A KNOWLEDGE SYSTEM FOR GROUP LEARNING AND DECISION MAKING

by

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

ABSTRACT

This paper gives a state-of-the-art report on the use of techniques based on personal construct psychology to automate knowledge engineering for group learning and decision making. It presents the concept of knowledge support systems as interactive knowledge engineering tools, and outlines the structure and key components of the KITTEN implementation on a network of Macintosh computers. It includes tools for interactive repertory grid elicitation and entailment analysis, tools for text analysis, behavioral analysis and schema analysis, that offer complementary and alternative approaches to knowledge acquisition. The KITTEN implementation integrates these tools around a common database with utilities designed to give multiple perspectives on the knowledge base.

INTRODUCTION

This paper describes a distributed system for human-computer interaction based on a network of computers. It aids group problem solving by enabling the participants to share in a construct elicitation process based on personal construct psychology, and in particular on repertory grid techniques. Since their development in the context of clinical psychology, these techniques have found a wide range of applications in education, management and expert systems development. A similar rationale applies: in education, the learner is attempting to acquire a specific construct system for the subject matter; in management, people with different construct systems are attempting to work together towards common objectives; in expert systems development, the knowledge engineer is attempting to make overt and encode the relevant construction system of an expert. The term **knowledge engineering** will be used in a broad sense to encompass both knowledge elicitation and transfer for knowledge-based systems, and also for the elicitation of knowledge structures for the purposes of learning and decision making.

The manual elicitation of construct systems using structured interview techniques is tedious, time-consuming, and requires skill, practice and an appropriate personality. In recent years computer-based techniques have been developed for on-line elicitation and analysis of construct systems. These have proved very successful and are now in widespread use [1, 2]. One of the most interesting developments of such systems has been to attempt to combine the construct systems of a number of people, make them accessible to all and hence share their knowledge and skills [3, 4, 5, 6]. The design, development and application of such systems requires an understanding of both the computer technology involved and the psycho-socio-dynamics of the users. The construct system design addresses this requirement by designing a specialist system for a domain where large-scale cognitive psychological and sociological theories already exist, implementing it on a network of Macintosh personal computers where an appropriate human-computer interface already exists.

PERSONAL CONSTRUCT PSYCHOLOGY

Kelly's [7] Personal Construct Psychology (PCP) provides a systemic model of the cognitive and behavioral processes of the individual, including inter-personal and social interaction. A person is modeled as **construing** her experience by dividing it up into chunks (elements) and classifying these using a system of bipolar constructs. Constructs are themselves construed by other constructs and hence form a hierarchical system. Construing is used to enable a person to anticipate the future as a hypothesized replication of the past. Constructs are changed as a result of errors in anticipation so that personal philosophies are seen as essentially fallibilist. A person will be more willing to change some constructs rather than other **core constructs**, giving a further order relation on the construct system.
In PCP, learning is seen as the development of the construct system, social processes are seen as the development of shared constructs, and problem-solving is seen as the active anticipation of the future. The theory provides complete foundations for all human psychological and social processes. Hence it is possible to build technologies based upon it which are very general in concept and application. For example, an expert system in PCP terms is a computer-based system in which the relevant construction system of an expert has been encoded in an active form [8]. In computational terms the elements in PCP may be seen as entities, the constructs as attributes and the relation between entities and attributes as values in an entity-attribute database. Kelly terms such a database a **repertory grid** and Fig.1 shows how a typical grid may be represented. The elements in this grid are Artificial Intelligence books used in a course. The three constructs illustrate the mix between subjective and objective possible in personal constructs.

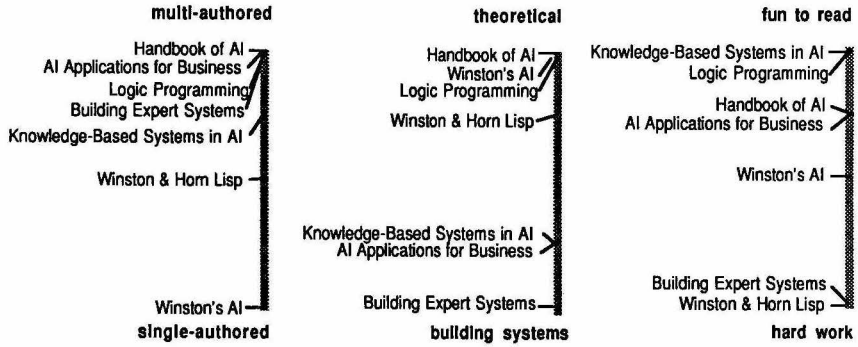


Fig.1 A repertory grid of AI books on three constructs

These considerations have heightened interest in the possibility of providing knowledge support systems to automate knowledge engineering as a process of direct interaction between domain experts and the computer. Knowledge engineering can be defined in very broad terms as: the acquisition, elicitation, structuring and encoding of knowledge for application in inferential, goal-directed, explanatory, decision and action support systems. Knowledge support systems have an even broader scope, encompassing both aids to knowledge engineering and support of human knowledge processes.

