Knowledge Engineering Techniques

B.R. GAINES
University of Calgary
Calgary, Alberta

DR. MILDRED L.G. SHAW
University of Calgary

The system analysis required in the development of expert systems is called knowledge engineering. Presented are the concepts of knowledge support systems as interactive knowledge engineering tools, the design criteria for such systems, and the outlines, structure and key components of KSS1 and its KITTEN implementation. KSS1 includes tools for interactive repertory grid elicitation and entailment analysis that are widely used for rapid prototyping of industrial expert systems. Also included are tools for text analysis, behavioral analysis and schema analysis, that offer complementary and alternative approaches toknowledge acquisition.

INTRODUCTION: KNOWLEDGE ENGINEERING

The initial success of expert system (ES) developments (Michie 1979, Gevarter 1983, Reitman 1984) and the development of a number of reasonably domain-independent software support systems for the encoding and application of knowledge (Hayes-Roth, Waterman & Lenat 1983) has opened up the possibility of widespread usage of ESs. In particular, the Japanese Fifth Generation Computing System (FGCS) development program (Moto-oka 1982, Gaines 1984) assumes this will happen and is targeted on knowledge processing rather than information processing. However, what Feigenbaum (1980) terms knowledge engineering, the reduction of a large body of knowledge to a precise set of facts and rules, has already become a major bottleneck impeding the application of ESs in new domains. We need to understand more about the nature of expertise in itself (Hawkins 1983) and to able to apply this knowledge to the elicitation of expertise in specific domains.

The problems of knowledge engineering have been stated clearly:

> "Knowledge acquisition is a bottleneck in the construction of expert systems. The knowledge engineer's job is to act as a go-between to help an expert build a system. Since the knowledge engineer has far less knowledge of the domain than the expert, however, communication problems impede the process of transferring expertise into a program. The vocabulary initially used by the expert to talk about the domain with a novice is often inadequate for problem-solving; thus the knowledge engineer and expert must work together to extend and refine it. One of the most difficult aspects of the knowledge engineer's task is helping the expert to structure the domain knowledge, to identify and formalize the domain concepts." (Hayes-Roth, Waterman & Lenat 1983)

This bottleneck is now the major impediment to the application of expert systems, and to the realization of their full industrial potential. The automation of knowledge acquisition and transfer has become a major research and technological goal. This paper presents some techniques for the automation of the knowledge engineering task with examples of their application and indications of their further development and integration.

AUTOMATING KNOWLEDGE ENGINEERING

The problem of knowledge elicitation from a skilled person is well-known in the literature of psychology. Bainbridge (1979) has reviewed the difficulties of verbal debriefing and notes that there is no necessary correlation between verbal reports and mental behavior, and that many psychologists feel strongly that verbal data are useless. However, this remark must be taken in the context of experimental psychologists working within a positivist, behavioral paradigm. Other schools of psychology see the problem not as methodological but as psychological and resulting from cognitive defences that impede internal communication for a variety of reasons (Freud 1914, Rogers 1967). Clinical psychologists in particular have developed techniques for making use of verbal interaction to identify underlying cognitive processes, for example through structured interviewing techniques. These can be used to by-pass cognitive defences, including those resulting from automization of skilled behavior. Welbank (1983) has reviewed many of these techniques in the context of knowledge engineering.

In an industrial context the application of techniques derived from clinical psychology to problems of hard system engineering is not an obvious or probable step to take, and the early approach to knowledge engineering was to "manage" it by inventing the profession of knowledge engineer as defined in the quotation above. It has been 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 be 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 lead 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. In 1980 we 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 (Gaines & Shaw 1980). In 1983 we reported further enhancements of these techniques and a preliminary experiment to validate them empirically as a knowledge engineering technique for priming expert systems (Shaw & Gaines 1983a). This work led to industrial studies of the methodology applied to the development of expert systems: to an advisory system for FMS (Gaines & Shaw 1986b); to a variety of industrial applications at Boeing Computer Services (Boose 1984, 1985, 1986); to studies at Lockheed Software Technology Center (Wahl 1986); to 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.

KNOWLEDGE SUPPORT SYSTEMS DESIGN CONSIDERATIONS

We see knowledge engineering 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.

We see knowledge support systems as having 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 1981a, 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.

