KNOWLEDGE ACQUISITION TECHNIQUES FOR KNOWLEDGE-BASED SYSTEMS

Mildred L. G. Shaw
Dept of Computer Science
University of Calgary
Calgary, Alberta
Canada T2N 1N4

ABSTRACT

Growing recognition of the significance of knowledge-based computing systems has focused attention on processes of knowledge acquisition and transfer. Commercial applications of expert systems are being impeded by the knowledge engineering bottleneck and have led to the development of rapid prototyping tools. These have proved practically useful but the range of application of existing tools is limited, and it is not clear what limitations are inherent in the prototyping tools. This paper introduces the concept of knowledge support systems as the integration of interactive knowledge acquisition systems and expert systems shells. A prototype knowledge support system that has been implemented is described with examples of some of the knowledge acquisition and application tools provided. It includes knowledge acquisition and transfer through interactive repertory grid elicitation; exchange techniques for understanding and agreement between experts; and different forms of analysis that have been widely used for comparing data between experts and for rapid prototyping of expert systems. An evaluation model for knowledge support systems is presented, and the results of initial validation studies are reported showing the extent to which experts agree with each other and with themselves at a later date. Preliminary results indicate that the measures used discriminate the differences within and between experts.

KEYWORDS: knowledge acquisition, knowledge support systems, rapid prototyping tools, validation.

INTRODUCTION

Much successful work has been done in the area of knowledge representation, but knowledge acquisition and transfer has long been thought of as the bottleneck in the process of building expert systems. It was possible that knowledge engineering might develop as a profession on a par with systems analysis and programming, and that an initial shortage of skilled knowledge engineers would cause problems which would be overcome eventually as the profession developed. However, this scenario now appears less and less likely. There is certainly a shortage of knowledge engineers and problems in developing applications, but doubts have been cast on the notion that human labor is the appropriate solution to the knowledge engineering problem.

The technology is now in a mass-market situation where many organizations see the need for expert systems. This has led to a growth in demand that is far more rapid than the growth in supply of trained and experienced knowledge engineers. In addition, the role of the knowledge engineer as an intermediary between the expert and the technology is being questioned not only on cost grounds but also in relation to its effectiveness. 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. Full exploitation of the potential of expert systems depends on the development of rapid prototyping systems directly usable by experts with the knowledge engineer acting only as manager, not intermediary.

PERSONAL CONSTRUCT PSYCHOLOGY

Kelly developed a systemic theory of human cognition based on the single primitive of a construct, or dichotomous distinction. For an individual, constructs are:

> "transparent templets which he creates and then attempts to fit over the realities of which the world is composed." (Kelly 1955)

He proposes that all of human activity can be seen as a process of anticipating the future by construing the replication of events:

> "Constructs are used for predictions of things to come, and the world keeps rolling on and revealing these predictions to be either correct or misleading. This fact provides a basis for the revision of constructs and, eventually, of whole construct systems." (Kelly 1955)

Hence his psychological model of man is strongly epistemological and concerned with the way in which man models his experience and uses this model to anticipate the future. The anticipation may be passive as in prediction, or active as in action.

Kelly developed his theory in the context of clinical psychology and hence was concerned to have techniques which used it to by-pass cognitive defenses and elicit the construct systems underlying behaviour. This is precisely the problem of knowledge engineering. His repertory grid (Shaw 1980) is a way of representing personal constructs as a set of distinctions made about elements relevant to the problem domain. In clinical psychology this domain will often be personal relationships and the elements may be family members and friends. In the development of expert systems the elements will be key entities in the problem domain such as oil-well sites or business transactions.

Repertory grids have been widely used: in clinical psychology (Shepherd & Watson 1982); to study processes of knowledge acquisition in education (Pope & Shaw 1981); and to study decision making by individuals and groups in management (Shaw 1980). PLANET (Shaw 1982) is an integrated suite of programs that operationalizes Kelly's work and may be used for the interactive elicitation and analysis of repertory grids. These programs have been widely used internationally in clinical psychology, education and management studies (Shaw 1981), and in knowledge engineering for expert systems (Shaw & Gaines 1983).