The PLANET system for repertory grid elicitation and analysis [1, 9, 10, 11] is a primitive knowledge support system satisfying requirements of domain independence and

direct use by the domain expert. Its foundations in personal construct psychology, which itself has strong systemic and cognitive science foundations [12, 13], are attractive in that the operation of a knowledge support system should be founded on well-developed and explicit theories of knowledge acquisition and elicitation. The following section describes our work on KITTEN, a knowledge support system that draws on many concepts and techniques for knowledge engineering, relating them all through personal construct psychology, and building a workbench of tools around a common database.

KITTEN

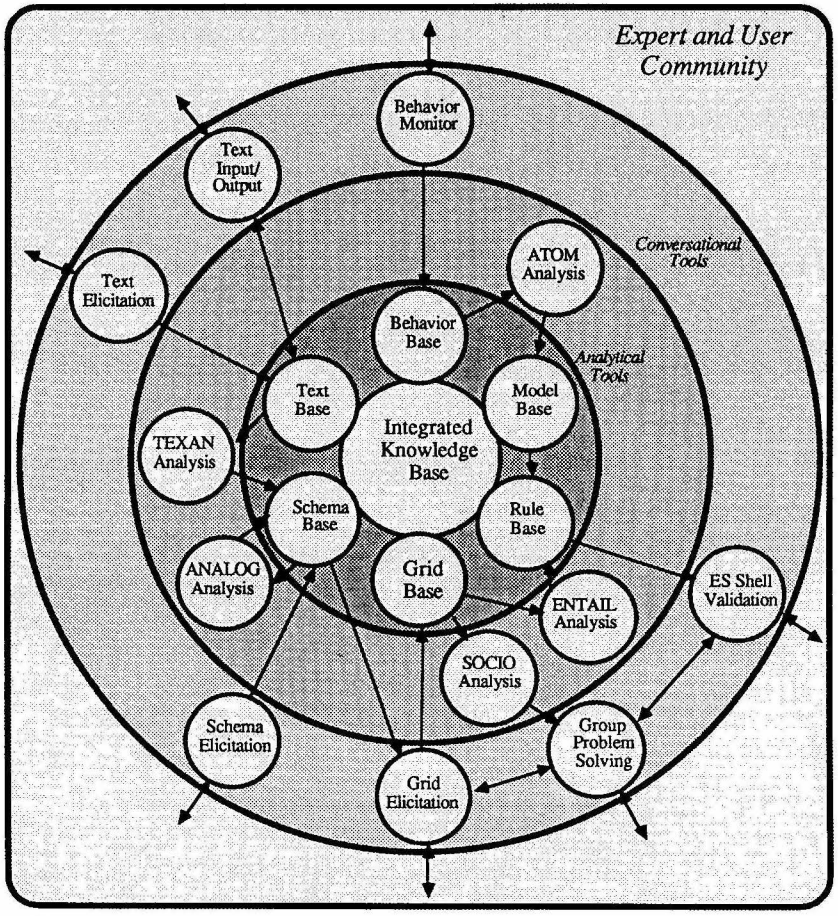


Fig.2 KITTEN—Knowledge Elicitation Tools for Experts and Novices

Figure 2 shows the structure of KITTEN which 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 KITTEN implementation is written in Pascal and currently runs both on a network of Apollo workstations and on a network of Macintosh computers.

A typical sequence is text input followed by text analysis through TEXAN which clusters associated words leading to a schema from which the user 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 ES shell [13].

An alternative route is to monitor the expert's behavior through a verbal protocol giving information used and decisions resulting and analyze this through ATOM which induces structure from behavior and again generates production rules [14].

These two routes can be combined. KITTEN attempts to make each stage as explicit as possible, and, in particular, to make the rule base accessible as natural textual statements rather than technical production rules. The ES shell being used in KITTEN currently is Nexpert [15] which gives a variety of textual and graphical presentations of the rule base enabling the user to see the impact of different fragments of knowledge.