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

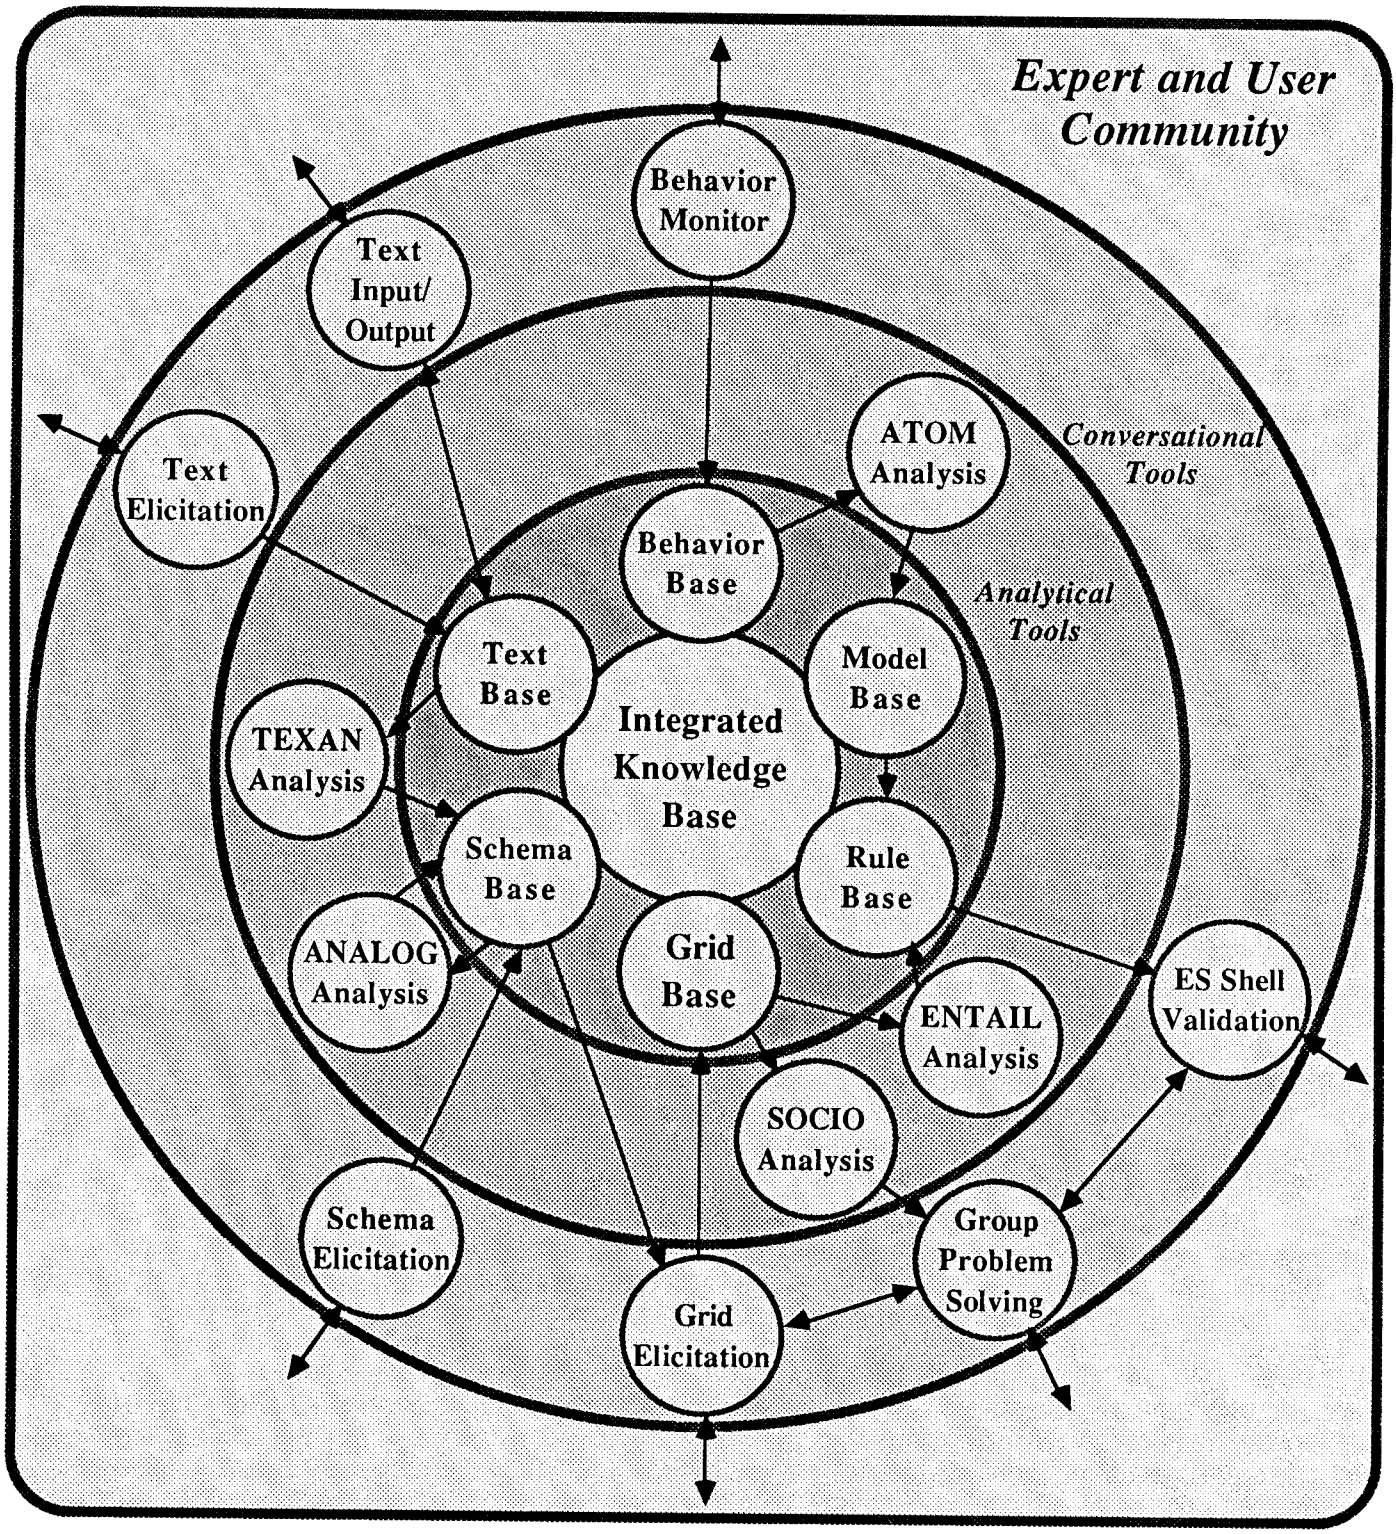

Fig.1 KSS1—KITTEN—
Knowledge Inititiation & Transfer Tools for Experts & Novices

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. Each of the tools included 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 in the KITTEN implementation of KSS1.

PEGASUS AND ETS: REPERTORY GRID ELICITATION

The elicitation of knowledge from experts through general, domain-independent techniques depends on the availability of a framework within which to represent knowledge. This framework needs to be well-founded in terms of both cognitive psychology and computer science. The *repertory grid* is a form of knowledge representation based on a systemic theory of human cognition founded on the single primitive of a *construct*, or dichotomous distinction (Gaines & Shaw 1981a, 1984, 1985). 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)

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)

The personal construct psychology 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.

Knowledge engineering may be seen as eliciting the constructs of an expert in the relevant subject area. The most common method used for eliciting a construct is the minimal context form or triad method. The elements are presented in groups of three, three being the least number which will produce both a similarity and a difference, and the subject is asked to say in what way two are alike and thereby different from the third. This is the emergent pole of the construct. The implicit pole may be elicited by the difference method (in what way does the singleton differ from the pair) or by the opposite method (what would be the opposite of the description of the pair).

As an example, thinking of three factors affecting factory layout, location of services, aisle space and area of department in what way are two alike and thereby different from the other one? We might first of all say that location of services, and aisle space are alike since they *affect flow*, whereas area of department does not. This is, then, the first construct with its two poles or opposite descriptions. Now all the elements in the set must be rated on this dimension as either being *affects flow*, or being *doesn't affect flow* as shown in Figure 2. This also shows the significance of the term personal in personal construct since it would not necessarily be a publicly agreed description that safety doesn't affect flow.