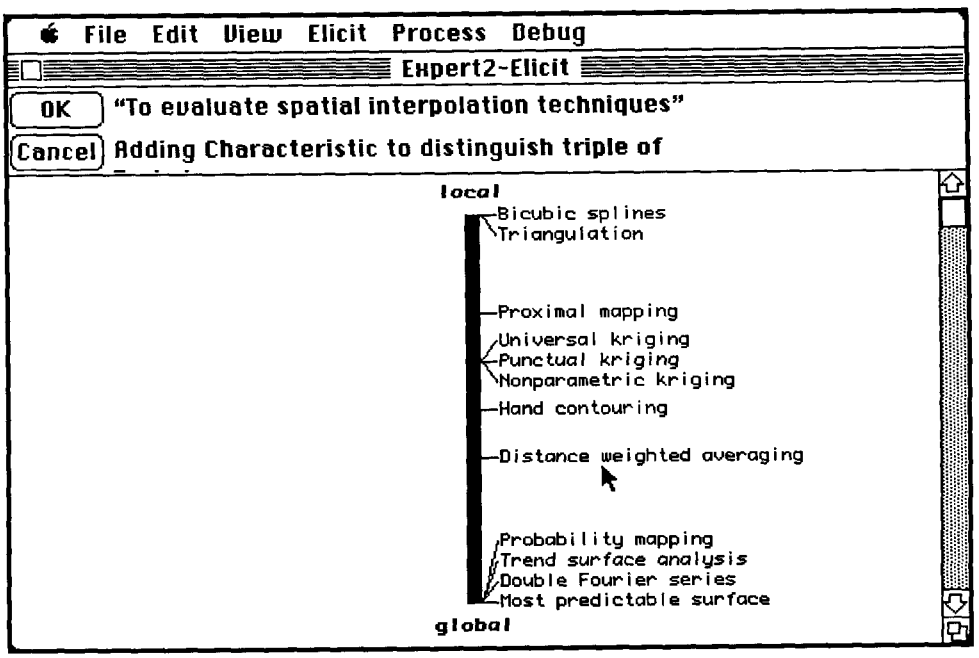

Figure 1. A KSS0 screen showing the construct local — global.

Kelly's personal construct psychology is important because it develops a complete psychology of both the normal and abnormal, which has strong systemic foundations. In the long term these foundations may be more important to knowledge engineering than the techniques currently based on them. However, this paper concentrates on the repertory grid as a technique for eliciting information from an expert.

REPERTORY GRIDS

A repertory grid is a two-way classification of data in which events are interlaced with abstractions in such a way as to express part of a person's system of cross-references between his personal observations or experience of the world (elements), and his personal constructs or classifications of that experience.

The elements are the things which are used to define the area of the topic, and can be concrete or abstract entities. For example, in the context of expertise about metal joining they might be types of rivet, or in expertise about medical diagnosis they might be symptoms. Before choosing the set of elements, the user must think carefully about the area of the topic and relate the elements to his purpose. The elements should be of the same type and level of complexity, and span the topic as fully as possible. It is usual to start with about six to twelve elements. The universe of discourse is determined by the elements which must be significant to the person in the context of the particular problem.

The constructs are the terms in which the elements are similar to or different from each other. Each construct therefore has two poles, each of which has a meaning with respect to its opposite. Any construct or dimension of thinking which is important to the subject is a valid construct. For example, to distinguish between people by saying that x and y are **blue-eyed** whereas b and c are **brown eyed** may be trivial, and not concerned with the important qualities of x, y, b, and c. However, if you are an eye specialist concerned with prescribing tinted contact lenses, this may be a significant construct. Thoughts and feelings, objective and subjective descriptions, attitudes and rules-of-thumb all constitute valid constructs if they contribute to the conceptual structures of the expert. The verbal description of the construct and the labelling of the poles need not be a publically agreed meaning in the outside world, but only a memory aid to the thinking process. The mapping of the elements onto the constructs produces the two-dimensional grid of relationships.

Constructs are elicited either in triples — in what significant way is one element similar to another, and in the same way different from a third; in pairs — in what significant way do these two elements differ; or just as they come to the mind of the expert. An example is shown of elements rated on the construct **local—global** using direct manipulation in the area of spatial interpolation techniques to produce contour maps, in Figure 1.

Figure 2 shows how the techniques are dragged on to the construct bar from the left. Both the construct labels and the technique names can be edited at any time. The placing of the techniques should be continued until the expert feels that their relative positions captures the idea he had in mind.

The mapping of the elements onto the constructs produces the two-dimensional grid of relationships which can be represented as a numeric data structure as shown in Figure 3. This structure may be viewed as a component of a database in entity-attribute form (Chen 1980): a repertory grid has elements as entities, constructs as attributes and allocations of elements to poles of constructs as values. For example, **Bicubic splines** is rated a 1 on the construct **local—global** specifying **local**, whereas **Probability mapping** is rated a 9, specifying **global**.

