KNOWLEDGE ACQUISITION : THE STATE OF THE ART

Mildred L. G. SHAW *

Introduction

The problem of expertise elicitation from a skilled person is well-known in the literature of psychology (Nisbett & Wilson, 1977, Broadbent, Fitzgerald and Broadbent 1986). Hawkins (1983) has analyzed the nature of expertise and emphasizes its severe limitations and dependence on critical assumptions which are often implicit. Dixon (1981) has surveyed studies showing that much human activity is not accessible to awareness. Collins (1985) has studied knowledge transfer processes among scientists and suggests that some knowledge may not be accessible through the expert, not only because he cannot express it, but also because he may not be aware of its significance to his activity. Bainbridge (1979, 1986) has reviewed the difficulties of verbal debriefing and notes that there is no necessary correlation between verbal reports and mental behavior and that many psychologists feel strongly that verbal data are useless.

The main problems identified in accessing an expert's knowledge are :

– «Expertise» may be fortuitous. Results obtained may be dependent on features of the situation which the experti is not controlling.

– Expertise may not be available to awareness. An expert may not be able to transmit the expertise by critiquing the performance of others because he is not able to evaluate it.

– Expertise may not be expressible in language. An expert may not be able to transmit the expertise explicitly because he unable to express it.

– Expertise may not be understandable when expressed in language. An apprentice may not be able to understand the language in which the expertise is expressed.

– Expertise may not be applicable even when expressed in language. An apprentice may not be able to convert verbal comprehension of the basis of a skill into skilled performance.

– «Expertise» expressed may be irrelevant. Much of what is learnt, particularly under random reinforcement schedules, is superstitious behavior that neither contributes nor detracts from performance.

– Expertise expressed may incomplete. There will usually be implicit situational dependencies that make explicit expertise inadequate for performance.

– «Expertise» expressed may incorrect. Experts may make explicit statements which do not correspond to their actual behavior and lead to incorrect performance (Gaines 1987a).

Hawkins (1983) has abstracted from industrial experience in developing mi-

* Knowledge Science Institute, University of Calgary, Alberta, Canada T2N 1N4.

neral exploration expert systems and proposed a model of human expertise relevant to expert systems. Figure 3 presents a summary of the essential features of this model :

- The expert first elicits data about the problem from the client ;
- He or she develops a minimal model that accounts for the data provided :
- He or she generates advice based on the model and feeds this back to the client :
- The client may accept the advice, or query it and, possibly, the model :
- The queries lead to further data elicitation, and repeat of the elicitation/modeling/advice/query cycle.

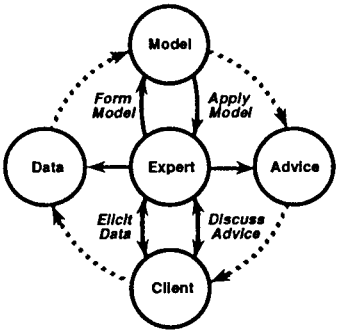

Fig.1 The negotiation cycle in expert-client interaction

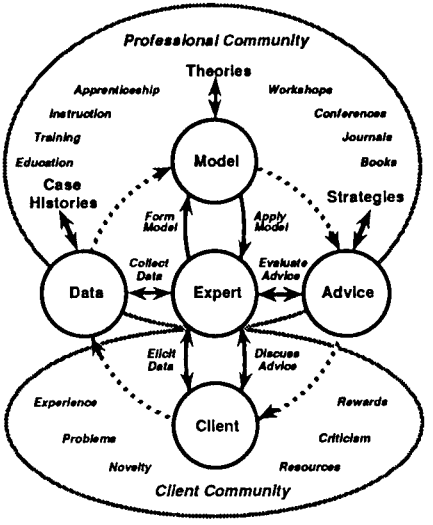

Fig.2 Extended knowledge acquisition process in expert-community interaction

Thus in Hawkins' model, the client plays an active role in further developing the model by providing more data until he or she is satisfied with the model and consequent advice. Expert advice giving and taking is part of a cycle of negociation around a process of model formation. The particular importance of this model of expert performance is the emphasis, it places on the embedded knowledge acquisition process. The expert is managing a data collection, modeling and decision cycle within a consultation. As Figure 1 shows, this implies that the expert is able to enhance his or her collection of case histories, models and advisory strategies through the consultation. Skilled performance is not the static utilization of expert knowledge and skills in Hawkins' model – it is instead a microcosm, and an essential part of, the wider processes of knowledge acquisition which are intrinsic to the role of an expert.

Access to the learning experiences of attempting to solve the problems of a client community is one major source of knowledge in the development of the expertise of an individual expert. The client community provides access to practical experience, a variety of problems, particularly novel ones that go beyond existing expertise, and it also manages the growth of expertise through systems of reward, criticism and access to resources. However, learning from experience is a slow and error-prone process, and socially significant areas of expertise become associated with professional communities that attempt to expedite learning and reduce errors through the sharing of experience,. Figure 2 extends the model of Figure 1 to include these other processes, such as apprenticeship, instruction, training, education, workshops, conferences, journals and books. Professional communities usually also play a major role in refining and directing the client community's reward, criticism and resource allocation systems.

The role of the professional community can be seen as an indirect enhancement of the basic processes of knowledge acquisition in the negotiation cycle of expert performance. The individual can acquire the data of others through access to reported case histories. He or she can enhance his or her repertoire of models through the published theories developed by others to account for their data and the data of many others. He or she can enhance his or her repertoire of advice through the published advisory strategies developed by others and evaluated by them, or yet others, against past case histories, theoretical foundations, and experience in use.

