analogy in complex domains". Unpublished MS.
- Killin, J L and Hickman, F R (1986b). "The role of phenomenological techniques of knowledge elicitation in

- complex domains". Unpublished paper presented at *Expert Systems '86 Conference*, Brighton, UK, 15–18 December 1986.
- Kitto, C M and Boose, J H (1987a). "Heuristics for expertise transfer: an implementation of a dialog manager for knowledge acquisition" *International Journal of Man-Machine Studies* **26** (2) 183–202.
- Kitto, C M and Boose J H (1987b). "Choosing knowledge acquisition strategies for application tasks" *WESTEX-87 Proceedings of the Western Conference on Expert Systems* (Cat. No. 87CH2463–8), Anaheim, CA, USA, 2–4 June 1987, Washington, DC, USA: IEEE, pp. 96–103.
- Klinker, G et al. (1987). "KNACK: report driven knowledge acquisition" *International Journal of Man-Machine Studies* **26** (1) 65–79.
- Kolodner, J L (1985). "Towards an understanding of the role of experience in the evolution from novice to expert" In: *Developments in expert systems*, Coombs, M J (Ed.). London: Academic Press, pp. 95–116.
- Kornell, J (1987). "Formal thought and narrative thought in knowledge acquisition" *International Journal of Man-Machine Studies* **26** 203–212.
- Kuipers, B J and Kassirer, J P (1984). "Causal reasoning in medicine: analysis of a protocol" *Cognitive Science* **8** 363–385.
- Kuipers, B J, Moskowitz, A J and Kassirer, J P (1988). "Critical decisions under uncertainty: representation and structure" *Cognitive Science* **12** (2) 177–210.
- LaFrance, M (1987). "The Knowledge Acquisition Grid: a method for training knowledge engineers" *International Journal of Man-Machine Studies* **26** 245–255.
- Lewis, C. (1982). "Using the 'thinking aloud' method in cognitive interface design". Research report RC9265 (No 40713). Yorktown Heights, NY: Thomas J Watson Research Center, IBM.
- Lehner, P E and Adelman, L (1987). "Biases in knowledge engineering" *Proceedings of the 1987 IEEE International Conference on Systems, Man and Cybernetics* (Cat. No. 87CH2503–1), Alexandria, VA, USA, 20–23 October 1987. Piscataway, NJ: IEEE, 1 pp. 376–377.
- Linster, M (1988). "KRITON: a knowledge elicitation tool for expert systems" In: *Proceedings of the European Knowledge Acquisition Workshop (EKAW'88)*, Boose, J H, Gaines, B R and Linster, M (Eds.). Sankt Augustin, West Germany: GMD, pp. 4.1–4.9.
- Littman, D C (1987). "Modelling human expertise in knowledge engineering: some preliminary observations" *International Journal of Man-Machine Studies* **26** (1) 81–92.
- McDermott, J (1986). "Making expert systems explicit" *Proceedings of the Tenth Congress of the International Federation of Information Processing Societies*, Dublin, Ireland, 1986. Amsterdam: Elsevier, pp. 539–544.
- McDermott, J (in press). "Preliminary steps towards a taxonomy of problem-solving methods" In: *Automating knowledge acquisition for expert systems*, Marcus, S (Ed.).
- McDonald, J E and Schvaneveldt, R W (1987). "The application of user knowledge to interface design". Memorandum in Computer and Cognitive Science MCCC-87-93. Las Cruces, NM: Computing Research Laboratory, New Mexico State University.
- Marcus, S (1987). "Taking backtracking with a grain of SALT" *International Journal of Man-Machine Studies* **26** (3) 383–398.
- Marcus, S (Ed.) (in press). *Automating knowledge acquisition for expert systems*.
- Marcus, S, Stout, J and McDermott, J (1988). "VT: An expert elevator designer that uses knowledge-based reasoning" *AI Magazine* **9** (1) 95–112.
- Merry, M (Ed.). (1985). "Expert Systems 85" *Proceedings of the Fifth Technical Conference of the British Computer Society Specialist Group on Expert Systems*, University of Warwick, UK, 17–19 December 1985, Cambridge: Cambridge University Press.
