

TECHNIQUES FOR KNOWLEDGE ACQUISITION AND TRANSFER

Mildred L G Shaw & Brian R Gaines
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 expert systems. It presents the concept of knowledge support systems as interactive knowledge engineering tools, states the design criteria for such systems, and outlines the structure and key components of KSS1 and its KITTEN implementation. KSS1 includes tools for interactive repertory grid elicitation and entailment analysis that have been widely used for rapid prototyping of industrial expert systems. It also includes 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.

KNOWLEDGE SUPPORT SYSTEMS

Problems of knowledge engineering have been recognized since the early days of 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 to 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 decline in costs of both hardware and software support for expert systems has brought the technology into a mass-market situation far more rapidly than originally envisioned;
- This has led to a growth in demand for expert systems that is proceeding far more rapidly than the growth in supply of trained and experienced knowledge engineers;
- The declining costs of expert system technology are also making the expense of human labour in tailoring the technology for particular applications appear to be the dominating constraint and an excessive cost;
- A move towards a labor-intensive activity such as knowledge engineering is contrary to all trends in industry;
- In particular it is contrary to the trend towards automatic programming techniques in the computing industry;
- 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.

The considerations of the previous section have heightened interest in the possibility of providing knowledge support systems (KSSs) to automate knowledge engineering as a process of direct interaction between domain experts and the computer. Gaines & Shaw (1980) proposed that personal construct psychology (Kelly 1955, Shaw 1980) could provide foundations for expert systems, particularly in systems that combined interactivity with database access and expert advice to provide decision support, and gave examples of algorithms and programs that extracted entailment rules from repertory grid data. Shaw & Gaines (1983) reported on further enhancements of these techniques and a preliminary experiment to validate them empirically as a knowledge engineering technique for priming expert systems. This work led to industrial studies of the methodology applied to the development of expert systems: Boeing Computer Services (Boose 1984, 1985, 1986) and Lockheed Software Technology Center (Wahl 1986) have reported success in applications; and validation has been reported in a statistics domain (Gammack & Young 1985).

This paper gives a state-of-the-art report on the use of techniques based on personal construct psychology to automate knowledge engineering for expert systems. It is based on four areas of advance since the previous paper:

- Improved techniques for the derivation of rules from repertory grid data which give: a natural knowledge representation for uncertain data combining fuzzy and probabilistic logics; and an information-theoretic measure of the significance of a derived rule (Gaines & Shaw 1986a);
- Widespread applications experience in prototyping expert systems using the methodology (Boose 1985, Gaines & Shaw 1986b);
- Improved interactive techniques for on-line knowledge engineering from groups of domain experts interacting through a computer network (Shaw 1986, Shaw & Chang 1986);
- KSS1, and its KITTEN implementation, a knowledge engineering workbench that provides next generation KSS facilities including textual analysis, induction of models

from behavior, multi-level and multi-expert repertory grid elicitation, and hierarchical construct laddering, to automate knowledge engineering for a wide range of problem domains.

KNOWLEDGE SUPPORT SYSTEM DESIGN CONSIDERATIONS

Knowledge engineering may be seen 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 even broader scope, encompassing both aids to knowledge engineering and support of human knowledge processes—in the long term the division between knowledge engineering tools and expert system shells will break down, and integrated systems will be necessary. The general requirements for a KSS are:

1. The KSS tools should be domain independent;
2. The KSS tools should be directly applicable by experts without intermediaries;
3. The KSS tools should be able to access a diversity of knowledge sources including text, interviews with experts, and observations of expert behavior;
4. The KSS system should be able to encompass a diversity of perspectives including partial or contradictory input from different experts;
5. The KSS system should be able to encompass a diversity of forms of knowledge and relationships between knowledge;
6. The KSS system should be able to present knowledge from a diversity of sources with clarity as to its derivation, consequences and structural relations;
7. Users of the KSS should be able to apply the knowledge in a variety of familiar domains and freely experiment with its implications;
8. The KSS should make provision for validation studies;
9. As much of the operation of the KSS as possible should be founded on well-developed and explicit theories of knowledge acquisition, elicitation and representation;
10. As the overall KSS develops it should converge to an integrated system.