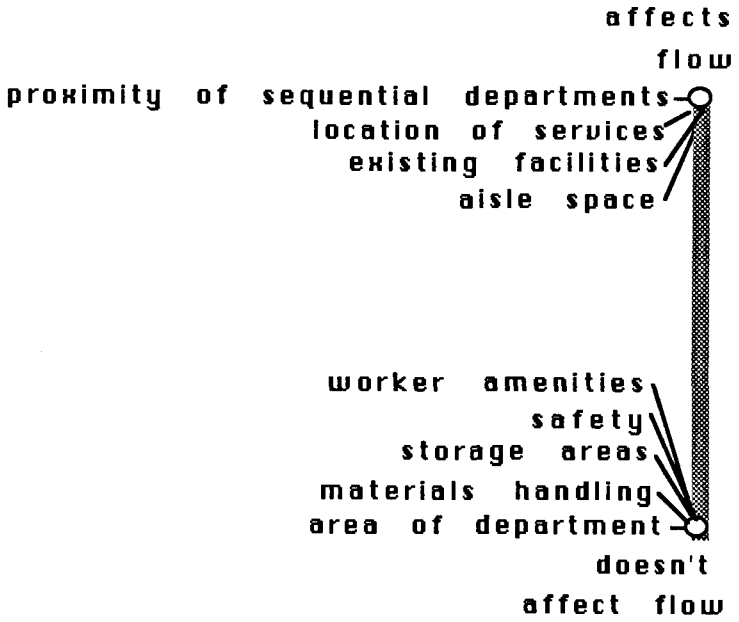

Fig. 2 A construct elicited about elements associated with factory layout

Then the second and subsequent constructs are elicited in exactly the same way using different triads each time. Figure 3 shows the grid of Figure 1 with a further construct elicited. The second construct shown here illustrates that constructs can be factual, imaginary, emotional, or whatever is important to the person generating the grid.