- Mettrey, W (1987). "An assessment of tools for building large knowledge-based systems" *AI Magazine* **8** (4) 81–89.
- Michalski, R S and Chilausky, R L (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.
- Michie, D (1986a). "Super Expert hits the UK". Report in *Expert Systems User* **2** (9) 32.
- Michie, D (1986b). "The superarticulacy phenomenon in the context of software manufacture" In: *Expert systems: automating knowledge acquisition*, Michie, D and Bratko, I. Wokingham: Addison-Wesley, pp. 21–44.
- Michie, D et al. (1984). "RULEMASTER: a second generation knowledge engineering facility" *First Conference on Artificial Intelligence Applications* (Cat. No. 84CH2107–1), Denver, CO, USA, 5–7 December 1984. Silver Spring, MD: IEEE, pp. 591–597. See also following reference.
- Michie, D, Muggleton, S, Reise, C and Zubrick, S (1984). "RULEMASTER: a second generation knowledge engineering facility". Radian Technical Report MI-R-623, December 1984. Austin, TX: Radian Corporation.
- Milne, R (1987). "Strategies for diagnosis" *IEEE Transactions on Systems, Man and Cybernetics SMC-17* (3) 333–339.

- Mingers, J (1987). "Expert systems—rule induction with statistical data" *Journal of the Operational Research Society* **38** (1) 39–47.
- Mitchell, A A (1987). "The uses of alternative knowledge acquisition procedures in the development of a knowledge-based media planning system" *International Journal of Man-Machine Studies* **26** 399–411.
- Mittal, S and Dym, C L (1985). "Knowledge acquisition from multiple experts" *AI Magazine* **5** (2) 32–36.
- Moller, J-U (1988). "Knowledge acquisition from texts" In: *Proceedings of the European Knowledge Acquisition Workshop (EKAW'88)*, Boose, J H, Gaines, B R and Linster, M (Eds.). Sankt Augustin, West Germany: GMD, pp. 25.1–25.16.
- Moralee, D S (Ed.) (1988). "Research and development in expert systems IV" *Proceedings of Expert Systems '87, the Seventh Annual Technical Conference of the British Computer Society Specialist Group on Expert Systems*, Brighton, UK, 14–17 December 1987. Cambridge: Cambridge University Press.
- Morik, K (1987a). "Acquiring domain models" *International Journal of Man-Machine Studies* **26** 93–104.
- Morik, K (1987b). "Knowledge acquisition and machine learning—the issue of modelling" *Proceedings of the first European Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Reading, UK, 2–3 September 1987. Reading: Reading University, section A4.
- Motta, E, Eisenstadt, M, West, M, Pitman, K and Evertsz, R (1986). "KEATS: The knowledge engineer's assistant (Alvey Project IKBS/20)". Final Project Report, HCRL Technical Report No. 20, December 1986. Milton Keynes, UK: Human Cognition Research Laboratory, The Open University.
- Motta, E, Eisenstadt, M, Pitman, K and West, M (1988). "Support for knowledge acquisition in the knowledge engineer's assistant". HCRL Technical Report No. 30, January 1988. Milton Keynes, UK: Human Cognition Research Laboratory, The Open University; also published in *Expert Systems* **5** (1) 6–28.
- Musen, M A, Fagan, L M, Combs, D M and Shortliffe, E H (1987). "Use of a domain model to drive an interactive knowledge-editing tool" *International Journal of Man-Machine Studies* **26** 105–121.
- Musen, M A (1988). "Conceptual models of interactive knowledge-acquisition tools" In: *Proceedings of the European Knowledge Acquisition Workshop (EKAW '88)*, Boose, J H, Gaines, B R and Linster, M (Eds.). Sankt Augustin, West Germany: GMD, pp. 26.1–26.15.
- Newell, A (1982). "The knowledge level" *Artificial Intelligence* **18** 87–127.
- Newell, A and Simon, H A (1972). *Human problem solving*, Englewood Cliffs, NJ: Prentice-Hall.