Knowledge engineering

Knowledge acquisition for practical systeme development has come to be termed *knowledge engineering*, following Feigenbaum's (1980) use of the term to describe the reduction of a large body of knowledge to a precise set of facts and rules. The term *knowledge engineer* has come to be used for the person responsible for such system development, and concise job descriptions for knowledge engineers have been given Hayes-Roth, Waterman & Lenat (1983). The basic model for knowledge engineering has been that the knowledge engineer mediates between the expert and knowledge base, eliciting knowledge from the expert, encoding it for the knowledge base, and refining it in collaboration with the expert to achieve acceptable performance. Figure 3 shows this

basic model with manual acquisition of knowledge from an expert followed by interactive application of the knowledge with multiple clients through an expert system shell :

- The knowledge engineer interviews the expert to elicit his or her knowledge ;
- The knowledge engineer encodes the elicited knowledge for the knowledge base ;
- The shell uses the knowledge base to make inferences about a particular case ;
- Clients use the shell to obtain advice about particular cases.

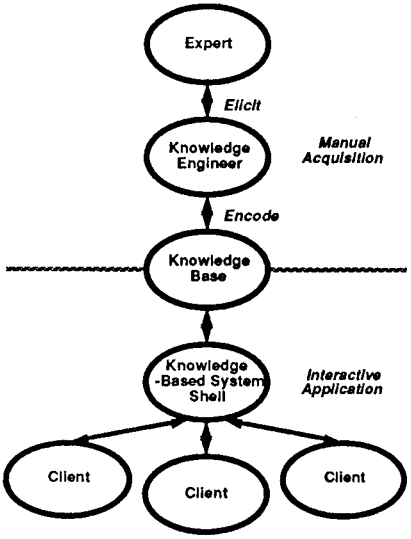

Fig.3 Basic model of knowledge acquisition by a knowledge engineer interviewing an expert

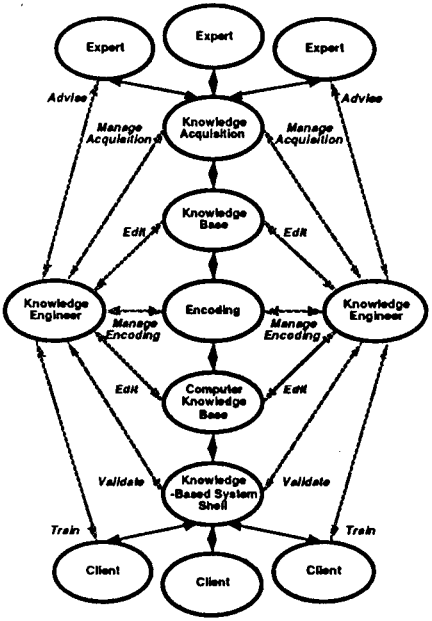

Fig.4 The knowledge engineers' roles in current system development using automated knowledge acquisition tools

Figure 3 may be compared with Figure 2 to show the introduction of the knowledge engineer and expert system as intermediaries between the expert and client, as socio-technical extensions to knowledge processes in society.

It seemed possible that knowledge engineering might develop as a profession on a par with systems analysis and programing, and that a major shortage of skilled knowledge engineers would cause problems initially that would be overcome eventually as the profession developed. However, this scenario now appears less and less lileky. There is certainly a shortage of knowledge engineers and major problems in developing applications, but doubts have been cast on the notion that human labor is the appropriate solution to the knowledge engineering problem.

Figure 4 shows how the role of the knowledge engineer has changed as such systems have come into use. Interactive knowledge acquisition and econding tools can greatly reduce the need for the knowledge engineer to act as an intermediary, but, in most applications, they leave a substantial role for the knowledge engineer. As shown in Figure 4, knowledge engineers have responsibility for :

- Advising the experts on the process of interactive knowledge elicitation;
- Managing the interactive knowledge acquisition tools, setting them up appropriately ;
- Editing the unencoded knowledge base in collaboration with the experts ;
- Managing the knowledge encoding tools, setting them up appropriately ;
- Editing the encoded knowledge base in collaboration with the experts ;
- Validating the application of the knowledge base in collaboration with the experts ;
- Setting up the user interface in collaboration with the experts and clients ;
- Training the clients in the effective use of the knowledge base in collaboration with the expert by developing operational and training procedures.

This use of interactive elicitation can be combined with manual elicitation and with the use of the interactive tools by knowledge engineers rather than, or in addition to, experts. Knowledge engineers can :

- Directly elicit knowledge from the expert ;
- Use the interactive elicitation tools to enter knowledge into the knowledge base.

Figure 4 shows multiple knowledge engineers since the tasks above may require the effort of more than one person, and some specialization may be appropriate. Multiple experts are also shown since it is rare for one person to have all the knowledge required, and, even if this were so, comparative elicitation from multiple experts is itself a valuable knowledge elicitation technique (Boose 1987, Shaw & Woodward 1987, Shaw & Gaines 1988).

Knowledge acquisition tools

A number of tools have been developed for the rapid prototyping of knowlege systems. The objective of systems such as PLANET (Shaw & Gaines 1983, 1986b, 1987a). ETS (Boose 1984, 1986), MOLE (Eshelman, Ehret, McDermott & Tan 1986), SALT (Marcus 1987), KITTEN (Shaw & Gaines 1986 a, 1987b), KNACK (Klinker Bentolila, Genetet, Grimes & McDermott 1986), KRITON

(Diederich, Ruhmann & May 1986), OPAL (Musen, Fagan, Combs & Shortliffe 1986), AQUINAS (Boose & Bradshaw 1987) and KSSO (Gaine 1987b) is to expedite the process of knowledge acquisition and transfer to knowledge-based systems.

Boose (1988) has surveyed a very wide range of knowledge acquisition tools and classified them along a variety of dimensions, analyzing the applicability of tools to tasks. He distinguishes three classes of tasks : analysis or classification tasks, synthesis or configuration tasks, and tasks combining analysis and synthesis. He notes that most knowledge acquisition tools currently are best suited to analysis tasks and are weak in acquiring expertise in synthesis.