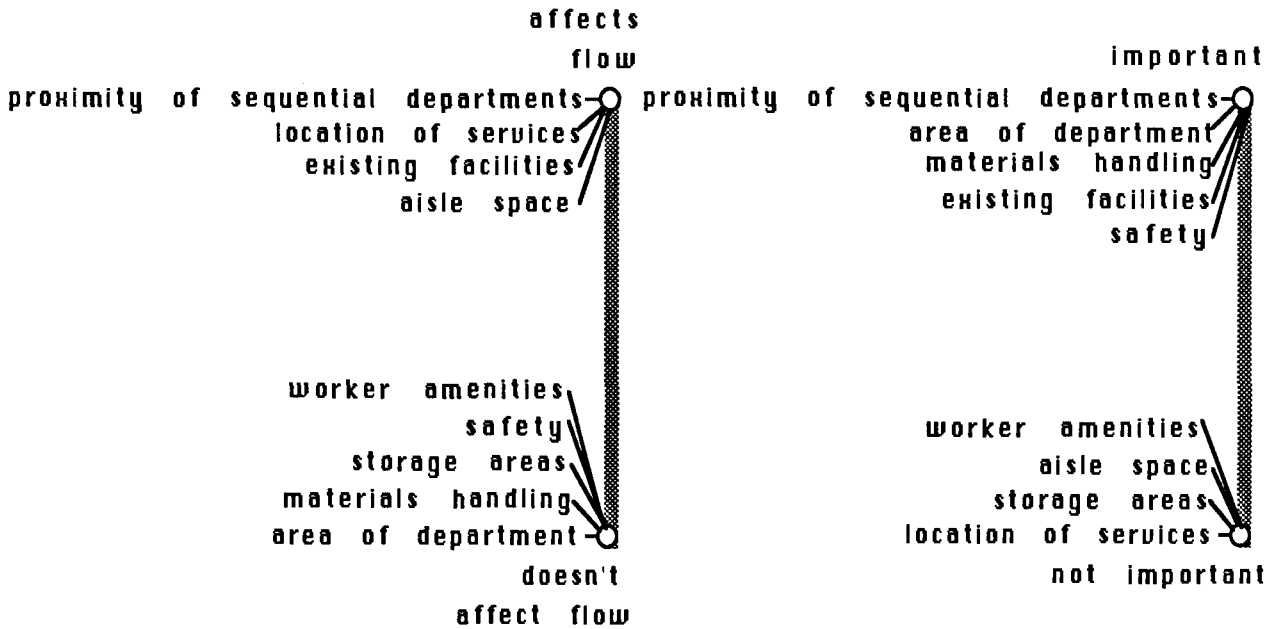

Fig. 3 Two constructs elicited about elements associated with factory layout

A scale allowing more distinctions than the pair 1 and 2 may be used as required. If a 1 to 5 scale is used then the above example might become the grid shown in Figure 4. Note how the use of a multi-point scale with an odd number of values allows for a central rating, in this case 3, which does not force the user to choose either pole. It may be desirable to apply more discrimination to this central rating and allow the subject the choice of the two possibilities: neither, that the element belongs to neither pole; or both, that the element belongs to both poles (Landfield 1976). It is also possible to extend these possibilities to allow separate rating on each pole (Shaw & Gaines 1980, Gaines & Shaw 1981b).