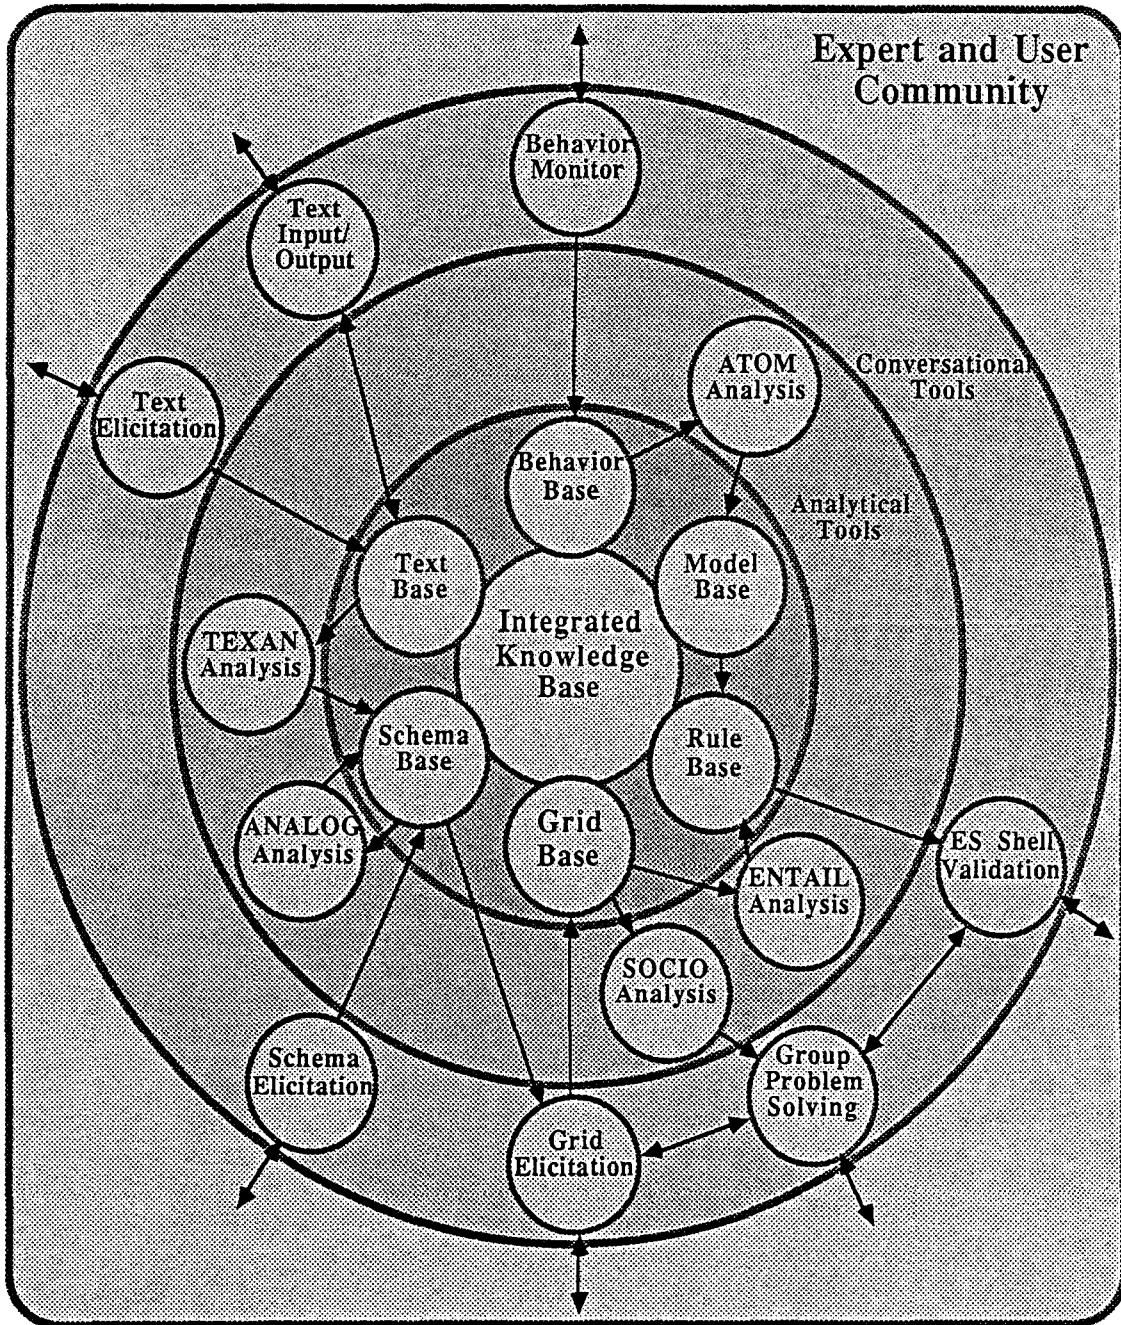
All of these requirements are subject to caveats—some domain dependence may be appropriate for efficiency in specific KSSs—some human intervention may be helpful or necessary when an expert is using a KSS—and so on. However, the broad design goals stated capture the key issues in KSS design currently.