- Niblett, T and Bratko, I (1987). "Learning decision rules in noisy domains" In: *Research and development in expert systems III*, Bramer, M A (Ed.). pp. 25–34.
- Nisbett, R E and Wilson, T D (1977). "Telling more than we can know: verbal reports on mental processes" *Psychological Review* **84** (3) 231–259.
- Olson, J R and Rueter, H H (1987). "Extracting expertise from experts: methods for knowledge acquisition" *Expert Systems* **4** (3) 152–168.
- Partridge, D (1985). "Input-expectation discrepancy reduction: a ubiquitous mechanism" *Proceedings of the Ninth International Joint Conference on Artificial Intelligence (IJCAI-85)*, Los Angeles, California, USA, 18–23 August 1985, pp. 267–273.
- Pask, G (1975). *Conversation, cognition and learning: a cybernetic theory and methodology*, Amsterdam: Elsevier.
- Patel, V L and Groen, G J (1986). "Knowledge based solution strategies in medical reasoning" *Cognitive Science* **10** 91–116.
- Pauker, S G, Gorry, G A, Kassirer, J P and Schwartz, M D (1976). "Towards the simulation of clinical cognition" *The American Journal of Medicine* **60** 981–998.
- Pavelin, C J and Wilson, M D (Eds.) (1987). "Knowledge acquisition for engineering applications". Rutherford Appleton Laboratory Report No. RAL-87-055. July 1987. Didcot, UK: Rutherford Appleton Laboratory.
- Perkin, R M E, Pitt, M and Price, A E (1987). "Expert systems: practical applications in a traditional industry" In: *Research and development in expert systems III*, Bramer, M A (Ed.). pp. 231–241.
- Prerau, D S (1987). "Knowledge acquisition in the development of a large expert system" *AI Magazine* **8** (2) 43–51.
- Quinlan, J R (1982). "Semi-autonomous acquisition of pattern-based knowledge" In: *Introductory readings in expert systems*, Michie, D (Ed.). New York, NY: Gordon and Breach, pp. 192–207.
- Quinlan, J R (1986). "Induction of decision trees" *Machine Learning* **1** (1) 81–106.
- Regoczei, S and Plantinga, E P O (1987). "Creating the domain of discourse: ontology and inventory" *International Journal of Man-Machine Studies* **27** 235–250.
- Reichgelt, H and van Harmelen, F (1986). "Criteria for choosing representation languages and control regimes for expert systems" *Knowledge Engineering Review* **1** (4) 2–17.
- Roycroft, A E and Loucopoulos, P (1985). "ACCI—an expert system for the apportionment of close companies' income" In: *Research and development in expert systems*, Bramer, M A (Ed.). pp. 127–139.
- Saffiotti, A (1987). "An AI view of the treatment of uncertainty" *Knowledge Engineering Review* **2** (2) 75–97.
- Sandberg, J, Winkels, R and Breuker, J A (1988). "Knowledge acquisition for intelligent tutoring systems" In:

- Proceedings of the European Knowledge Acquisition Workshop (EKAW '88)*, Boose, J H, Gaines, B R and Linster, M (Eds.). Sankt Augustin, West Germany: GMD, pp. 27.1–27.12.
- Schreiber, G, Bredeweg, B, Davoodi, M and Wielinga, B J (1987). "Towards a design methodology for KBS". Deliverable D8, Esprit Project 1098, November 1987, Amsterdam: University of Amsterdam.
- Schreiber, G, Breuker, J A, Bredeweg, B and Wielinga, B J (1988). "Modelling in KBS development" In: *Proceedings of the European Knowledge Acquisition Workshop (EKAW '88)*, Boose, J H, Gaines, B R and Linster, M (Eds.). Sankt Augustin, West Germany: GMD, pp. 7.1–7.15.
- Schvaneveldt, R W, Durso, F T and Dearholt, D W (1985). "PATHFINDER: scaling with network structures". Memorandum in Computer and Cognitive Science MCCS-85-9, Computing Research Laboratory, New Mexico State University, Las Cruces, NM, USA.