The general requirements for a system to automate knowledge acquisition for knowledge-based systems are that the tools in a knowledge acquisition system are domain independent and directly applicable by experts without intermediaries. They should be able to access a diversity of knowledge sources including text, interviews with experts, and observations of expert behaviour and be able to encompass a diversity of perspectives including partial or contradictory input from different experts. Moreover, the knowledge acquisition system should be able to encompass a diversity of forms of knowledge and relationships between knowledge and be able to present knowledge from a diversity of sources with clarity as to its derivation, consequences and structural relations. Users of the knowledge acquisition system should be able to apply the knowledge in a variety of familiar domains and freely experiment with its implications. It should make provision for validation studies and as much of the operation of the knowledge acquisition system as possible should be founded on well-developed and explicit theories of knowledge acquisition, elicitation and representation. As the overall knowledge acquisition develops it should converge to an integrated system.

It is particularly important in describing tools that they are presented in the context of user interaction, since their primary function is to operate directly with experts. It is impossible to describe all the knowledge acquisition tools currently in the literature, and the references cited at the beginning of this section and specialist collections should be consulted for details and examples of use (Gaines & Boose 1988, Boose & Gaines 1988, Marcus 1988). The next section takes one particular tool that covers the range of methodologies discussed earlier, and uses it to illustrate the major features of a wide variety of tools.

An example of an integrated knowledge acquisition system

The problems of eliciting knowledge from experts as skilled performers of a task have been described. Clinical psychologists see the problem as one of cognitive defences that impede internal communication, and have developed techniques of verbal interaction to identify underlying cognitive processes (Freud 1914, Kelly 1955, Rogers 1967). These can be used to by-pass cognitive defences, including those resulting from automization of skilled behavior, and Welbank (1983) has surveyed psychological techniques for doing this.

One major class of tools for knowledge acquisition has been developed

around the 'personal construct psychology' of George Kelly (1955) and, in particular, his 'repertory grid' technique for eliciting conceptual systems as individual models of the world. In computational terms the grid may be seen as small entity-attribute database in which a domain is characterized by an expert in terms of critical objects within it and the properties of theses objects relevant to problem-solving within the domain. The elicitation of an expert's conceptual system necessarily involves approximation since a complete system may involve indefinitely complex relations. However, as a start one wishes to elicit the major distinctions that an individual uses in a domain, the terminology used for them, and the relation of such distinctions and terminology to those of others. Entity-attribute grid elicitation is an effective method for eliciting major distinctions and terminology in a domain (Shaw 1981).

The grid technique is an extensional approach in which individuals are asked to specify a set of entities in a domain, then make distinctions among them, naming the distinctions and classifying all the specified entities in terms of them. The extension of a distinction determined in this way is only an approximation to the underlying concept since critical entities may be missing in the classification. However, both manual and computer-based elicitation techniques attempt to prompt the individual for additional entities to discriminate between extensionally related distinctions (that is making the same, or similar, classifications). Knowledge acquisition is essentially a negotiation process leading to approximations to conceptual structures that are adequate for some practical purpose.

The application of interactive repertory grid elicitation to expert system development was first proposed by Gaines and Shaw (1980), taken up and demonstrated effective in industrial applications by Boose (1984), and has become in widespread use through a number of derivative system developments (Gammack & Young 1985, Diederich, Ruhmann & May 1986, Wahl 1986, Garg-Janardan & Salvendy 1987).

KSSO : Knowledge Support System Zero

Figure 5 shows the architecture of one system that integrates a wide range of knowledge acquisition methodologies. KSSO is written in Pascal and runs on the Apple Macintosh family of computers to provide a highly interactive and graphic knowledge acquisition environment. At the heart of KSSO is an object-oriented knowledge base in which knowledge is formally represented as a multiple-inheritance structure of classes, objects, properties, values, and relations. Such a structure generalizes the entity-attribute grids used in several early knowledge acquisition systems and has proved both general and powerful in a variety of applications.

The elicitation tools in KSSO are based on Shaw's (1980) computer-based interviewing techniques extended through the use of graphical rather than numerical data entry. Teh visualization and direct manipulation of knowledge structures through the graphic and click-and-drag facilities of modern workstations is an important development at the user interface giving experts, knowledge engineers and clients improved access to the knowledge structures.

– **Elicit** accepts specifications of entities within a domain and provides an interactive graphical elicitation environment within which the experts can distinguish entities to derive their attributes. The resultant class is continuously analyzed to provide feedback prompting the expert to enter further entities and attributes.

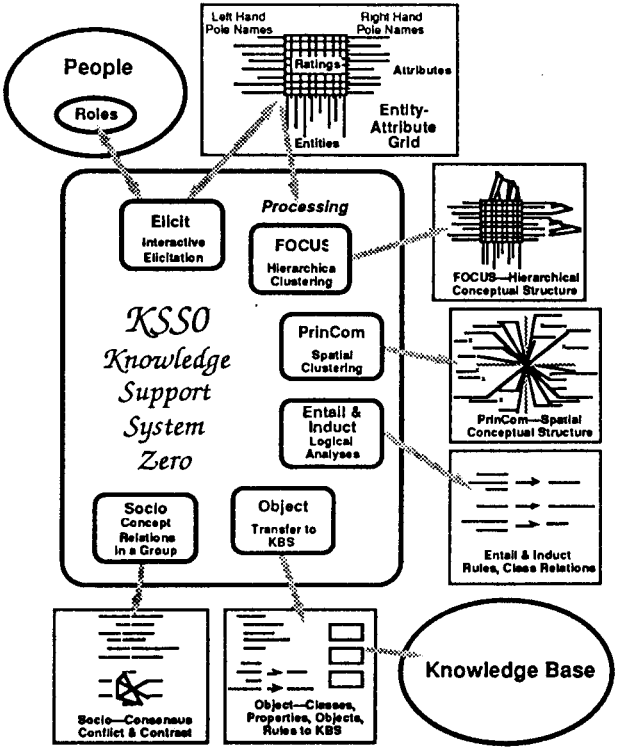

Fig.5 KSS0, an integrated knowledge acquisition and support system

– **Exchange** allows interactive exchange and comparison of entities and attributes from multiple experts in a domain.