The PLANET system for repertory grid elicitation and analysis (Shaw 1980, 1982, Shaw & Gaines 1986b,c) is a primitive KSS satisfying requirements 1 and 2 for domain independence and direct use. Its foundations in personal construct psychology, which itself has strong systemic and cognitive science foundations (Gaines & Shaw 1981, Shaw & Gaines 1986a), are attractive in terms of requirement 9. Boose (1985) in evaluating ETS has noted the limitations of basic repertory grid techniques in terms of requirement 5—that the methodology is better suited for analysis than for synthesis problems, for example, debugging, diagnosis, interpretation and classification rather than design and planning, and that it is difficult to apply to deep causal knowledge or strategic knowledge—and is attempting to overcome these use grid hierarchies in NeoETS (Bradshaw & Boose 1986). The TEIRESIAS extension to MYCIN is an early form of KSS providing debugging support for an expert system using basic analogical reasoning (Davis & Lenat 1982). The development of KSSs has become a major area of activity recently, for example, MORE (Kahn, Nowlan & McDermott 1985), SALT (Marcus, McDermott & Wang 1985), SEAR (van de Brug, Bachant & McDermott 1985), and MOLE (Eshelman & McDermott 1986).

The following section describes our work on KSS1, a knowledge support system that draws on many concepts and techniques for knowledge engineering to begin to encompass requirements 3 through 8, while attempting to satisfy 9 by relating them all through personal construct psychology, and 10 by building a workbench of tools around a common database.

KSS1: A KNOWLEDGE SUPPORT SYSTEM

Figure 1 shows the structure of KSS1 (in its workstation implementation termed KITTEN: Knowledge Initiation & Transfer Tools for Experts and Novices). KSS1 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 on a coupled IBM AT and Apple Macintosh to combine processing power and interactivity.



**Figure 1 KSS1—KITTEN—
Knowledge Initiation and Transfer Tools for Experts and Novices**

The KSS1 structure is best understood by following sequences of activity that lead to the generation of a rule base and its loading into an ES 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 ES shell (Gaines & Shaw 1986a).

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 (Gaines 1977).

These two routes can be combined. KSS1 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 (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 KSS1 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).

The KITTEN implementation is an initial KSS1 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.

ENTAIL: ENTAILMENT ANALYSIS

PLANET and ETS access the expert's personal construct system by interactively eliciting a repertory grid of constructs classifying elements characterizing to part of the domain of expertise. A repertory grid, such as that of Figure 2, may be viewed as a component of a database in entity-attribute form with elements as entities, constructs as attributes and ratings of elements on poles of constructs as values.

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 (Gaines & Shaw 1986a).

Figure 3 is an ENTAIL II analysis of the grid of Figure 2. 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, $L1 \rightarrow L9$ 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.