- Schweickert, R et al. (1987). "Comparing knowledge elicitation techniques: a case study" *Artificial Intelligence Review* 1 (4) 245–253.
- Shadbolt, N R, Burton, A M, Hedgecock, A P and Rugg, G (1987). "Knowledge elicitation: an evaluation of techniques across experts". Technical report 87-3-1, Artificial Intelligence Group, Department of Psychology, University of Nottingham, UK.
- Shaw, M L G (Ed.) (1981). *Recent advances in personal construct psychology*, London: Academic Press.
- Shaw, M L G and Gaines, B R (1986). "Interactive elicitation of knowledge from experts" *Future Computing Systems* 1 (2) 150–190. (Alternatively, see the authors' chapter in Kidd (1987) for a near-identical account).
- Shaw, M L G and Gaines, B R (1987). "KITTEN: Knowledge initiation and transfer tools for experts and novices" *International Journal of Man-Machine Studies* 27 251–280.
- Silva, A H and Regan, D C (1986). "Using cognitive psychology techniques for knowledge acquisition" *Proceedings of the 1986 IEEE International Conference on Systems, Man and Cybernetics* (Cat. No. 86CH2364-8), Atlanta, GA, USA, 14–17 October 1986. New York, NY: IEEE 1986, 1 pp. 530–535.
- Slatter, P E (1987). *Building expert systems: cognitive emulation*, Chichester: Ellis Horwood.
- Smart, G (1987). "How to make expert systems . . . and stay in business" *Third International Expert Systems Conference*, London, UK, 2–4 June 1987. Oxford: Learned Information, pp. 215–222.
- Soloway, E, Bachant, J and Jensen, K (1987). "Assessing the maintainability of XCON-in-RIME: coping with the problems of a very large rule base" *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, Seattle, Washington, USA, August 1987. Los Altos, CA: Morgan Kaufmann 2 pp. 824–829.
- Sorenson, T and Nordhus, O (1988). "LOKE: a drill bit selection system" In: *Research and development in expert systems IV*, Moralee, D S (Ed.). pp. 32–41.
- Sowa, J F (1984). *Conceptual structures: information processing in mind and machine*, Reading, MA: Addison-Wesley.
- Sparck Jones, K (1985). "Natural language interfaces for expert systems: an introductory note" In: *Research and development in expert systems*, Bramer, M A (Ed.). Cambridge: Cambridge University Press, pp. 85–94.
- Spradley, J P (1979). *The ethnographic interview*, New York, NY: Holt, Rinehart and Winston.
- Stamper, R K (1988). "Pathologies of AI: responsible use of artificial intelligence in professional work" *AI & Society* 2 (1) 3–16.
- Sweller, J., Mawer, R F and Howe, W (1982). "Consequences of history-cued and means-end strategies in problem solving" *American Journal of Psychology* 95 455–483.
- Systems Designers Scientific (1986). *Final report for the Alvey RESCU Awareness Club*, Camberley: Systems Designers plc.
- Thomas, M (1988). "Structured techniques for the development of expert systems" In *Research and development in expert systems IV*, Moralee, D S (Ed.). Cambridge: Cambridge University Press, pp. 258–266.
- Trimble, E G and Cooper, C N (1987). "Experience of knowledge acquisition for expert systems in construction" In: *Knowledge acquisition for engineering applications*, Pavelin, C J and Wilson, M D (Eds.). pp. 7–13.
- Tversky, A and Kahneman, D. (1974). "Judgement under uncertainty: heuristics and biases" *Science* 185 1124–1131.
- Twine, S (1988). "From information analysis towards knowledge analysis" In: *Proceedings of the European Knowledge Acquisition Workshop (EKAW '88)*, Boose, J H, Gaines, B R and Linster, M (Eds.). Sankt Augustin, West Germany: GMD, pp. 6.1–6.15.
- van Terheyden, A G R and Chalcraft, D A (1987). "Combining inductive and deductive reasoning" *Computer-Aided Engineering Journal* 4 (1) 24–28.