KNOWLEDGE SUPPORT SYSTEMS

Knowledge support systems can be seen as integrated interactive knowledge acquisition systems and expert systems shells. They are seen as having a broad scope, encompassing both aids to knowledge acquisition and transfer, knowledge representation, and support of human knowledge processes (Shaw & Gaines 1987). It is desirable that the tools in a knowledge support 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 support system should be able to encompass a

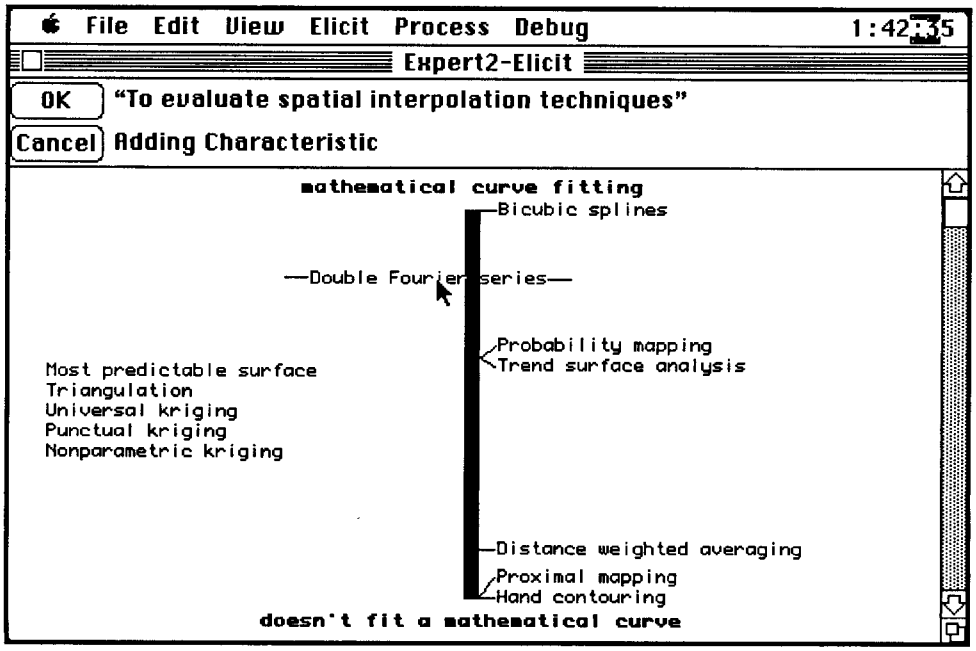

Figure 2. A KSS0 screen showing direct manipulation of elements.

Entities: 12, Attributes: 18, Range: 1 to 9, Purpose: To evaluate spatial interpolation techniques

		1	2	3	4	5	6	7	8	9	10	11	12		
qualitative and quantitative	1	8	8	6	9	1	9	8	8	9	4	4	4	1	quantitative
local	2	9	9	6	3	5	1	9	9	1	4	4	4	2	global
autocorrelation not considered	3	1	1	5	4	7	2	3	1	2	9	9	9	3	autocorrelation considered
doesn't honour data points	4	2	1	3	5	9	9	1	1	9	8	7	7	4	honours data points
multiple variables considered	5	9	9	9	9	9	9	9	1	9	9	9	9	5	usually one variable considered
mathematical curve fitting	6	4	4	8	9	9	1	2	4	9	5	8	8	6	doesn't fit a mathematical curve
nonparametric	7	9	8	6	1	1	3	6	8	3	8	8	1	7	parametric
interval or ratio data	8	9	1	1	5	5	1	1	1	1	1	1	1	8	nominal data
requires periodicities	9	6	6	9	9	9	9	1	6	9	9	9	9	9	doesn't require periodicities
doesn't fit a trend	10	9	9	1	1	8	1	7	9	1	6	1	1	10	fits a trend to the data
heavy computing load	11	7	6	7	8	9	4	4	5	3	1	2	3	11	no computing load
assumes isotropic surface	12	1	1	4	3	8	9	1	1	9	7	6	6	12	assumes anisotropic surface
not as susceptible to clusters	13	8	8	6	9	3	4	7	8	5	1	2	2	13	estimates susceptible to clusters
doesn't incorporate geologic model	14	2	2	3	1	9	1	2	2	1	6	6	6	14	incorporates geologic model
interpretive	15	9	9	5	9	1	7	9	9	7	3	4	4	15	representative
not very important	16	4	5	2	1	9	1	1	6	7	8	9	9	16	very important
not very effective	17	4	5	3	1	9	2	4	6	7	9	8	8	17	very effective
not widely used	18	3	8	7	5	9	3	3	4	6	2	2	1	18	widely used

	1	2	3	4	5	6	7	8	9	10	11	12
	Probability mapping	Trend surface analysis	Distance weighted averaging	Proximal mapping	Hand contouring	Bicubic splines	Double Fourier series	Most predictable surface	Triangulation	Universal kriging	Punctual kriging	Nonparametric kriging

Figure 3. A conventional repertory grid of spatial interpolation techniques. A 1 represents the left construct pole and a 9 the right.