The visual analysis tools consist of an interactive interface to represent the abstractions derived from those entities in terms of hierarchical clusters using Shaw's (1980) FOCUS algorithm, and relational diagrams such as a non-hierarchical conceptual maps derived through principal components analysis (Slater

1976). The objectives are to validate the raw domain knowledge and suggest further structure at a higher level through interactive topological induction (Rappaport & Gaines 1988) :

– **Focus** hierarchically clusters entities and attributes within a sub-domain prompting the experts to add higher-level entities structuring the domain.

– **PrinCom** spatially clusters entities and attributes within a sub-domain prompting the experts to add higher-level entities structuring the domain.

The group comparison tools consist of an interactive interface to represent the relations between the terminologies and conceptual systems of different experts, or experts and clients. The objectives are to determine the consensus, conflict, correspondence and contrast between different conceptual systems (Shaw & Gaines 1988) :

– **Socio** compares the structures for the same sub-domain generated by different experts, or the same expert at different times or from varying perspectives.

The inductive part consists in the derivation of constraints within the conceptual structures through logical entailment analyses (Gaines and Shaw 1985, Quinlan 1987, Cendrowska 1987). The objective is to suggest further structure at a higher level that translates into class inclusions or rules in NEXPERT :

– **Entail and Induct** induce logical entailments enabling the attributes of an entity, or the evaluations of a decision-making situation in a domain, to be derived from other attributes (Gaines and Shaw 1986).

The generative part consists in the transformation of the knowledge analysis made by the previous tools into a formalism understandable by the NEXPERT inference mechanisms (Rappaport 1987) :

– **Object** formats the specifications of sub-domains as classes, of entities as objects, of attributes as properties, and of entailments as methods, and transfers them to the performance tool.

Knowledge Acquisition in Action : Defining the Domain

The system is designed to be very clean and simple to use with a substantial entertainment value. The objective is to make interaction with the system attractive and fun while retaining a professional working environment. Help is available at all times through context sensitive, interactive 'advices buttons'. The expert should never be at a loss, need help he or she cannot get, or have to access a manual.

Figure 6 is an attribute screen from KSSO in a knowledge elicitation sequence concerned with spatial interpolation techniques for contour mapping (Shaw and Woodward 1987) and shows the techniques agreed on by all the geographers being placed on the attribute *local – global*. Some techniques have already been rated, some are waiting in a list on the left, and *proximal mapping* is being dragged on to the attribute bar. Both the attribute labels and the techniques can be edited any time, and also the placing of the techniques. Figure 7 shows two matched attributes, with both sets of ratings shown, so that they can be compared. Again, any one is moveable if the expert wants to adjust anything.

Techniques may also be matched, as shown in the screen in Figure 8. The technique at the top has a marker on each attribute sticking up, and the one at the bottom a marker sticking down. Again, any marker is moveable if the expert wants to change anything.

Knowledge Acquisition in Action : Visualization

The elicitation screens shown above all use the graphic capabilities of the workstation to encourage and support the experts' visualization of the data being entered. The rating of an entity on an attribute is seen in relation to the ratings of other entities to allow immediate consistency checks to be made. The analysis tools continue this emphasis on visualization and validation by presenting the conceptual structures underlying the data entered in graphic form.

At any time during the elicitation experts can ask for the structure of their entity-attribute conceptual system to be analyzed and displayed. The basic data structure of the grid that has been entered can be shown with entity names, attribute names, rating values, and shading to indicate low, medium and high values.

Figure 9 shows the hierarchical clustering of the final entities and attributes in based on the matches between them using Shaw's (1980) FOCUS algorithm. Note how this has tended to pull similarly shaded areas together. The clusters can be used to suggest the introduction of hierarchical relations in the conceptual framework. As can be seen on the attribute tree at the top right, there are three clusters at 75 %, and one attribute which is not part of any cluster. From the bottom right tree of techniques, it can be seen that the one most different from the others is *hand contouring*. Many other conclusions can be drawn from this diagram, but it is sufficient here just to give a flavour of what the expert experienced in using KSSO.

Figure 10 shows an alternative mapping into two-dimensions based on a principal components analysis (Slater 1976). This can again be used to identify groupings leading to more abstract attributes. Both forms of clustering are also important in allowing the expert to examine the semantic consistency of the data he or she has entered – if the way entities and attributes cluster does not make sense then it is a clear indication that something is wrong in the data entered or the experts' conceptualization of the domain.

KSSO also contains facilities for extracting conceptual clusters by analyzing word associations in textual material. This is useful for defining knowledge islands, and pump-priming the related elicitations with entities and attributes. The Socio tool in KSSO allows grids elicited in the same domain from multiple experts to be compared and consensual, conflicting, corresponding and contrasting relations between attributes to be displayed in graphical form (Shaw & Gaines 1988).

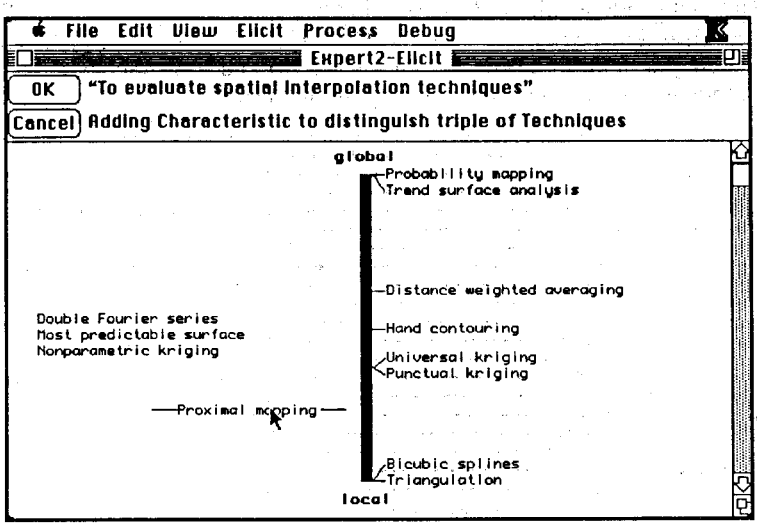

Figure 6. KSS0 attribute rating screen.