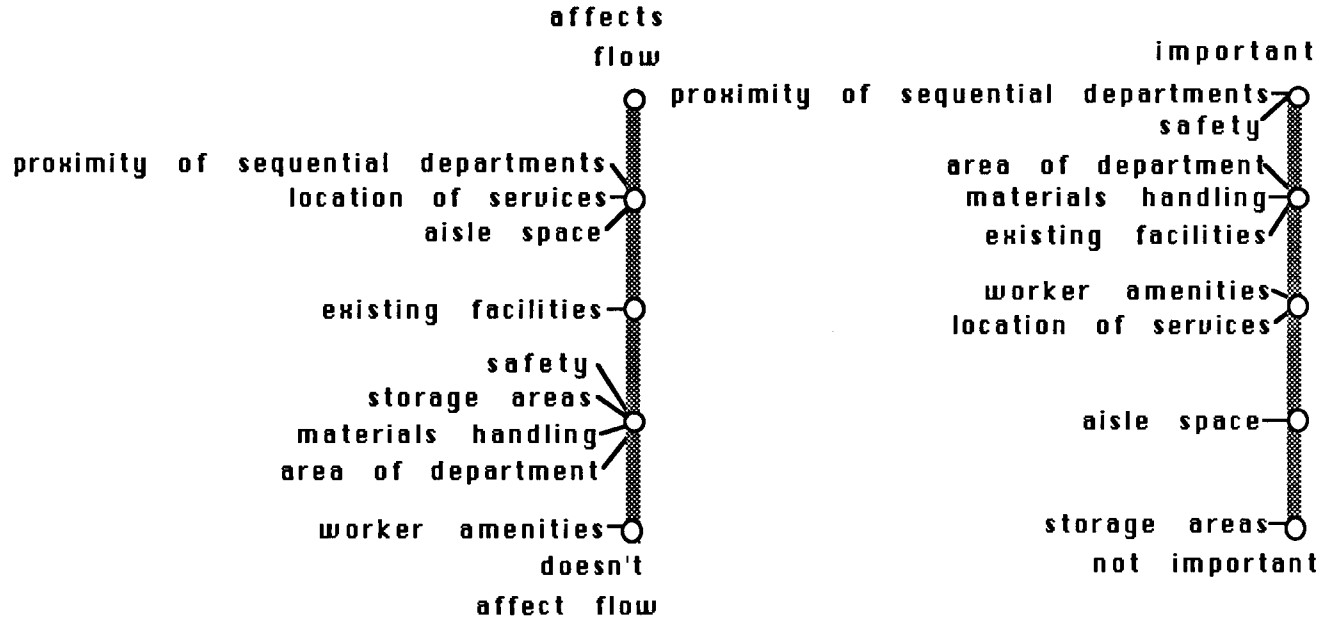

Fig. 4 Scalar constructs elicited about elements associated with factory layout

Figure 5 shows a fully developed grid on factory layout planning (Adler & Matheson 1986) represented in numeric form for use as a data structure for computer processing. A repertory grid 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.

```
                                 * 1  2  3  4  5  6  7  8  9 *
                                 ******************************
                affects flow  1 * 2  4  4  2  4  3  2  4  5 *  1 doesn't affect flow
                                 *                            *
affects department location  2 * 1  3  4  3  5  2  4  3  4 *  2 doesn't affect department location
                                 *                            *
                 inflexible  3 * 2  2  4  2  3  1  4  2  4 *  3 flexible
                                 *                            *
            space dependent  4 * 4  2  4  5  3  1  2  3  3 *  4 space independent
                                 *                            *
         internal interface  5 * 1  3  2  5  2  5  1  3  3 *  5 external interface
                                 *                            *
                constraining  6 * 1  3  4  2  4  1  3  2  4 *  6 not constraining
                                 *                            *
                 inexpensive  7 * 3  3  5  1  3  1  2  5  2 *  7 can be expensive
                                 *                            *
                   important  8 * 1  2  2  3  5  2  4  1  3 *  8 not important
                                 *                            *
       human not considered  9 * 5  5  1  5  2  3  5  3  5 *  9 human considered
                                 *                            *
           can be automated 10 * 1  1  5  5  2  3  4  3  5 * 10 can't be automated
                                 *                            *
doesn't improve environment 11 * 2  3  2  2  4  1  4  5  5 * 11 improves environment
                                 *                            *
        workload unaffected 12 * 4  1  5  1  3  2  1  2  3 * 12 affects workload
                                 ******************************
                                   *  *  *  *  *  *  *  *  worker amenities
                                   *  *  *  *  *  *  *  safety
                                   *  *  *  *  *  *  aisle space
                                   *  *  *  *  *  existing facilities
                                   *  *  *  *  storage areas
                                   *  *  *  location of services
                                   *  *  materials handling
                                   *  area of department
                                   proximity of sequential departments
```

Fig. 5 Repertory grid on factory layout

In the past the elicitation of repertory grids has been a task requiring a skilled psychologist who can draw out a person's constructs whilst not imposing his own ideas or personality. This has involved the development of a variety of techniques and methodologies for grid elicitation (Kelly 1955, Fransella & Bannister 1977). The coming of powerful, low-cost personal computers has made it possible to program some of the elicitation techniques used by people in a form where they can be applied automatically. It is now possible to codify some of the techniques used by people in a form where they can be administered through a conversational interactive computer program (Shaw 1980, Shaw & Gaines 1981a,b,c, 1983b, 1984).

PLANET (Shaw 1982, Shaw & Gaines 1986) and ETS (Boose 1985) 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. They both use interactive prompt-response sequences to elicit the expert's ratings of elements on constructs. In KSS1 construct are elicited using the natural click-and-drag techniques of the window/mouse interaction rather than numeric rating scales as shown in Fig. 6.

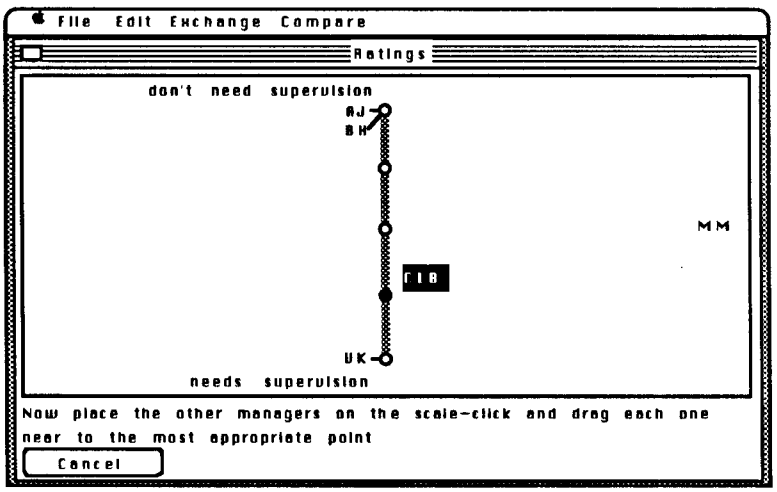

Fig. 6 Construct elicitation in KSS1 using click and drag for rating

An important feature of interactive grid elicitation is that on-line analysis may be used to give feedback to the expert about his construct system, for example that two elements have not been adequately discriminated, or that two constructs are giving very similar discriminations. This feedback may be used to elicit additional constructs and elements. Figure 7 shows the natural representation of the construct match screen supporting element elicitation in KSS1.