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 support 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 support system as possible should be founded on well-developed and explicit theories of knowledge acquisition, elicitation and representation. As the overall knowledge support system develops it should converge to an integrated system.

Clearly, these are general requirements and may not apply in particular cases. For example, some domain dependence may be appropriate for efficiency in specific knowledge support systems; some human intervention may be helpful or necessary when an expert is using a knowledge support system. However, in general these requirements capture the key issues in knowledge support system design.

PLANET (Shaw 1982, Shaw & Gaines 1986) has been developed and widely used for some years for knowledge elicitation. This paper briefly reviews the techniques involved and describes the most recent developments of a comprehensive integrated rapid prototyping system, KSS0, which attempts to satisfy the above design requirements for a knowledge support system.

RELATED WORK

ETS and Aquinas (Boose 1984, Boose & Bradshaw 1987) are related quite closely to PLANET and KSS0, having been developed in response to the initial success of the PLANET system. KSS0 and Aquinas are integrated prototyping systems providing knowledge acquisition tools encompassing a diversity of forms of knowledge and relationships between knowledge. KSS0 can access a wide range of knowledge sources including text, interviews with experts, and observations of expert behaviour. Aquinas can present knowledge from multiple sources with clarity as to its derivation, consequences and structural relations. Both systems can encompass a diversity of perspectives including partial or contradictory input from different experts. Users of these knowledge acquisition tools are able to apply the knowledge in a variety of familiar domains and freely experiment with its implications. These systems offer the capability to expedite the prototyping stage of complex knowledge-based systems, motivating the experts to be closely involved in all aspects of the system development by giving them a supportive and comprehensible environment. All four systems use repertory grid techniques to elicit the expert's conceptual structures, and use different versions of ENTAIL (Gaines & Shaw 1986) to produce rules for an expert system shell.

MOLE (Eshelman, Ehret, McDermott, & Tan, 1987) is an expert system shell that has been used effectively to build heuristic classification systems. It interviews domain experts, exploits assumptions of exhaustiveness and exclusivity to determine the most likely candidates among competing hypotheses, and has many methods for analyzing and refining the knowledge base for consistency and adequacy. It distinguishes between three types of knowledge—covering, differentiating, and combining knowledge—in its problem-solving, and uses a system of "support values" to represent uncertainty.

CSRL (Bylander & Mittal, 1986) is another successful tool with a similar problem-solving approach, although it is difficult for domain experts to learn and use. Gruber and Cohen (1987) have developed a system called MUM that embodies architectural principles that facilitate knowledge acquisition. It allows experts to specify local evidence combining knowledge; however it is limited to a symbolic representation of uncertainty.

An innovative "learning apprentice" program, ODYSSEUS, has been developed by Wilkins (1987) to help experts refine and debug knowledge bases for the HERACLES heuristic classification shell. However, its use is appropriate only during the knowledge acquisition "end-game"; that is, it requires that a reasonable knowledge base has already been created.

The following section describes the work on KSS0, a knowledge support system that draws on many concepts and techniques for knowledge engineering to begin to encompass some of the requirements discussed above, attempting to relate them all through the theoretical base of personal construct psychology, and build a workbench of tools around a common database.

KSS0: A KNOWLEDGE SUPPORT SYSTEM

KSS0: Knowledge Support System Zero consists of a: knowledge base; various analytical tools for building and transforming the knowledge base; and a number of conversational tools for interacting with the knowledge base. The KSS0 implementation is written in Pascal and currently runs on a network of Macintosh workstations. The KSS0 structure is best understood by following sequences of activity that lead to the generation of a rule base and its loading into an expert system shell.

A typical sequence is text input followed by text analysis through TEXAN which clusters associated words leading to a schema from which the expert can select related elements and initial constructs with which to commence grid elicitation. The resultant grids are analyzed by ENTAIL which induces the underlying knowledge structure as production rules that can be loaded directly into an expert system shell (Gaines & Shaw 1986).

An alternative route is to monitor the expert's behaviour through a verbal protocol giving information used and decisions resulting and analyze this through ATOM which induces structure from behaviour using a search over a model space ordered by complexity and goodness of fit, and again generates production rules (Gaines 1977). A version of ATOM is incorporated in KSS0 that takes a set of sequences of arbitrary symbolic data and generates a set of production rules that will reconstruct it. These can be loaded into the expert system shell to give a simulator of the behavioural system.

These two routes can be combined. KSS0 attempts to make each stage as explicit as possible to the expert, and, in particular, to make the rule base accessible as natural textual statements rather than technical production rules. The expert system shell being used in KSS0 currently is Nexpert (Roy 1986) which gives a variety of textual and graphical presentations of the rule base enabling the expert to see the impact of different fragments of knowledge.