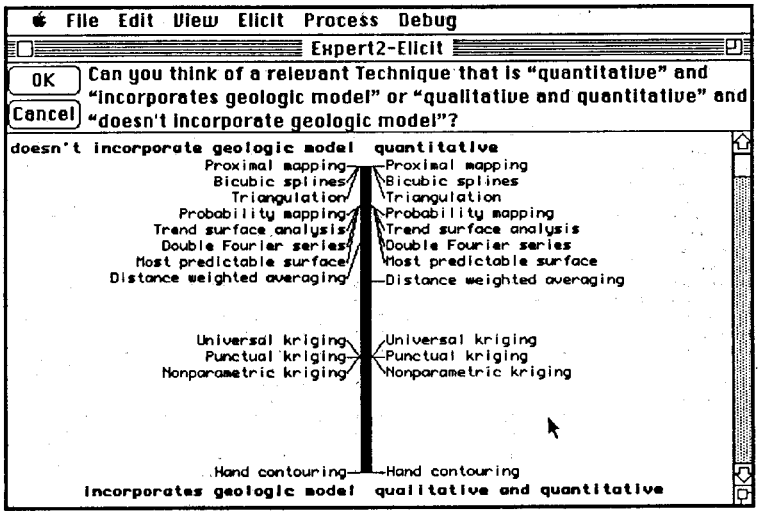

Figure 7. KSS0 showing matching attributes.

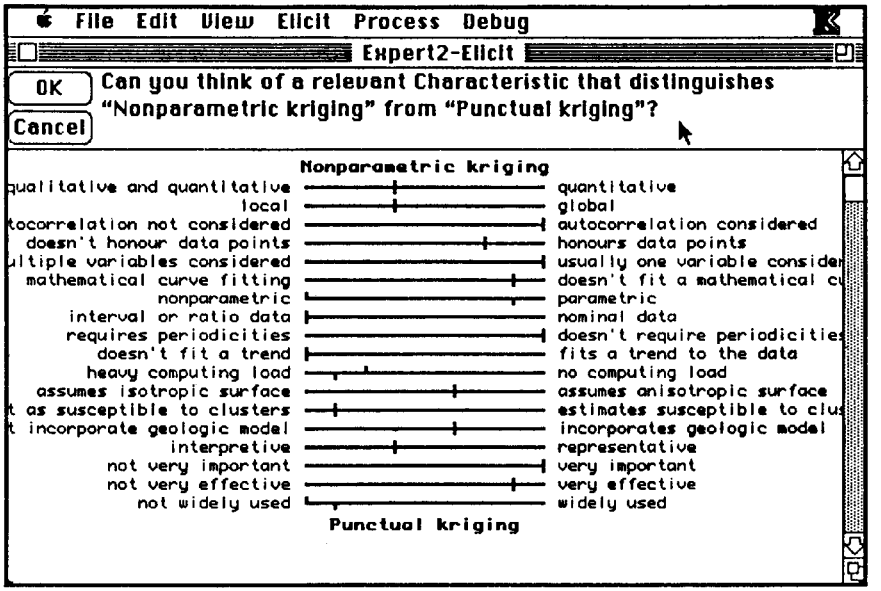

Figure 8. KSS0 showing matching techniques.

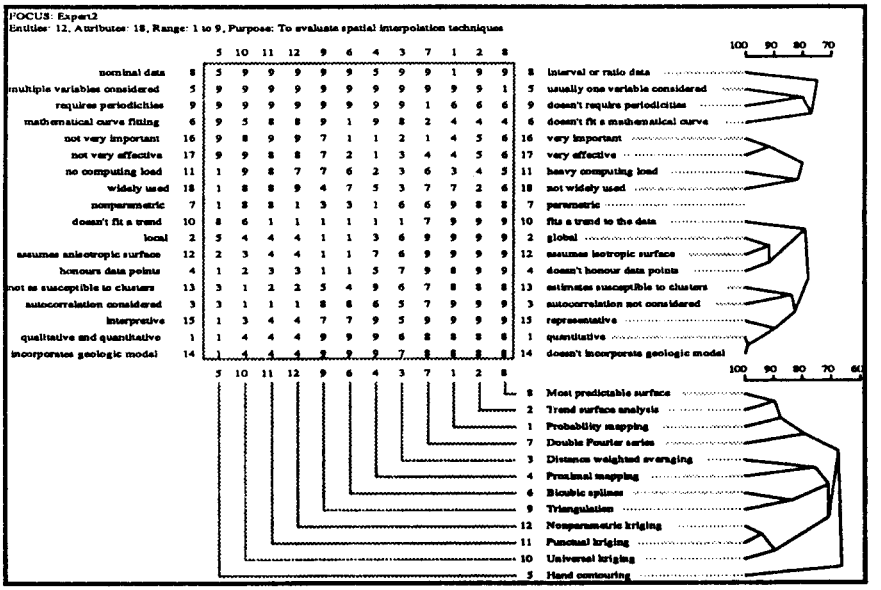

Figure 9. KSS0 hierarchical cluster analysis.