The data of Figure 2 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 3 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.

	*	1	2	3	4	5	6	7	8	9	10	*	

intelligent	1	*	1	4	5	3	3	5	2	3	5	*	1 dim
		*										*	
willing	2	*	1	2	4	5	1	1	4	3	1	2	2 unwilling
		*										*	
new boy	3	*	1	2	3	5	4	4	4	1	4	3	3 old sweats
		*										*	
little supervision reqd	4	*	3	1	4	5	2	1	5	2	2	3	4 need supervision
		*										*	
motivated	5	*	1	1	4	5	2	2	5	3	3	2	5 less motivated
		*										*	
reliable	6	*	3	2	2	5	1	1	5	1	2	3	6 not so reliable
		*										*	
mild	7	*	3	4	5	2	2	3	1	5	4	5	7 abrasive
		*										*	
ideas men	8	*	1	1	5	4	2	3	1	3	4	4	8 staid
		*										*	
self starters	9	*	2	1	5	5	1	3	5	3	4	5	9 need a push
		*										*	
creative	10	*	1	1	5	5	2	3	4	3	4	5	10 non-creative
		*										*	
helpful	11	*	4	3	4	2	3	5	1	4	5	5	11 unhelpful
		*										*	
professional	12	*	1	2	3	3	2	1	5	2	4	4	12 less professional
		*										*	
overall rating high	13	*	2	1	3	4	1	2	5	2	3	4	13 overall rating low
		*										*	
messers	14	*	2	2	5	4	3	5	1	5	3	1	14 tidy

		*	*	*	*	*	*	*	*	*	*	*	S10
		*	*	*	*	*	*	*	*	*	*	*	S9
		*	*	*	*	*	*	*	*	*	*	*	S8
		*	*	*	*	*	*	*	*	*	*	*	S7
		*	*	*	*	*	*	*	*	*	*	*	S6
		*	*	*	*	*	*	*	*	*	*	*	S5
		*	*	*	*	*	*	*	*	*	*	*	S4
		*	*	*	*	*	*	*	*	*	*	*	S3
		*	*	*	*	*	*	*	*	*	*	*	S2
		*	*	*	*	*	*	*	*	*	*	*	S1

Figure 2 Repertory grid elicited on staff appraisal

Figure 4 shows Nexpert in operation loaded with the entailments of Figure 3. 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 proper inferences that are missing. He can then edit the rules and test the revized system using Nexpert's facilities.

In some applications the data in the grid about elements may also be relevant and can be loaded in as a set of rules based on the ratings of the elements on the constructs (Boose 1984), for example, that S1 is intelligent. In other applications the relations between the constructs of different experts may be significant and can be derived by combining multiple grids (Bradshaw & Boose 1986).

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
L10 → 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

Figure 3 ENTAIL analysis of repertory grid on staff appraisal

The screenshot displays the Nextpert expert system interface. At the top, a menu bar includes 'File', 'Edit', 'Expert', 'Encyclopedia', 'Inspector', 'Report', and 'Windows'. Below the menu, a 'Network' window shows a list of terms with associated rule numbers (r.51 to r.50) and arrows pointing to a 'RULES' window. The 'RULES' window displays 'Rule 22' with the following logic: 'If There is evidence of: less-motivated Then need-supervision is confirmed as a likely conclusion'. Below the rules, a 'Hypothesis' window shows 'intelligent' and a 'Conclusions' window shows 'intelligent is already known as true'. A 'Question' window is open, asking 'Is there evidence of: reliable' with buttons for 'TRUE', 'FALSE', and 'NOTKNOWN'. Below the question window are 'Explanation Facilities' buttons for 'WHY?' and 'Apropos'. A footer note states: 'Question window asks the user for information. Click the answer or type in the value.'

Figure 4 Inference rules derived by ENTAIL in Nextpert expert system shell

TEXAN: TEXT ANALYSIS

Repertory grid techniques depend on eliciting elements and constructs from experts that are representative of a domain and comprehensive in their classification. The interactive 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 (Shaw 1980). However, this structural feedback is only applicable when a grid has been partially completed and the initial selection of elements has had no computer-based support.