The group problem-solving component of KSS0 is particularly important because it goes beyond the stereotype of an **expert** and **users**, and allows the system to be used to support an interactive community in their acquisition and transfer of knowledge and mutual understanding. The SOCIO analysis allows members of a community to explore their agreement and understanding with other members, and to make overt the knowledge network involved (Shaw 1980, 1981, 1988). Much expertise only resides within the social context of cooperating individuals and requires elicitation across the group. The SOCIO analysis program supports group elicitation techniques in which the construct systems of a number of experts are compared. Grids are elicited separately but then exchanged in two ways: an expert can place elements on a colleague's constructs from his own point of view, and the analysis system then allows him to explore their agreement; or he can attempt to place them from his colleague's point of view and hence explore his understanding. SOCIO supports several users in seeing the relationship of their points of view to those of others; exploring

differing terminology for the same constructs; becoming aware of differing constructs having the same terminology; extending their own construct systems with those of others; providing others with constructs they have found valuable; and exploring a problem-solving domain using the full group resources.

The KSS0 implementation is an initial prototype offering a workbench with minimal integration of the knowledge base, but each of the tools has already proven effective, and their combination is proving very powerful in stimulating experts to think of the knowledge externalization process from a number of different perspectives.

VALIDATION OF A KNOWLEDGE SUPPORT SYSTEM

A model for evaluation of a knowledge support system and different types of validation has been discussed by Shaw and Woodward (1988). There are two distinct aspects of validation, those concerned with objective aspects of the system and those concerned with subjective aspects of the system. These distinctions represent the various aspects of operation of knowledge support systems. The objective aspects are those concerning the performance of the expert system or knowledge base, elicited from the expert or experts, against objective criteria. A knowledge support system supports the knowledge processes of the expert, and allows her to develop a system which actually performs effectively in the real world, rather than merely confirming the opinion of the expert.

The subjective aspects relate to the way in which the knowledge acquisition system acquires the knowledge that the expert has. It is clearly impossible to separate the knowledge acquisition tool from the expert, so the term **knowledge acquisition system** is used to mean the knowledge acquisition tool, **and** the expert:

knowledge acquisition system = knowledge acquisition tool + expert

Figure 4 shows the model developed for this study. Only the subjective aspects of validation were investigated.

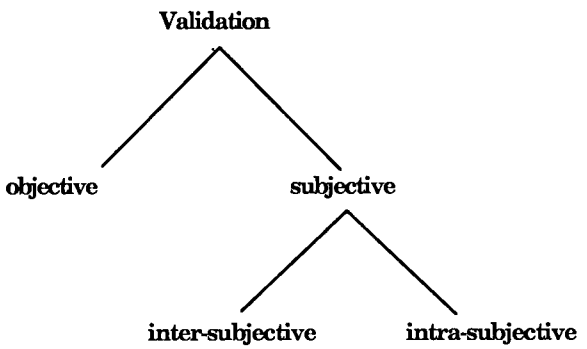

Figure 4. The evaluation model for the study

It is now thought to be essential to have a group of experts involved in building a knowledge base, so it was necessary to see if there was mutual consistency and/or construing across experts. Also, whether there was any consistency of an expert over a time interval, both in the terminology used, and how it was used. Finally, an investigation was made as to whether the analysis performed by the system on each expert's knowledge was seen as correct, or if any of it was incorrect.

A measure of diachronic reliability (Chow, 1987) appears most suitable for evaluating the **replication** of knowledge support systems. This measure reflects the degree of similarity, over time, of the calculated output. A knowledge support system reflects a high degree of replication in the sense that the output of the knowledge support system, the results of the algorithms, reflect a high degree of similarity over time. Strictly speaking, this approach to replication is strongly influenced by the reliability of the user of the knowledge support system. However, the knowledge support system can be viewed as an aid to the user which helps the user maintain consistency over time given the same task demands. Also, knowledge support systems are not used in isolation but are used by individuals so any measure of a knowledge support system must be done with those who will use the system. If replication measures are low, the system cannot be considered consistent or helpful, unless this can be seen to be a product of the raw data from the experts.

Although **replication** is a possible requirement of a knowledge support system, other basic requirements need attention. The knowledge support system must also demonstrate various forms of **consistency**. In this sense, consistency refers to the extent to which the knowledge support system produces what it is expected to produce. At the subjective level of evaluation, consistency is defined by two measures. The first is **intra-subjective consistency**. Based on the concepts of content and face validity (Travers, 1969; Campbell & Stanley, 1963), this type of consistency is defined as the extent to which the knowledge support system output is correct as judged by the expert. Knowledge support systems produce output which has a high degree of similarity with what is apparently expected. The absence of this form of consistency indicates that the system may be performing its calculation incorrectly and/or that the output bears little resemblance to what would be expected by the designer, by the user of the tool or by the domain expert. This result has serious implications for any subsequent use of knowledge support system output.