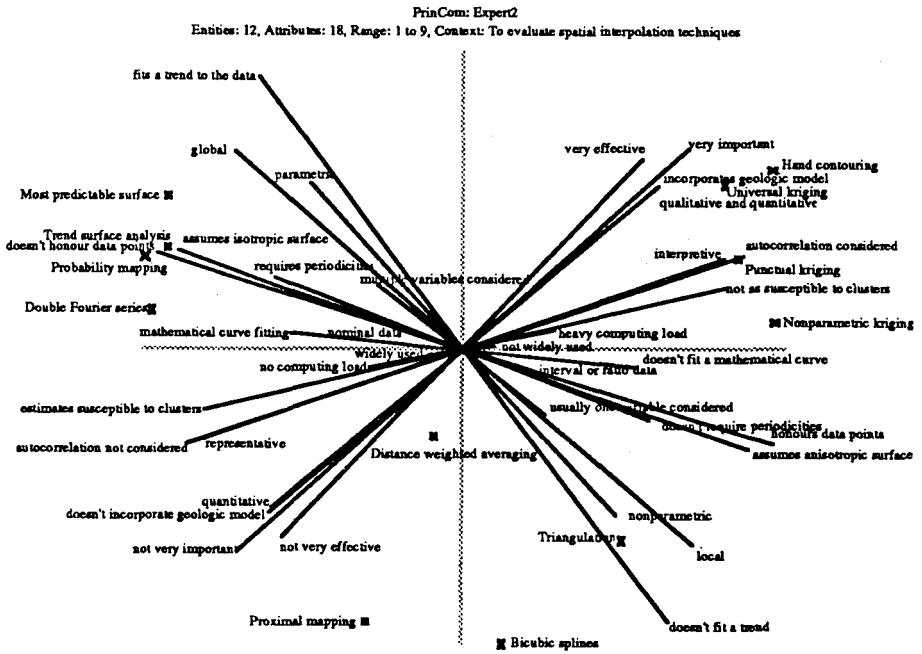

Fig.10 Spatial cluster analysis

Conclusions : Impediments to knowledge-based system development

Expert systems technology, like many other computing developments, has been 'oversold', and the expectations of what may be achieved through the use of knowledge-based system technology have been raised too high. There are still major problems in the development of knowledge-based systems, and it is appropriate to conclude this chapter by summarizing these problems – Figure 11 gives a synopsis of the major impediments to knowledge-based system development.

The top box in Figure 11 summarizes what may be termed the problems of *professionalism* in developing knowledge-based systems :

– *Theory* : the lack of scientific and engineering foundations for knowledge-based systems is a major impediment to the development of professional standards for expert system development. We understand remarkably little about the nature of expertise even though it one of the most valued commodities in our society. We also have little fundamental understanding of knowledge, the ways in which it is represented in human thought processes, the ways in which it is used in human problem solving, and the ways in which it is trasmitted between people.

– *Techniques* : the lack of well-established techniques for knowledge engineering is a major problem for new development teams. There is some folklore generally mutually contradictory, a few rules of thumb based on experience, but so far no in-depth studies of what is actually done and how successful it has been.

– *Tools* : the lack of commercially available tools for knowledge acquisition and transfer is also a major problem. There are now many research reports of such tools but few of those developed are available as products.

– *Training* : the training situation is improving all the time, largely because it has become the major secondary industry to the supply of expert system shells. A range of courses is being offered by commercial and university organizations. It is still a problem to get good quality courses targeted to specific requirements and on the schedules required.

The middle box in Figure 11 summarizes problems of *knowledge acquisition and transfer* in developing knowledge-based systems. Knowledge acquisition and transfer processes in human society are complex and poorly understood. Neurology, psychology, linguistics, education, sociology, anthropology, philosophy and systems theory all have significant contributions to make. However, integrating them, disambiguating jargon, and combining different objectives and perspectives are major problems. There is a variety of shor-term problems arising from these weaknesses in foundations :

– *System Specification* : the formal specification of software that has made software engineering a professional discipline which can ensure quality within well-defined budgets and schedules is notably lacking for expert system developments. We seem at the stage of conventional software development of twenty years ago. This should change as expert system developments move out of the laboratories into the hands of experienced software development groups trained in software management. There is little in the literature to support and expedite this change currently, and we need to develop specification disciplines for expert systems.

– *Access to Expertise* : the problems of getting experts to make the time commitments necessary to the development of expert systems should never be underestimated. Successful projects inevitably seem to have involved the experts in such a way that they become personally committed to the success of the system.

– *Acquisition of Skills* : communicating the basis of skilled performance between people, and between people and computers, is very difficult.

– *Representation of Knowledge* : encoding knowledge within the representation schemes of current shells can present major difficulties. Production rules as used in early expert systems offer very limited facilities for encoding simple knowledge structures at a single level. Frames and their generalization through objects allow conceptual hierarchies to be better expressed. Neither rules nor frames are adequate to encode the procedural or temporal knowledge which is often a key component of expertise.