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 (Smith 1976). 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 associated elements and constructs.

Figure 5 shows a TEXAN clustering of an evaluation study of data logging, analysis and presentation methodologies for human performance evaluation in complex systems (Gaines & Moray 1985). Figure 6 shows an independent mapping of the main knowledge islands for an expert system design based on the analyzed report (Gaines 1986). The TEXAN analysis was done some time after the production of Figure 6, and the shading of Figure 5 shows the relationship of some of the groupings in the schema with the knowledge islands. There is not a one-to-one correspondence but this, and similar analyses, show that basic text analysis can focus attention on salient features of the domain and pump-prime the knowledge elicitation process.

In the long term more sophisticated text analysis techniques may be used to derive knowledge from text without human intervention. However, for many domains the knowledge is not yet that explicit and pump-priming of elicitation from experts will remain a significant requirement.

ANALOG: SCHEMA ANALYSIS

The groupings of Figures 5 and 6 when combined with the construct classifications of repertory grids as in Figure 2 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 (Gaines & Shaw 1982) and generates maximal sub-graph isomorphisms between two classificatory data structures. It may be regarded as a generalization of the copy-edit process being used in the encoding of commonsense knowledge in CYC (Lenat, Prakash & Shepherd 1986). 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.

ATOM: BEHAVIOR ANALYSIS

Michalski and Chilausky (1980) 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

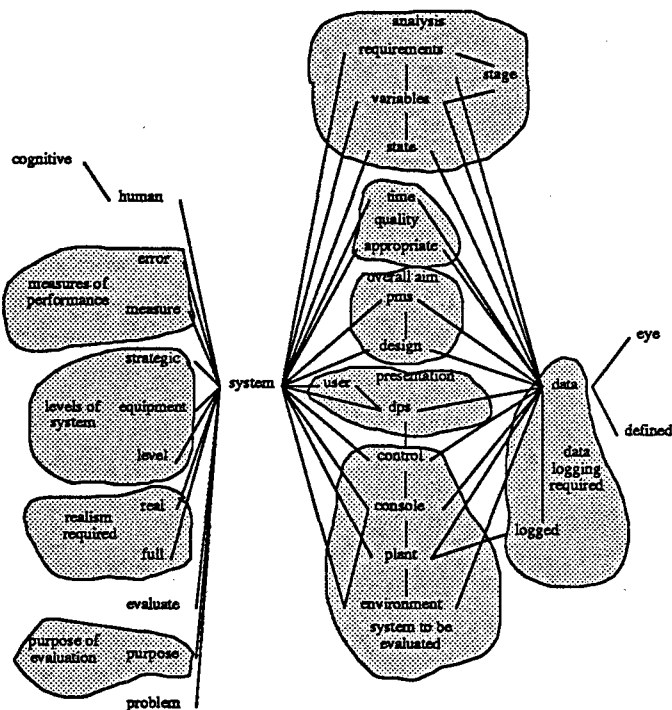


Figure 5 TEXAN clustering of word associations from text with annotation showing knowledge islands