Another, complementary, measure of consistency at this level of evaluation is that of **inter-subjective consistency**. Based on the concept of **concurrent validity** (Chow, 1987; Travers, 1969) this type of consistency is a measure of prediction of events which occur at roughly the same time. If two or more individuals, with similar experience in an area, produce similar output using a knowledge support system, then the system can be said to afford a degree of inter-subjective consistency. The extent to which two or more people produce the the same output represents the degree of consistency.

However, inconsistency among experts is not unexpected as experts may well not agree. In many areas, each expert has a unique perspective on the topic and often believes that there is only one way that the topic should be considered. This is not necessarily an undesirable situation, as long as the system is able to reflect this. If a task is carefully defined, specifically outlined and agreed upon as important, then a certain degree of similarity can be expected in the outcome of the knowledge support system. This suggests the selection and development of a set of test cases with specific task requirements. These test cases can then be used to determine the degree of inter-subjective consistency demonstrated by various experts, while using the system. In the same way as for replication, the knowledge support system can be seen as an aid to users.

THE RESULTS

KSS0 has been evaluated against the model and the results reported in the two domains of spatial interpolation techniques to produce contour maps and in trouble-shooting and maintenance of valves for oil and gas pipelines. Some preliminary results are described on validation experiments to show the extent to which this system can replace standard interviewing techniques.

Two separate studies were completed using specific task demands from these two separate and quite different domains. The first study involved three researchers in a university geography department each of whom is an expert in the field of spatial interpolation techniques to produce contour maps. In the second study, a trainer and an engineer in the field of gas pipeline valve maintenance provided their expertise. These two

individuals were chosen because they were both considered knowledgeable in the area, but also because they brought a different background to the field: one from a field-operator perspective, the other from that of an engineer. The trainer had been involved in training gas pipeline operators in proper maintenance of pipeline valves for 15 years and the engineer owned a company which sold valves. Each group elicited and exchanged grids on the first occasion, then 2 to 3 months later elicited new constructs and ratings using the old set of elements, and elicited new ratings using the old set of elements and constructs. Details of the design are given in Shaw and Woodward (1988).

Two measures were used in this study, **consistency-with** another and **construed-by** another. **Consistency** is measured using **exchange grids**. Do the experts see the topic in a similar way at the same time? By exchanging elements and constructs they are able to view the topic from the perspective of the other. To do this involves an understanding of what the other's elements and constructs actually **mean** in the domain, or **can I construct a point of view which makes sense of them?** Not only that, but **do I agree with the other's construction of the topic, having understood the perspective of the other?** So **consistency** is defined as the degree of match between one expert's ratings of his/her own elements and constructs and another expert's ratings of the first expert's elements and constructs. The patterning of the elements on a particular construct is matched against the patterning of the same elements on that same construct in the other grid.

Consistency over time can also be measured in this way but using the **same** expert's grid on the two occasions with the same element and construct labels as before, but new ratings. Only elements and constructs with the same labels in both grids can contribute to this measure of consistency. Operationally, this measure was calculated using the EXCHANGE measure based on the matching algorithm that is described in Shaw (1980) and used in FOCUS, CORE and SOCIOGRIDS. It ranges from 100 for perfect match to 0 for maximum reversal. This measure is picking out the differing use of terminology in the two grids. The following formula was used:

$$\textbf{G consistency-with G' at criterion} = 100 \times \frac{\text{(number of constructs in G matched greater than criterion by same constructs in G')}}{\text{(number of constructs in G)}}$$

Construing is akin to consistency. The entity labels must be the same in both grids, but separate sets of constructs can be used. In this case, the measure picks out the constructs in the second grid with the best match to the patterning of the elements on the one being considered in the first grid. In this case, the same construct may be chosen more than once in the second grid if it matches more closely than any other to more than one construct in the first. Again, the same matching algorithm was used. This measure is picking out the best match between the constructs in the two grids, regardless of terminology. The following formula was used:

$$\textbf{G construed-by G' at criterion} = 100 \times \frac{\text{(number of constructs in G matched greater than criterion by any constructs in G')}}{\text{(number of constructs in G)}}$$

Operational definitions for the various construing and consistency measures were used to evaluate the knowledge support system. Intra-subjective understanding was defined as the degree of construing of the earlier grid by the later one. Those constructs which were matched on the two grids at a criterion of 80 were used for this calculation, and this was calculated for each expert. The results ranged from 63 to 86 in the first study, and 83 to 100 in the second. This difference in values between the two studies may reflect the fact that the three experts in the first study are academic researchers at a university, and hence tended to be more provisional, exploratory and accepting of different ways of considering the problem. They were at no time fixed in their views of the topic. In the second study, however, the experts were professionals using their expertise in the field, and hence much more decisive and sure of what they were doing and how it should be done.