The bottom box in Figure 11 summarizes problems of *system performance*

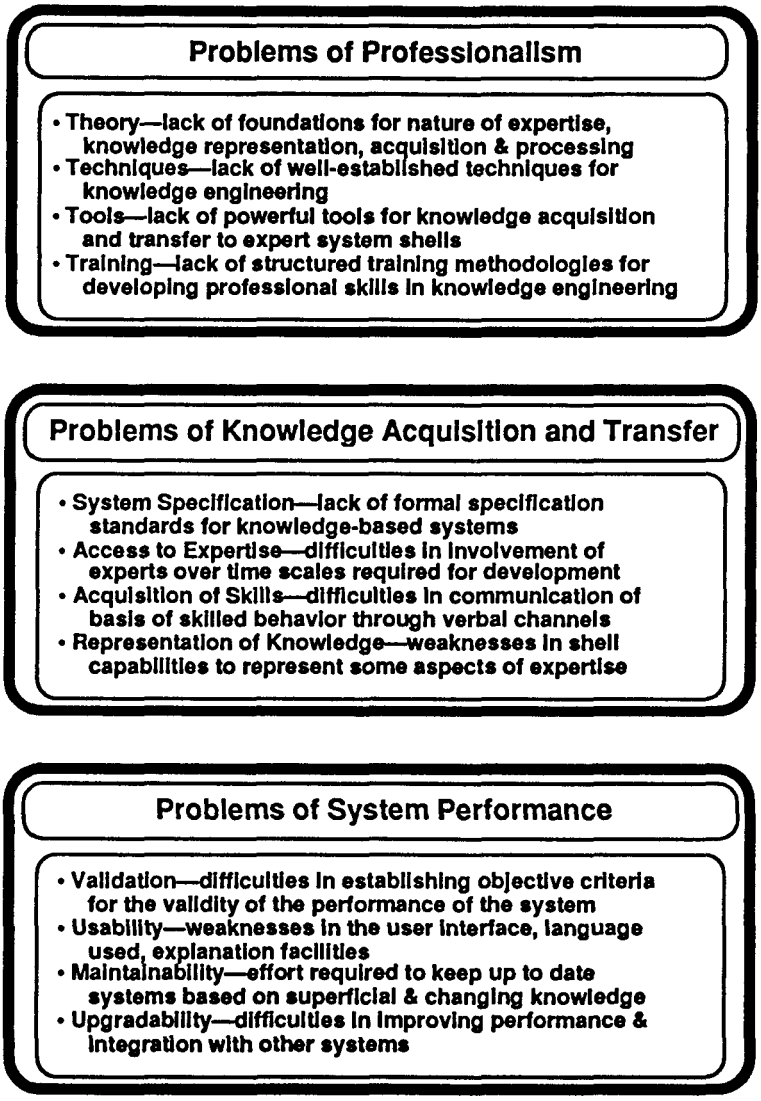

Fig.11 Impediments to knowledge-based system development

Problems of Knowledge Acquisition and Transfer

- System Specification—lack of formal specification standards for knowledge-based systems
- Access to Expertise—difficulties in involvement of experts over time scales required for development
- Acquisition of Skills—difficulties in communication of basis of skilled behavior through verbal channels
- Representation of Knowledge—weaknesses in shell capabilities to represent some aspects of expertise

Problems of System Performance

- Validation—difficulties in establishing objective criteria for the validity of the performance of the system
- Usability—weaknesses in the user interface, language used, explanation facilities
- Maintainability—effort required to keep up to date systems based on superficial & changing knowledge
- Upgradability—difficulties in improving performance & integration with other systems

in developing knowledge-based systems. Rapid prototyping is a significant practical objective and techniques for accelerated initial development without loss of long-term quality are now common in conventional data processing. However, expert system development emphasizes weaknesses in some areas of performance which we tend to take for granted in standard information technology. There are problems with :

Validation : because we can rarely define and quantify expertise, the validation of expert systems can be very difficult. Do we subjectively validate against the opinions of the expert or do we attempt objective validation of performance against performance standards, accepting that the experts may not themselves satisfy such criteria. How, do we determine whether the knowledge base captures the expert's relevant knowledge, accurately, consistently and completely ? There are very few validation studies in the literature.

Usability : the usability of expert systems raises many problems. If we accept the jargon of the expert and the rules of inference based on it, how do we ensure that the operators will use this jargon in the same way as the expert so that the inferences remain valid. It is such considerations that make nonsense of the popular stereotype that access to encoded expertise can make experts of all of us. Understanding the questions, advice and explanations from an expert system generally requires an operator with the same professional background as the expert.

Maintainability : the superficiality of the knowledge encoded in current expert systems raises major problems of maintenance. Keeping the knowledge base up to date may require an ongoing effort at the same level as the original development, a situation very different from most conventional data processing. The usual problems of extensive maintenance not introducing problems of its own are at least as severe for expert systems as for standard data processing.

Upgradability : many expert systems, and the shells supporting them have been designed as standalone packages with little scope for integration into large-scale data processing environments. These systems will be difficult to upgrade and integrate with other software.

This emphasis on problems impeding knowledge-based system development is not intended to be negative. In any major technological advance there are significant problems and it is the recognition and classification of these that indicates the growing maturity of the field. There is intensive, world-wide research in all the problem areas noted above, and growing experience in how to overcome them. By recognizing the problems project managers may be able to target their objectives to avoid some, and allocate their resources to overcome others.

Acknowledgements

I am grateful to Brian Gaines for access to his material on this topic.

REFERENCES

ADLER J. & MATHESON J., Expert knowledge elicitation using personal construct theory : obtaining plant layout expertise with 'PLANET', *BASc Thesis*, Department of Industrial Engineering, University of Toronto April 1986.

BAINBRIDGE L., Verbal reports as evidence of the process operator's knowledge, *International Journal of Man-Machine Studies, 11* (4), 411-436, July 1979.

BAINBRIDGE L., Asking questions and accessing knowledge, *Future Computing Systems 1* (2), 143-149, 1986.

BOOSE J.H., Personal construct theory and the transfer of human expertise, *Proceedings AAAI-84*, 27-33, California : American Association for Artificial Intelligence, 1984.

BOOSE J.H., *Expertise Transfer for Expert System Design*, New York : Elsevier, 1986.

BOOSE J.H., Rapid Acquisition and Combination of Knowledge from Multiple Experts in the Same Domain, *Future Computing Systems 1* (2), 191-216, 1987.

BOOSE J.H., A survey of knowledge acquisition techniques and tools. Boose J.H. & Gaines B.R. (eds), *Proceedings of the Third AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop*, pp. 3-1-3-23, Banff, November 1988.

BOOSE J.H. & BRADSHAW J.M., 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, 1987.

BOOSE J.H. & GAINES B.R. (Eds), *Knowledge Acquisition Tools for Expert Systems*, London, Academic, Press, 1988.

BROADBENT D.E., FITZGERALD P. & BROADBENT M.H.P., Implicit and explicit knowledge in the control of complex systems, *British Journal of Psychology 77*, 33-50, 1986.

CENDROWSKA J., An algorithm for inducing modular rules, *International Journal of Man-Machine Studies 27* (4), 349-370, October 1987.

COLLINS H.M., *Changing Order : Replication and Induction in Scientific Practice*, London : SAGE, 1985.

DIEDERICH J., RUHMANN I. & MAY M., KRITON : A knowledge acquisition tool for expert systems. Boose J.H. & Gaines B.R. (Eds), *Proceedings of the First AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop*, pp. 12-0-12-12, Banff, November, 1986.

DIXON N., *Preconscious Processing*, Chichester : Wiley, 1981.