The group problem-solving component of KITTEN is particularly important because it 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 [1, 3].

The KITTEN implementation is an initial prototype offering a workbench with minimal integration of the knowledge base, but each of 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. The following sections describe and illustrate some of the tools.

SOME KITTEN TOOLS

PLANET accesses the expert's personal construct system by interactively eliciting a repertory grid of constructs classifying elements characterizing part of the domain of expertise. Figure 3 shows a repertory grid on staff appraisal. The entailment analysis of a repertory grid treats each pole of a construct as a fuzzy predicate to which the elements have degrees of membership given by their ratings, and induces the logical implications between these predicates. The original ENTAIL program produced all entailments consistent with the grid and allowed the expert to prune any that seemed spurious before using them as inference rules in an expert system. ENTAIL II rank orders entailments in terms of the uncertainty reduction they induce in the distribution of the data, and hence tends to reject spurious entailments [16].

Figure 4 is an ENTAIL II analysis of the grid of Figure 3. The entailments are shown with three values in the range from 0 to 1: first, the truth value of the hypothesis; second, the probability of the hypothesis being true; and third, the information content (uncertainty reduction generated) of asserting the hypothesis. For example, L1→L9 has a truth value of 0.80, a probability of 1.00, and an information content of 0.29. The information content measures the significance of the hypothesis and is used to ensure that trivial entailments consistent with the data are pruned.

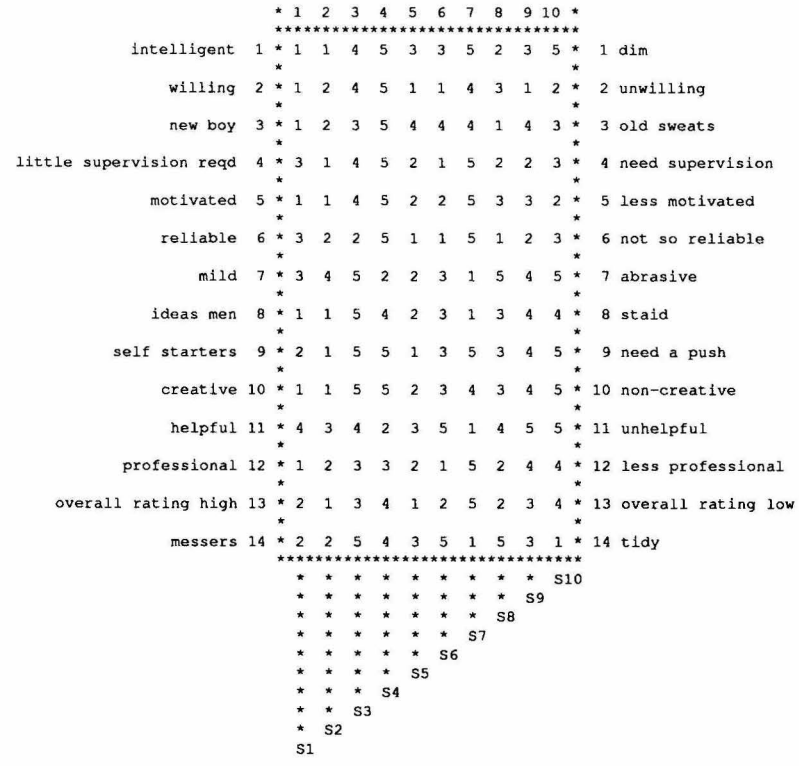


Fig.3 Repertory grid elicited on staff appraisal

The data of Figure 3 may be regarded as that of an expert on staff appraisal concerned with deriving his overall rating (construct 13) from behavioral assessments such as intelligent and creative. The ENTAIL analysis of Figure 4 shows that L1, L4, L6, L9, L10 and L12 imply L13, that intelligent, creative, reliable and professional self-starters requiring little supervision receive a high overall rating, whereas R2, R4, R5, R6, R9 and R12 imply R13, that being unwilling, less motivated, not so reliable, less professional, needing supervision and needing a push leads to a low overall rating.