- Waldron, V R (1985). "Process tracing as a method for initial knowledge acquisition" *Proceedings of the Second IEEE Conference on Artificial Intelligence Applications: The Engineering of Knowledge-Based Systems* (Cat. No. 85CH2215-2), Miami Beach, FL, USA, 11–13 December 1985, Washington DC: IEEE Comput. Soc. Press 1985, pp. 661–665.
- Waterman, D A (1986). *A guide to expert systems*, Reading, MA: Addison-Wesley.
- Welbank, M (1983). *A review of knowledge acquisition techniques for expert systems*, Ipswich: Martlesham Consultancy Services.

- Welbank, M (1987). "Perspectives on knowledge acquisition" In: *Knowledge acquisition for engineering applications*, Pavelin, C J and Wilson, M D (Eds.). pp. 14–19.
- White, A P (1987). "Probabilistic induction by dynamic path generation in virtual trees" In: *Research and development in expert systems III*, Bramer, M A (Ed.). pp. 35–46.
- Wielinga, B J and Breuker, J A (1985). "Interpretation of verbal data for knowledge acquisition" In: *Advances in artificial intelligence*, O'Shea, T (Ed.). Amsterdam: North-Holland, pp. 3–12.
- Wielinga, B J and Breuker, J A (1986). "Models of expertise" *Proceedings of Seventh European Conference on Artificial Intelligence (ECAI '86)*, Brighton, July 1986 1 306–318.
- Wilkins, D C, Buchanan, B G and Clancey, W J (1984). "Inferring an expert's reasoning by watching". Heuristic Planning Project Report No. HPP-84-29. Department of Computer Science, Stanford University, Stanford, CA, USA.
- Wilson, M D (1987). "Task analysis for knowledge acquisition" In: *Knowledge acquisition for engineering applications*, Pavelin, C J and Wilson, M D (Eds.). pp. 68–83.
- Winograd, T and Flores, C F (1986). *Understanding computers and cognition: a new foundation for design*, Norwood, NJ: Ablex.
- Winston, P H (1982). "Learning new principles from precedents and exercises" *Artificial Intelligence* 19 321–350.
- Winston, P H (1986). "'Models rule', Winston warns". Report in *Expert systems User* 2 (9) 5.
- Wood, S (1987). "Expert systems for theoretically ill-formulated domains" In: *Research and development in expert systems III*, Bramer, M A (Ed.). pp. 132–139.
- Woods, D D (1986). "Cognitive technologies: the design of joint human-machine cognitive systems" *AI Magazine* 6 (4) 86–92.
- Wright, G and Ayton, P (1987). "Eliciting and modelling expert knowledge" *Decision Support Systems* 3 (1) 13–26.
- Wyatt, J C and Emerson, P A (1988). "'Pragmatic' knowledge engineering: a new approach for difficult domains" *Conference on Human and Organisational Issues of Expert Systems*, Stratford-on-Avon, UK, 4–6 May 1988 (sponsored by ICL and The Ergonomics Society).
- Young, R M (1988). "Role of intermediate representations in knowledge elicitation" In: *Research and development in expert systems IV*, Moralee, D S (Ed.). Cambridge: Cambridge University Press, pp. 287–288.
- Young, R M and Gammack, J G (1987). "Role of psychological techniques and intermediate representations in knowledge elicitation" *Proceedings of the first European Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Reading, UK, 2–3 September 1987. Reading: Reading University, section D7.
- Zimolong, B, Nof, S Y, Eberts, R E and Salvendy, G (1987). "On the limits of expert systems and engineering models in process control" *Behaviour and Information Technology* 6 (1) 15–36.

Postscript

The papers cited by Candy and Lunn (1988) (in revised format), Eklund and Kuczora (1988), Hayball and Barlow (1988) and Wyatt and Emerson (1988) are to be included with some other papers presented at the ICL/Ergonomics Society Conference on Human and Organisational Issues in Expert Systems, in a forthcoming book to be edited by Dianne Berry and Anna Hart. No more citation details were available as this paper went to press.