Intra-subjective agreement was defined as the degree of consistency between one expert's ratings of his/her own elements and constructs in the earlier grid with his/her ratings of the elements and constructs in the later grid. These scores ranged in the first study from 79 to 94, indicating that the experts were using their terminology in much the same way from one occasion to the other and that KSS0 was able to reflect this consistency. In the second study both were 67, indicating some difference in the use of terminology.

Inter-subjective agreement was defined as the degree of consistency between one expert's ratings of his/her own elements and constructs in the first grid with another expert's ratings of the first expert's elements and constructs; that is the extent to which the experts agree. These scores ranged in the first study from 8 to 33.showing that the experts disagree in their terminology and in how they view the topic quite extensively. This can also be seen by inspecting the raw data. In the second study the scores range from 13 to 100. Here it is interesting to note that the trainer could use the terminology of the engineer totally, but the engineer could only use a small part of the trainer's terminology.

Inter-subjective understanding was defined as the degree of construing of each expert's grid by another expert. These values are much higher showing that the experts have similar distinctions about the topic even though they differ greatly in how they use the terminology. In the first study these values range from 31 to 71, and in the second from 67 to 100. This suggests that KSS0 shows a facility for allowing experts to compare their understanding and to determine a level of consistency. Whether the threshold criteria are high enough to produce the required level of knowledge-base performance or too high to reach the criterion of completeness is a matter for further study.

Finally, intra-subjective perspective consistency was defined as the degree of consistency of the rules produced by the ENTAIL algorithm with what the expert apparently expected. The ratings of these statements were then compared to the entailment output from the earlier grid. These show whether the expert finds the rules meaningful, and rates the rules as correct or significant in the same way as the ENTAIL algorithm (Gaines & Shaw 1986, Shaw & Gaines 1987). They range from 82 to 92, showing a high level of expected rules appearing in the KSS0 output. This indicates a good basis for using the ENTAIL produced rules as input to an expert system shell.

There are many more ways in which the data are being processed, especially in terms of the rules derived from the various grids. At this stage the similarities and differences between experts and within experts over a 2 to 3 month time interval have been investigated. Clearly, the content of these similarities and differences can be identified and related to external criteria.

CONCLUSIONS

Growing recognition of the significance of knowledge-based computing systems has focused attention on processes of knowledge acquisition and transfer. Commercial applications of expert systems are being impeded by the knowledge engineering bottleneck and have led to the development of rapid prototyping tools. These have proved practically useful but the range of application of existing tools is limited, and it is not clear what

limitations are inherent in the prototyping tools, what arises from the shells and their knowledge representation and inferencing procedures, and what arises from our lack of understanding of the underlying processes of knowledge acquisition and transfer.

This paper has introduced the concept of knowledge support systems as integrated interactive knowledge acquisition systems and expert systems shells. A prototype knowledge support system that has been implemented on a network of Macintosh computers has been described with examples of some of the knowledge acquisition and application tools provided. It includes knowledge acquisition and transfer through interactive repertory grid elicitation; exchange techniques for understanding and agreement between experts; and different forms of analysis that have been widely used for comparing data between experts and for rapid prototyping of expert systems. Such systems offer the capability to expedite the prototyping stage of complex knowledge-based systems, motivating the experts to be closely involved in all aspects of the system development by giving them a supportive and comprehensible environment.

An evaluation model for knowledge support systems was presented, and the results of initial validation studies reported. Measures have been put forward which tease out differences in terminology between experts, and show the content of this difference. To a knowledge engineer, it seems that if experts talk about their topic of expertise in the same terms then they mean the same thing, and if they talk in different terms they mean different things. However, it can be seen from the results of this study that the knowledge engineer, or even another expert, may be mistaken. The experts in the first study were clearly accustomed to disagreeing on terminology, a fact which may not always be made explicit. The consistency measure used for inter-subjective agreement showed up the divergence in use of terminology between the experts (8 to 33), but when used for one expert alone for intra-subjective agreement gave much higher values (79 to 94). This measure, then, shows a good range of discrimination in the two cases within experts and between experts. This is not so apparent in the second study where the experts were professionals using their skills and knowledge every day in an industrial situation. A similar rationale can be applied in the construed-by measures of inter- and intra-subjective understanding.