Figure 5 shows Nexpert in operation loaded with the entailments of Figure 4. Interaction with Nexpert enables the expert to see the derived rules in action. He can determine their consequences with test data, analyze new hypothetical cases, and see the inter-relations between rules presented graphically. The logging and explanation facilities of Nexpert enable him to track down spurious inferences that may arise with the rules derived by ENTAIL, or

Entail	Truth	Prob.	Inf.	(Cutoff 0.17) Implication Usually
L 1 →L 9	0.80	1.00	0.29	intelligent→self starters
L 9 →L13	1.00	1.00	0.29	self starters→overall rating high
R 9 →R 1	0.80	1.00	0.28	need a push→dim
L10 →L 8	1.00	1.00	0.28	creative→ideas men
L 1 →L10	0.80	1.00	0.26	intelligent→creative
R 8 →R10	1.00	1.00	0.26	staid→non-creative
L10 →L 9	0.80	1.00	0.26	creative→self starters
R13 →R 6	0.80	1.00	0.26	overall rating low→not so reliable
L 9 →L10	0.80	1.00	0.24	self starters→creative
R10 →R 1	0.80	1.00	0.24	non-creative→dim
L 1 →L 1	0.80	1.00	0.23	creative→intelligent
R13 →R 9	1.00	1.00	0.23	overall rating low→need a push
R 4 →R13	0.80	1.00	0.22	need supervision→overall rating low
R 5 →R 4	0.80	1.00	0.22	less motivated→need supervision
R 5 →R13	0.80	1.00	0.22	less motivated→overall rating low
R 9 →R10	0.80	1.00	0.22	need a push→non-creative
L 1 →L 3	0.80	1.00	0.21	intelligent→new boy
L 6 →L13	0.80	1.00	0.21	reliable→overall rating high
R10 →R 9	0.80	1.00	0.20	non-creative→need a push
R 1 →R 6	0.60	1.00	0.19	dim→not so reliable
R 1 →R10	0.80	1.00	0.19	dim→non-creative
R 9 →R 4	0.60	1.00	0.19	need a push→need supervision
R 9 →R12	0.60	1.00	0.19	need a push→less professional
R 9 →R13	0.60	1.00	0.19	need a push→overall rating low
R12 →R13	0.80	1.00	0.19	less professional→overall rating low
R13 →R 4	0.80	1.00	0.19	overall rating low→need supervision
R13 →R12	0.80	1.00	0.19	overall rating low→less professional
L 4 →L 5	0.80	1.00	0.18	little supervision reqd→motivated
L 4 →L 9	0.60	1.00	0.18	little supervision reqd→self starters
R 6 →R 4	0.80	1.00	0.18	not so reliable→need supervision
R 6 →R13	0.80	1.00	0.18	not so reliable→overall rating low
L12 →L 9	0.60	1.00	0.18	professional→self starters
L13 →L 4	0.80	1.00	0.18	overall rating high→little supervision reqd
L13 →L 5	0.80	1.00	0.18	overall rating high→motivated
L13 →L 9	0.60	1.00	0.18	overall rating high→self starters
L 1 →L 8	0.80	1.00	0.17	intelligent→ideas men

Fig.4 ENTAIL analysis of repertory grid elicited on staff appraisal

proper inferences that are missing. He can then edit the rules and test the revized system using Nexpert's facilities.

Repertory grid techniques depend on eliciting elements and constructs from experts that are representative of a domain and comprehensive in their classification. The interative elicitation program PEGASUS in PLANET uses online analysis of the grid to feed back comments to the expert which stimulate the addition of elements and constructs to achieve comprehensiveness [1]. However, this structural feedback is only applicable when .a grid has been partially completed and the initial selection of elements has been made. TEXAN is a text analysis program designed to pump-prime the grid elicitation process when a manual or text book is available that the expert regards as having reasonable coverage of the domain.

It uses techniques that were originally designed to map subject matter concepts against student concepts in computer-managed instruction systems [17]. The text is fully indexed by all non-noise words grouped by their stems, and a coupling matrix of word associations is calculated using a simple distance-in-text measure. The high-frequency associations in the text are clustered and presented to the expert as a prototypical schema for the subject area which he can edit for spurious words and associations, and then use to suggest knowledge islands and