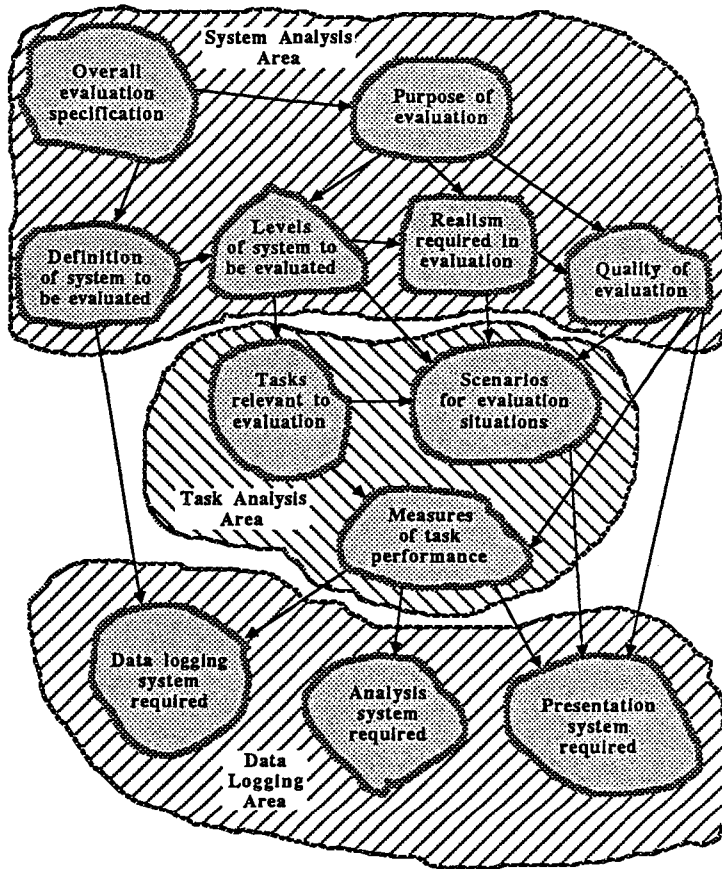


Figure 6 Knowledge islands specified by expert for performance measuring system

in ENTAIL, models are evaluated in terms of the uncertainty reduction induced by the model in the distribution of the modeled behavior (Gaines 1976, 1977, 1979). 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 (1977) and interactively elicited by Stevens (1985).

SOCIO: MUTUAL AGREEMENT AND UNDERSTANDING ANALYSIS

This paper has 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 (Shaw 1986). It follows the paradigm suggested by Chang (1986) 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.

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 2.




Figure 7 Rating elements on a construct spatially using click and drag

Figure 8 shows the natural representation of the construct match screen supporting element elicitation.

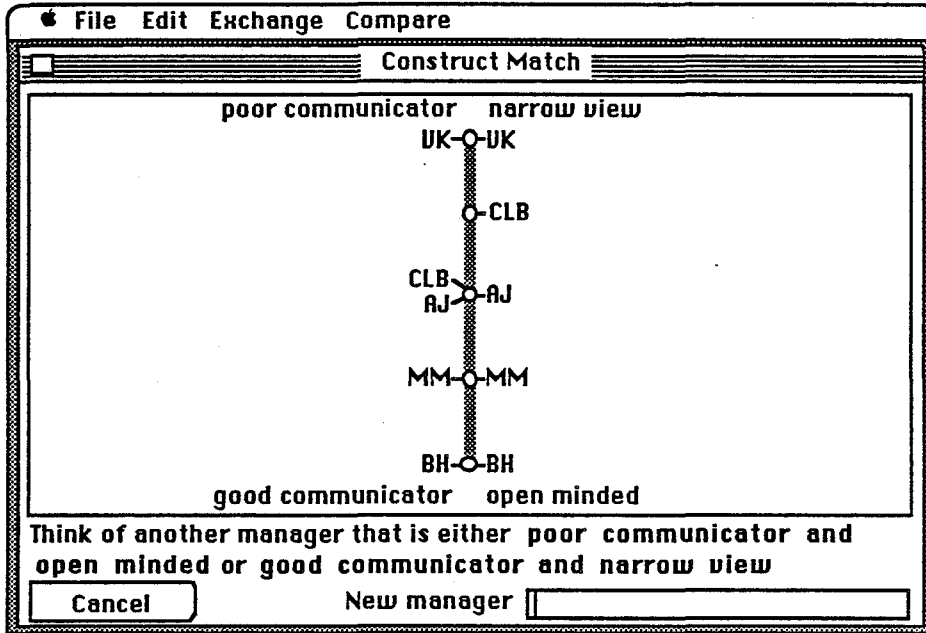


Figure 8 Eliciting an element by breaking a construct match

Figure 9 shows the presentation of agreement when one participant has made his own assignment of elements to constructs elicited from another participant.