'SHELMAN L., EHRET D., McDERMOTT J. & TAN M., MOLE : A tenacious knowledge acquisition tool. Boose J.H. & Gaines B.R. (Eds), *Proceedings of the First AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop*, pp. 13-0-13-12, Banff, November 1986.

FEIGENBAUM E.A., Knowledge Engineering : the Applied Side of Artificial Intelligence, *Report STAN-CS-80-812*, Department of Computer Science, Stanford University, 1980.

FREUD S., *Psychopathology of Everyday Life*, London : Benn, 1914.

GAINES B.R., An overview of knowledge acquisition and transfer, *International Journal of Man-Machine Studies 26* (4), 453-472, April, 1987a.

GAINES B.R., Rapid prototyping for expert systems. Oliff M.D., Ed., *Intelligent Manufacturing : Proceedings from First International Conference on Expert Systems and the Leading Edge in Production Planning and Control*, pp. 45-73, Menol Park, California, Benjamin Cummins, 1987b.

GAINES B.R. & BOOSE J.H., Eds. *Knowledge Acquisition for Knowledge-Based Systems*, London, Academic Press, 1988.

GAINES B.R. & SHAW M.L.G., New directions in the analysis and interactive elicitation of personal construct systems, *International Journal of Man-Machine Studies 13* (1), 81-116, July 1980.

GAINES B.R. & SHAW M.L.G., From fuzzy sets to expert systems, *Information Science 36* (1-2), 5-16 July, 1985.

GAINES B.R. & SHAW M.L.G., Induction of inference rules for expert systems, *Fuzzy Sets and Systems*, 8 (3), 315-328, April, 1986.

GAMMACK J.G. & YOUNG R.M., Psychological techniques for eliciting expert knowledge. Bramer M., Ed. *Research and Development in Expert Systems*, pp. 105-116, Cambridge University Press, 1985.

GARG-JANARDAN C. & SALVENDY G., A conceptual framework for knowledge elicitation, *International Journal of Man-Machine Studies*, *26* (4), 521-531, April,1987.

HAWKINS D., An analysis of expert thinking, *International Journal of Man-Machine Studies*, *18* (1), 1-47, January, 1983.

HAYES-ROTH F., WATERMAN D.A. & LENAT D.B., Eds., *Building Expert Systems*, Reading Massachusetts : Addison-Wesley, 1983.

KELLY G.A., *The Psychology of Personal Constructs*, New York, Norton, 1955.

KLINKER G., BENTOLILA J., GENETET S., GRIMES M. & McDERMOTT J., KNACK-reportdriven knowledge acquisition. Boose J.H. & Gaines B.R. (Eds), *Proceedings of the First AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop*, pp. 23-0-23-13, Banff, November, 1986.

MARCUS S., Taking backtracking with a grain of SALT, *International Journal of Man-Machine Studies 26* (4), 383-398, April 1987.

MARCUS S., *Automating Knowledge Acquisition for Expert Systems*, New York : Kluwer, 1988.

MUSEN M.A., FAGAN L.M., COMBS D.M. & SHORTLIFFE E.H., Using a domain model to drive an interactive knowledge editing tool. Boose J.H. & Gaines B.R. (Eds), *Proceedings of the First AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop*, pp. 33-0-33-11. Banff, November 1986.

NISBETT R.E. & WILSON T.D., Telling more than we can know : verbal reports on mental processes, *Psychological Review 84*, 231-259.

RAPPAPORT A., Multiple-problem subspaces in the knowledge design process, *International Journal of Man-Machine Studies 26* (4), 435-452, April, 1988.

RAPPAPORT A. & GAINES B.R., Integration of acquisition and performance systems, *Proceedings of the Third AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop*, pp. 25-1-25-20, Banff, November, 1988.

QUINLAN J.R., Simplifying decision trees, *International Journal of Man-Machine Studies 27* (3), 221-234, September, 1987.

ROGERS C.R., *On Becoming a Person : A Therapist's View of Psychotherapy*, London : Constable, 1967.

SHAW M.L.G., *On becoming a Personal Scientist*, London : Academic Press, 1980.

SHAW M.L.G., *Recent Advances in Personal Construct Technology*, London : Academic Press, 1981.

SHAW M.L.G. & GAINES B.R., A computer aid to knowledge engineering, *Proceedings of British Computer Society Conference on Expert Systems*, 263-271, December, Cambridge, 1983.

SHAW M.L.G. & GAINES B.R., Techniques for knowledge acquisition and transfer, Boose J.H. & Gaines B.R. (Eds), *Proceedings of the First AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop*, pp. 39-0-39-13, Banff, November, 1986.

SHAW M.L.G. & GAINES B.R., Interactive elicitation of knowledge from experts, *Future Computing Systems, 1* (2), 151-190, 1986.

SHAW M.L.G. & GAINES B.R., An interactive knowledge elicitation technique using personal construct technolgy. Kidd A., Ed. *Knowledge Elicitation for Expert Systems : A Practical Handbook*, pp. 109-136, Plenum Press, 1987a.

SHAW M.L.G. & GAINES B.R., KITTEN : Knowledge Initiation & Transfer Tools for Experts & Novices, *International Journal of Man-Machine Studies 27* (3), 251-280, September, 1987b.

SHAW M.L.G. & GAINES B.R., A methodology for recognizing consensus, correspondence, conflict and contrast in a knowledge acquisition system. Boose J.H. & Gaines B.R. (Eds), *Proceedings of the Third AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop*, pp. 30-1-30-19. Banff, November, 1988.

SHAW M.L.G. & WOODWARD J.B., Validation of a knowledge support system. Boose J.H. & Gaines B.R. (Eds), *Proceedings of the Second AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop*, pp. 18-0-18-15, Banff, October, 1987.

SLATER P., Ed., *Dimensions of Intrapersonal Space : Volume 1*, London : John Wiley.

WAHL D., An application of declarative modeling to aircraft fault isolation and diagnosis. Luker P.A. & Adelsberger H.H., Eds, *Intelligent Simulation Environments*, pp. 25-28, January 1986, La Jolla, California : Society for Computer Simulation.

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