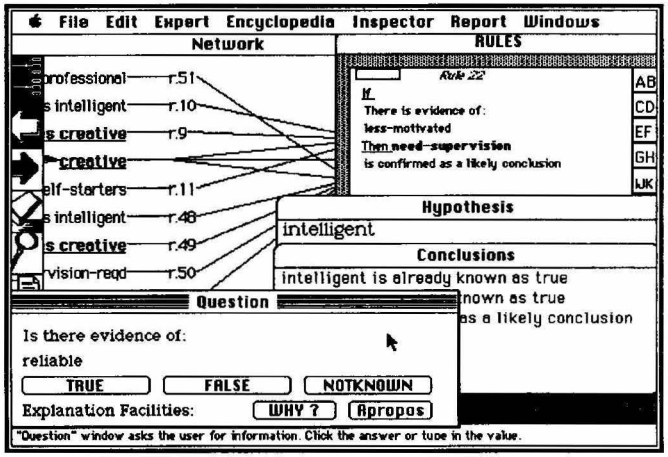


Fig.5 Inference rules derived by ENTAIL in use in Nexpert expert system shell

associated elements and constructs. Figure 6 shows a TEXAN clustering of an evaluation study of data logging, analysis and presentation methodologies for human performance evaluation in complex systems [18].

Michalski and Chilausky [19] have demonstrated that inductive modeling of an expert's behavior may produce effective rules when those elicited by interview techniques are clearly inadequate. ATOM is an algorithm for inducing the structure of a system from its behavior using a search over a model space ordered by complexity and goodness of fit. As in ENTAIL, models are evaluated in terms of the uncertainty reduction induced by the model in the distribution of the modeled behavior [14, 20, 21]. A version of ATOM has been incorporated in KITTEN 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 ES shell to give a simulator of the behavioral system. This has proved effective with inter-personal interaction data such as that analyzed by Mulhall [22] and interactively elicited by Stevens [23].

These groupings, when combined with the construct classifications of repertory grids as in Figure 3 may be viewed as schema structuring a knowledge domain. ANALOG is a program that maps schema to schema based on their structure without regard to content. It is based on a theory of analogy that explicates analogies as pullbacks of faithful functors between categories [26] and generates maximal sub-graph isomorphisms between two classificatory data structures. ANALOG produces meaningful results on artificial examples and grids in related domains. It will also find meaningless analogies between unrelated domains which cannot be rejected by information-theoretic statistical procedures such as those used in ENTAIL and ATOM. It seems likely that effective application of ANALOG depends on the expert pump-priming the matching with known or hypothesized relations and the program extending these rather than attempting to generate them completely.

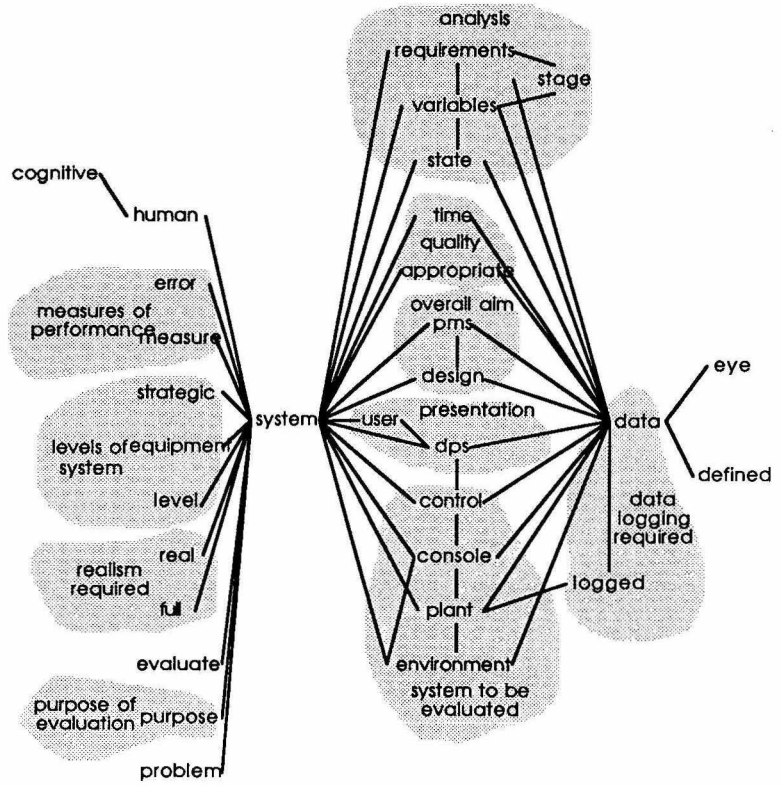


Fig.6 TEXAN clustering of word associations from text with annotation showing knowledge islands

SOCIO: MUTUAL AGREEMENT AND UNDERSTANDING ANALYSIS

We have already emphasized the need for knowledge elicitation methodologies to cope with a group of experts as well as the individual. 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 users are compared. Grids are elicited separately but then exchanged in two ways: a user 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.