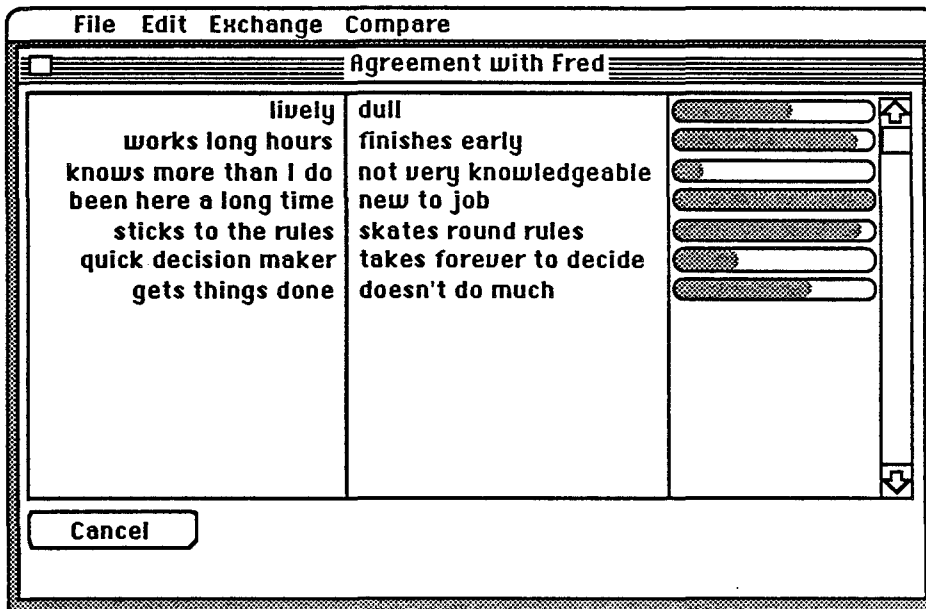


Figure 9 Presentation of degrees of agreement when one participant has rated elements using another's constructs

CONCLUSIONS: STEPS TOWARD AN INTEGRATED KSS

This paper has presented the concept of knowledge support systems as interactive knowledge engineering tools, stated the design criteria for such systems, and outlined the structure and key components of KSS1 and its KITTEN implementation. KSS1 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 has 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 building the tools around a common database, providing access to the same data in each of its intermediate forms, and providing conversion utilities between different data forms.

The objective of integrating the tools has raised a number of new and significant questions. ENTAIL transforms a repertory grid to a set of production rules—is it possible, and useful, to convert production rules to a repertory grid? Technically the result is a possible world of grids that might have generated the rules, and the capability does prove useful, particularly given the other grid analysis tools available in KITTEN. Similar considerations apply to the transformations between other forms of knowledge representation. The next generation of knowledge support tools will be increasingly flexible in handling all aspects of knowledge acquisition, representation, processing and presentation. They will not be optimized with a particular knowledge representation, uncertainty calculation, inference mechanism, and so on, that are in some sense right. Rather they will provide a wide range of perspectives on the knowledge base, preserving source data and chains of derivative processes, so that users can freely explore the knowledge or follow a very specific path according to their choices and needs.

ACKNOWLEDGEMENTS

Financial assistance for this work has been made available by the National Sciences and Engineering Research Council of Canada. 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

- 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. (1985). A knowledge acquisition program for expert systems based on personal construct psychology. *International Journal of Man-Machine Studies* 20(1), 21-43 (January).
- Boose, J.H. (1986). Rapid acquisition and combination of knowledge from multiple experts in the same domain. *Future Computing Systems*, In press.
- Bradshaw, J.M. & Boose, J.H. (1986). NeoETS. *Proceedings of North American Personal Construct Network Second Biennial Conference*. pp. 27-41. University of Calgary: Department of Computer Science (June).
- Chang, E. (1986). Participant systems. *Future Computing Systems*, in press.
- Davis, R. & Lenat, D.B. (1982). *Knowledge-Based Systems in Artificial Intelligence*. New York: McGraw-Hill.
- Eshelman, L. & McDermott, J. (1986). MOLE: a knowledge acquisition tool that uses its head. *Technical Report*. Carnegie-Mellon University: Department of Computer Science.
- Gaines, B.R. (1976). Behaviour/structure transformations under uncertainty. *International Journal of Man-Machine Studies*, 8(3), 337-365 (May).
- Gaines, B.R. (1977). System identification, approximation and complexity. *International Journal of General Systems*, 3, 145-174.