In the first study there was clearly inconsistency among the experts, and this was reflected by the results from the knowledge support system. This was therefore a good population for study, as it highlighted a controversy and/or disagreement which would be likely to occur between experts and users of an expert system, even if it did not occur among the experts themselves.

ACKNOWLEDGEMENTS

Financial assistance for this work has been made available by the Natural Sciences and Engineering Research Council of Canada. The KSS0 system is made available by the Centre for Person Computer Studies. I would like to thank my colleagues Brian Gaines and Brian Woodward who were jointly involved with me in conducting this research.

REFERENCES

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

Boose, J.H. & 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**(1), 3-28.

Bylander, T. & Mittal, S. (1986). CSRL: A language for classificatory problem-solving and uncertainty handling, **AI Magazine** (August).

Campbell, D.T. & Stanley, J.C. (1963) **Experimental and Quasi-Experimental Designs for Research.** Chicago:Rand McNally and Company.

Chen, P.P., Ed. (1980). **Entity-relationship approach to systems analysis and design**. New York: North Holland.

Chow, S.L. (1987). **Experimental Psychology: Rationale, Procedures and Issues**. Calgary: Detselig Enterprises.

Eshelman, L., Ehret, D., McDermott, J. & Tan, M. (1987). MOLE: A tenacious knowledge acquisition tool. **International Journal of Man-Machine Studies, 26**(1), 41-54..

Gaines, B.R. (1977). System identification, approximation and complexity. **International Journal of General Systems, 3**, 145-174.

Gaines, B.R. (1987) Rapid prototyping for expert systems. Oliff, M. (Ed). **Proceedings of International Conference on Expert Systems and the Leading Edge in Productions, Planning and Control.** pp. 213-241. University of South Carolina.

Gaines, B.R. & Shaw, M.L.G. (1981). New directions in the analysis and interactive elicitation of personal construct systems. Shaw, M.L.G., Ed. **Recent Advances in Personal Construct Technology**. pp. 147-182. London: Academic Press.

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

Gruber, T. & Cohen, P. (1987). Design for acquisition: Principles of knowledge system design to facilitate knowledge acquisition. **International Journal of Man-Machine Studies, 26**(2), 143-159.

Johnson, P.E. (1986).Cognitive models of expertise. **Symposium on Expert Systems and Auditor Judgment.** University of Southern California.

Kelly, G.A. (1955). **The Psychology of Personal Constructs**. New York: Norton.

Landfield, A. (1976). A personal construct approach to suicidal behaviour. Slater, P., Ed. **Dimensions of Intrapersonal Space: Volume 1**. pp. 93-107. London: John Wiley.

Pope, M.L. & Shaw, M.L.G. (1981). Personal construct psychology in education and learning. Shaw, M.L.G., Ed. **Recent Advances in Personal Construct Technology**. pp. 105-114. London: Academic Press.

Roy, J. (1986). Expert systems in Nexpert. **MacTutor, 2**(2), 48-51 (February).

Shaw, M.L.G. (1980). **On Becoming a Personal Scientist**. London: Academic Press.

Shaw, M.L.G. (1981). **Recent Advances in Personal Construct Technology**. pp. 31-44. London: Academic Press.

Shaw, M.L.G. (1982). PLANET: some experience in creating an integrated system for repertory grid applications on a microcomputer. **International Journal of Man-Machine Studies, 17**(3), 345-360.

Shaw, M.L.G. (1988). An interactive knowledge-based system for group problem-solving. **IEEE Transactions on Systems, Man and Cybernetics**, (to appear).

Shaw, M.L.G. & Gaines, B.R. (1980). Fuzzy semantics for personal construing. **Systems Science and Science**. pp. 146-154. Kentucky: Society for General Systems Research.

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

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

Shaw, M.L.G. & Gaines, B.R. (1986b). Techniques for Knowledge Acquisition and Transfer. **Proceedings of the Workshop on Knowledge Acquisition for Knowledge-Based Systems,** November, Banff, 39-0 — 39-13.

Shaw, M.L.G. & Gaines, B.R. (1987). Techniques for knowledge acquisition and transfer. **International Journal of Man-Machine Studies, 27**(3), 251-280.

Shaw, M.L.G. & Woodward, J.B. (1988) Validation in a knowledge support system: consistency and.construing with multiple experts. **International Journal of Man-Machine Studies,** in press.

Shepherd, E. & Watson, J.P., Eds. (1982). **Personal Meanings**. London: John Wiley.

Travers, R.M.V. (1969) **An Introduction to Educational Research (Third Edition).** MacMillan: London.

Wilkins, D.C. (1987). Knowledge base debugging using apprenticeship learning techniques. **International Journal of Man-Machine Studies,** 27(3), 281-294.