The SOCIO program has been used to develop a Participant Construct System, PCS, which supports multiple interacting users in group problem solving activities through terminals on a network of Macintoshes [24]. It follows the paradigm suggested by Chang [25] of a

participant system in which computer-based communication between multiple users is essential to the performance of their tasks, rather than merely incidental to their use of a timeshared computer or computer network. PCS supports its 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;
- exploring a problem-solving domain using the full group resources.

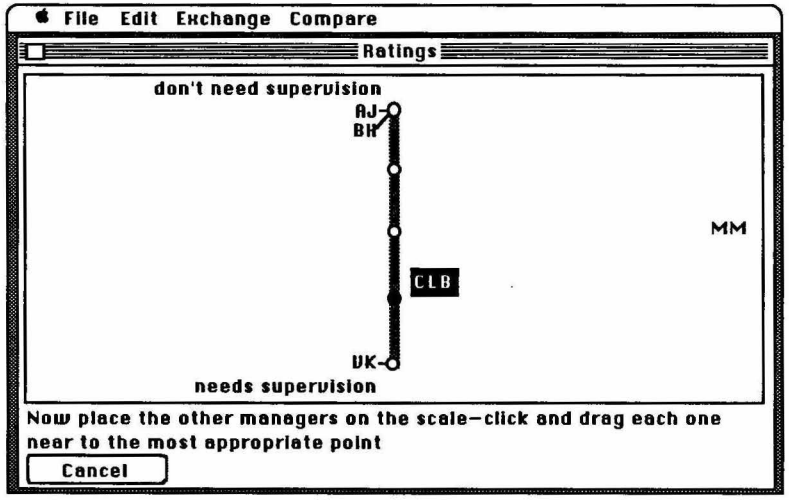


Fig.7 Rating elements on a construct spatially using click and drag

The system uses the Macintosh pull-down menus to add, delete or change elements or constructs. Figure 7 shows a construct being elicited using the natural click-and-drag techniques of the window/mouse interaction rather than numeric rating scale of Figure 3. In the example in Figure 7, the user had been presented with a triad of managers AJ, BH, and VK. He is asked in what way two are similar and in the same way different from the other. This person has said that AJ and BH are alike as they don't need supervision, whereas VK needs supervision. He then places each of the other managers in turn on the scale in an appropriate place. This is his own terminology and his own scale for this construct. If he finds that his original triad AJ, BH and VK need to be moved to accommodate the others he can do that. In this case he puts CLB partly up the scale, and MM partly down the scale, because MM doesn't need much supervision and CLB needs a bit more but not as much as VK He then continues with the process, adding more constructs from triads. As the elicitation proceeds he can bring up the Status window to see who else is participating on the system, and how many elements and constructs each has elicited. He may also select additional elements from those entered by other participants.