- Gaines, B.R. (1979). Sequential fuzzy system identification. *International Journal of Fuzzy Sets and Systems*, 2(1), 15-24 (January).
- Gaines, B.R. (1986). Development of performance measures for computer-based man-machine interfaces: Application to previous SHINMACS evaluation. **Technical Report DCIEM-PER-SUP:MAR86.**
- Gaines, B.R. & Moray, N. (1985). Development of performance measures for computer-based man-machine interfaces. **Technical Report DCIEM-PER-FIN:JUL85.**
- Gaines, B.R. & Shaw, M.L.G. (1980). New directions in the analysis and interactive elicitation of personal construct systems. *International Journal of Man-Machine Studies* 13(1) 81-116 (July).
- 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).
- 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.
- Gaines, B.R. & Shaw, M.L.G. (1986a). Induction of inference rules for expert systems. *Fuzzy Sets and Systems*, 8(3), 315-328 (April).
- Gaines, B.R. & Shaw, M.L.G. (1986b). Knowledge engineering for an FMS advisory system. Lenz, J.E. (Ed.) **Proceedings of the Second International Conference on Simulation in Manufacturing: AMS'86.** 51-61. Bedford, UK: IFS Conferences (June).
- Gammack, J.G. & Young, R.M. (1985). Psychological techniques for eliciting expert knowledge. Bramer, M., Ed. **Research and Development in Expert Systems,** pp. 105-116. Cambridge University Press.
- Kahn, G., Nowlan, S. & McDermott, J. (1985). MORE: an intelligent knowledge acquisition tool. **Proceedings of the Ninth International Joint Conference on Artificial Intelligence.** pp. 581-584. California: Morgan Kaufmann.
- Kelly, G.A. (1955). **The Psychology of Personal Constructs.** New York: Norton.
- Lenat, D., Prakash, M. & Shepherd, M. (1986). CYC: Using common sense knowledge to overcome brittleness and knowledge acquisition bottlenecks. *AI Magazine* 6(4), 65-85.
- Marcus, S., McDermott, J. & Wang, T. (1985). Knowledge acquisition for constructive systems. **Proceedings of the Ninth International Joint Conference on Artificial Intelligence.** pp. 637-639. California: Morgan Kaufmann.
- 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.
- Mulhall, D.J. (1977). The representation of personal relationships: an automated system. *International Journal of Man-Machine Studies*, 9(3), 315-335 (May).
- 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). Conversational heuristics for eliciting shared understanding. Shaw, M.L.G., Ed. **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. (1986). PCS: a knowledge-based interactive system for group problem solving. **Proceedings of 1986 International Conference on Systems, Man and Cybernetics,** to appear (October).

- Shaw, M.L.G. & Chang, E. (1986). A participant construct system. **Proceedings of North American Personal Construct Network Second Biennial Conference**. pp. 131-140. University of Calgary: Department of Computer Science (June).
- 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). 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.
- Shaw, M.L.G. & Gaines, B.R. (1986b). Interactive elicitation of knowledge from experts. **Future Computing Systems**, 1(2), to appear.
- Shaw, M.L.G. & Gaines, B.R. (1986c). An interactive knowledge elicitation technique using personal construct technology. Kidd, A., Ed. **Knowledge Elicitation for Expert Systems: A Practical Handbook**. To appear. Plenum Press.
- 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).
- 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).
- van de Brug, A, Bachant, J. & McDermott, J. (1985). Doing R1 with style. **Proceedings of the Second Conference on Artificial Intelligence Applications**. IEEE 85CH2215-2, pp. 244-249. Washington: IEEE Computer Society Press.
- Wahl, D. (1986). 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). La Jolla, California: Society for Computer Simulation.