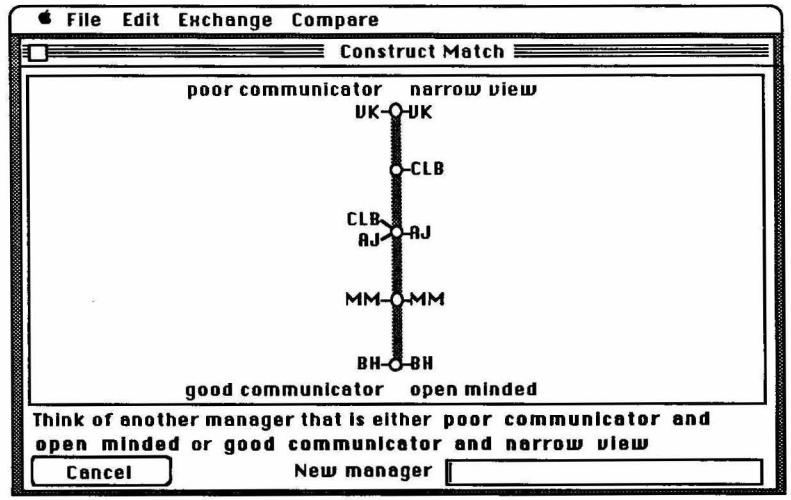


Fig.8 Eliciting an element by breaking a construct match

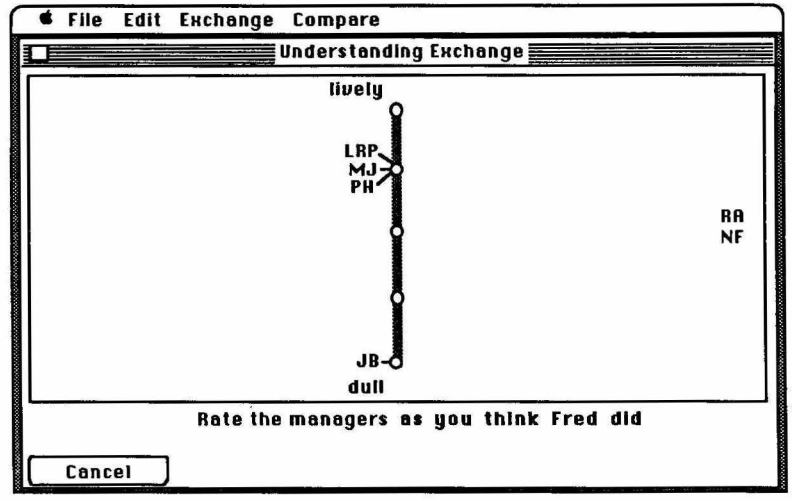


Fig.9 Understanding exchange

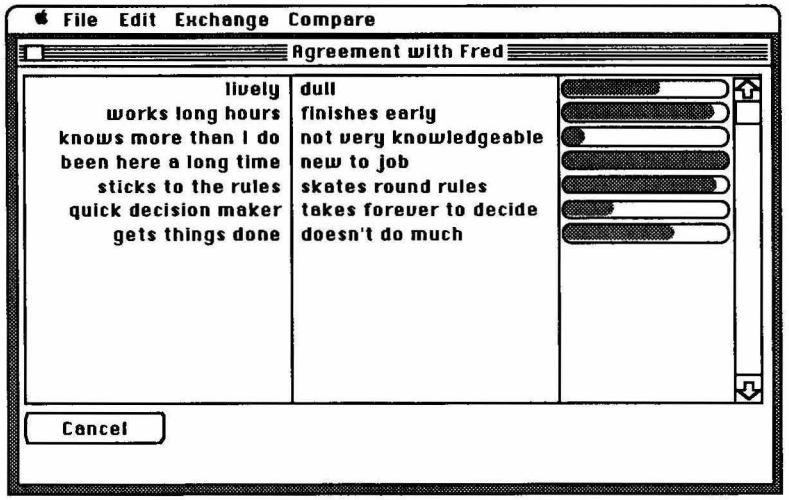


Fig.10 Presentation of degrees of agreement when one participant has rated elements using another's constructs

As he continues to elicit elements and constructs, calculations are made as to the similarity of the placing of each element on each construct. This is shown by 'thermometers' like those in Figure 10. Figure 8 shows the natural representation of the construct match screen supporting element elicitation.

There are also two forms of construct **exchange** available to a participant to enable him to compare his perspective on management effectiveness with those of other participants. He can rate a chosen participant's elements on his constructs either as: he, himself, would have done; or as he thinks the other participant did. These two procedures enable him to compare either his **agreement** with the other participant or his **understanding** of the other participant. He can select either of these forms, or do both one after the other. Figure 9 shows an **understanding exchange** with Fred. He has to rate Fred's elements on Fred's constructs **as he thought Fred had rated them.** The screens look the same for each construct, just the positioning of the elements is different .

Figure 10 shows the presentation of **agreement** when one participant has made his own assignment of elements to constructs elicited from another participant. Such **exchange grid** procedures are known to aid individuals in coming to understand their agreements and disagreements with the way others perceive the world, and to help the to understand and use the constructs of others.[1]

CONCLUSIONS

This paper has presented the concept of knowledge support systems as interactive knowledge engineering tools and outlined the structure and key components of the KITTEN implementation. KITTEN consists of a set of knowledge engineering tools, some of which already have track records of successful use in knowledge acquisition studies. In developing KITTEN the integrity of each of these tools have been preserved, enabling each to be utilized effectively in a stand-alone mode. However, the first steps have also been made towards an integrated knowledge support system by the tools being built around a common database, providing access to the same data in each of its intermediate forms, and providing conversion utilities between different data forms. It incorporates exchange grid procedures of exploring agreement and understanding in a group, and shows the commonality of decision making and problem solving in a specific domain.

ACKNOWLEDGEMENTS

Financial assistance for this work has been made available by the National Sciences and Engineering Research Council of Canada. Most of this work has been carried out jointly with Brian Gaines, University of Calgary. The PCS implementation is being carried out by Alberta Research Council. We are grateful to John Boose and Jeff Bradshaw of Boeing AI Center, and Ernie Chang of Alberta Research Council for stimulating discussions relating to knowledge support systems.

REFERENCES

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

[2] Shaw, M.L.G., Ed. (1981). **Recent Advances in Personal Construct Technology**. London: Academic Press.

[3] Shaw, M.L.G. (1981). Conversational heuristics for eliciting shared understanding. Shaw, M.L.G., Ed. **Recent Advances in Personal Construct Technology**. pp. 31-44. London: Academic Press.

[4] Shaw, M.L.G. (1985). Communities of knowledge. Epting, F. & Landfield, A.W., Eds. **Anticipating Personal Construct Psychology**. pp. 25-35. Lincoln: University of Nebraska Press.

[5] Shaw, M.L.G. & Gaines, B.R. (1981). The personal scientist in the community of science. Reckmeyer, W.J., Ed. **General Systems Research and Design: Precursors and Futures**. pp. 59-68. Society for General Systems Research (January).

[6] Boose, J.H. (1986). Rapid acquisition and combination of knowledge from multiple experts in the same domain. **Future Computing Systems 1**(2), In press.

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

[8] 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.

[9] 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.

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

[11] Shaw, M.L.G. & Gaines, B.R. (1986). An interactive knowledge elicitation technique using personal construct technology. Kidd, A., Ed. **Knowledge Elicitation for Expert Systems: A Practical Handbook**. To appear. Plenum Press.

[12] Gaines, B.R. & Shaw, M.L.G. (1981). A programme for the development of a systems methodology of knowledge and action. Reckmeyer, W.J., Ed. **General Systems Research and Design: Precursors and Futures**. pp. 255-264. Society for General Systems Research (January).

[13] Shaw, M.L.G. & Gaines, B.R. (1986). A framework for knowledge-based systems unifying expert systems and simulation. Luker, P.A. & Adelsberger, H.H., Eds. **Intelligent Simulation Environments**, 38-43 (January). La Jolla, California: Society for Computer Simulation.

[14] Gaines, B.R. (1977). System identification, approximation and complexity. **International Journal of General Systems, 3**, 145-174.
[15] Roy, J. (1986). Expert systems in Nexpert. **MacTutor, 2**(2), 48-51 (February).
[16] Gaines, B.R. & Shaw, M.L.G. (1986). Induction of inference rules for expert systems. **Fuzzy Sets and Systems**, 8(3), 315-328 (April).
[17] Smith, R.A. (1976). Computer-based structural analysis in the development and administration of educational materials. **International Journal of Man-Machine Studies, 8**(4), 439-463 (July).
[18] Gaines, B.R. & Moray, N. (1985). Development of performance measures for computer-based man-machine interfaces. **Technical Report** DCIEM-PER-FIN:JUL85.
[19] Michalski, R.S. & Chilausky, R.L. (1980). Knowledge acquisition by encoding expert rules versus computer induction from examples—A case study involving soyabean pathology. **International Journal of Man-Machine Studies, 12**, 63-87.
[20] Gaines, B.R. (1976). Behaviour/structure transformations under uncertainty. **International Journal of Man-Machine Studies, 8**(3), 337-365 (May).
[21] Gaines, B.R. (1979). Sequential fuzzy system identification. **International Journal of Fuzzy Sets and Systems, 2**(1), 15-24 (January).
[22] Mulhall, D.J. (1977). The representation of personal relationships: an automated system. **International Journal of Man-Machine Studies, 9**(3), 315-335 (May).
[23] Stevens, R.F. (1985). An on-line version of the personal relations index psychological test. **International Journal of Man-Machine Studies, 23**(5), 563-585 (November).
[24] Shaw, M.L.G. (1986a). **Participant Construct System User Manual**. ARC-CPCS-PCS-MANUAL:FEB86. Calgary: Alberta Research Council (February).
[25] Chang, E. (1986). Participant systems. **Future Computing Systems**, in press.
[26] Gaines, B.R. & Shaw, M.L.G. (1982). Analysing analogy. Trappl, R., Ricciardi, L. & Pask, G., Eds. **Progress in Cybernetics and Systems Research**. Vol. IX, pp. 379-386. Washington: Hemisphere.