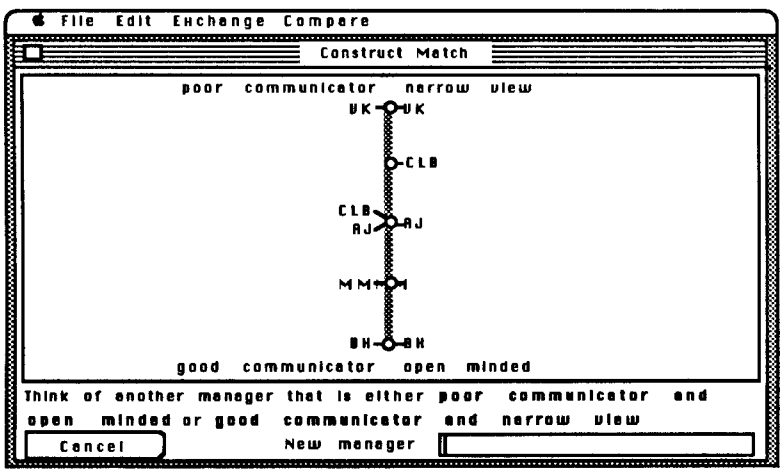

Fig. 7 Element elicitation in KSS1 through feedback of construct matches

ENTAIL: ENTAILMENT ANALYSIS

Once knowledge is represented in a repertory grid it may be analyzed to generate rules suitable for priming an expert system shell. 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 (Gaines & Shaw 1980). 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).

		1	2	3	4	5	6	7	8	9	10	
intelligent	1	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 man	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

Fig.8 Repertory grid elicited on staff appraisal

Figure 9 is an ENTAIL II analysis of the grid of Figure 8 elicited in a study of personnel selection. 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 implies 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 8 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 9 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.

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
R 2 → R 4	0.80	1.00	0.17	unwilling→need supervision
R 2 → R 5	0.80	1.00	0.17	unwilling→less motivated
R 2 → R13	0.80	1.00	0.17	unwilling→overall rating low
R 3 → R 1	0.80	1.00	0.17	old sweats→dim
L 4 → L13	0.80	1.00	0.17	little supervision reqd→overall rating high
L12 → L13	0.80	1.00	0.17	professional→overall rating high
L13 → L12	0.80	1.00	0.17	overall rating high→professional

Fig.9 ENTAIL analysis of repertory grid on staff appraisal

Figure 10 shows Nexpert in operation loaded with the entailments of Figure 9. 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.

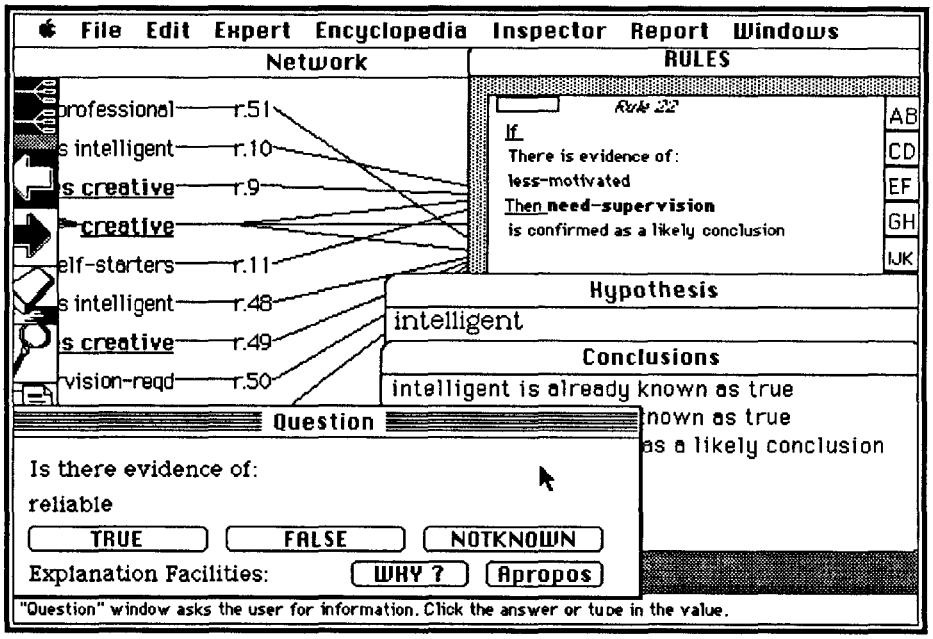

Fig.10 Inference rule, derived by ENTAIL from repertory grid on staff appraisal, in use in Nexpert expert system shell

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 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 (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 11 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 12 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 12, and the shading of Figure 11 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.

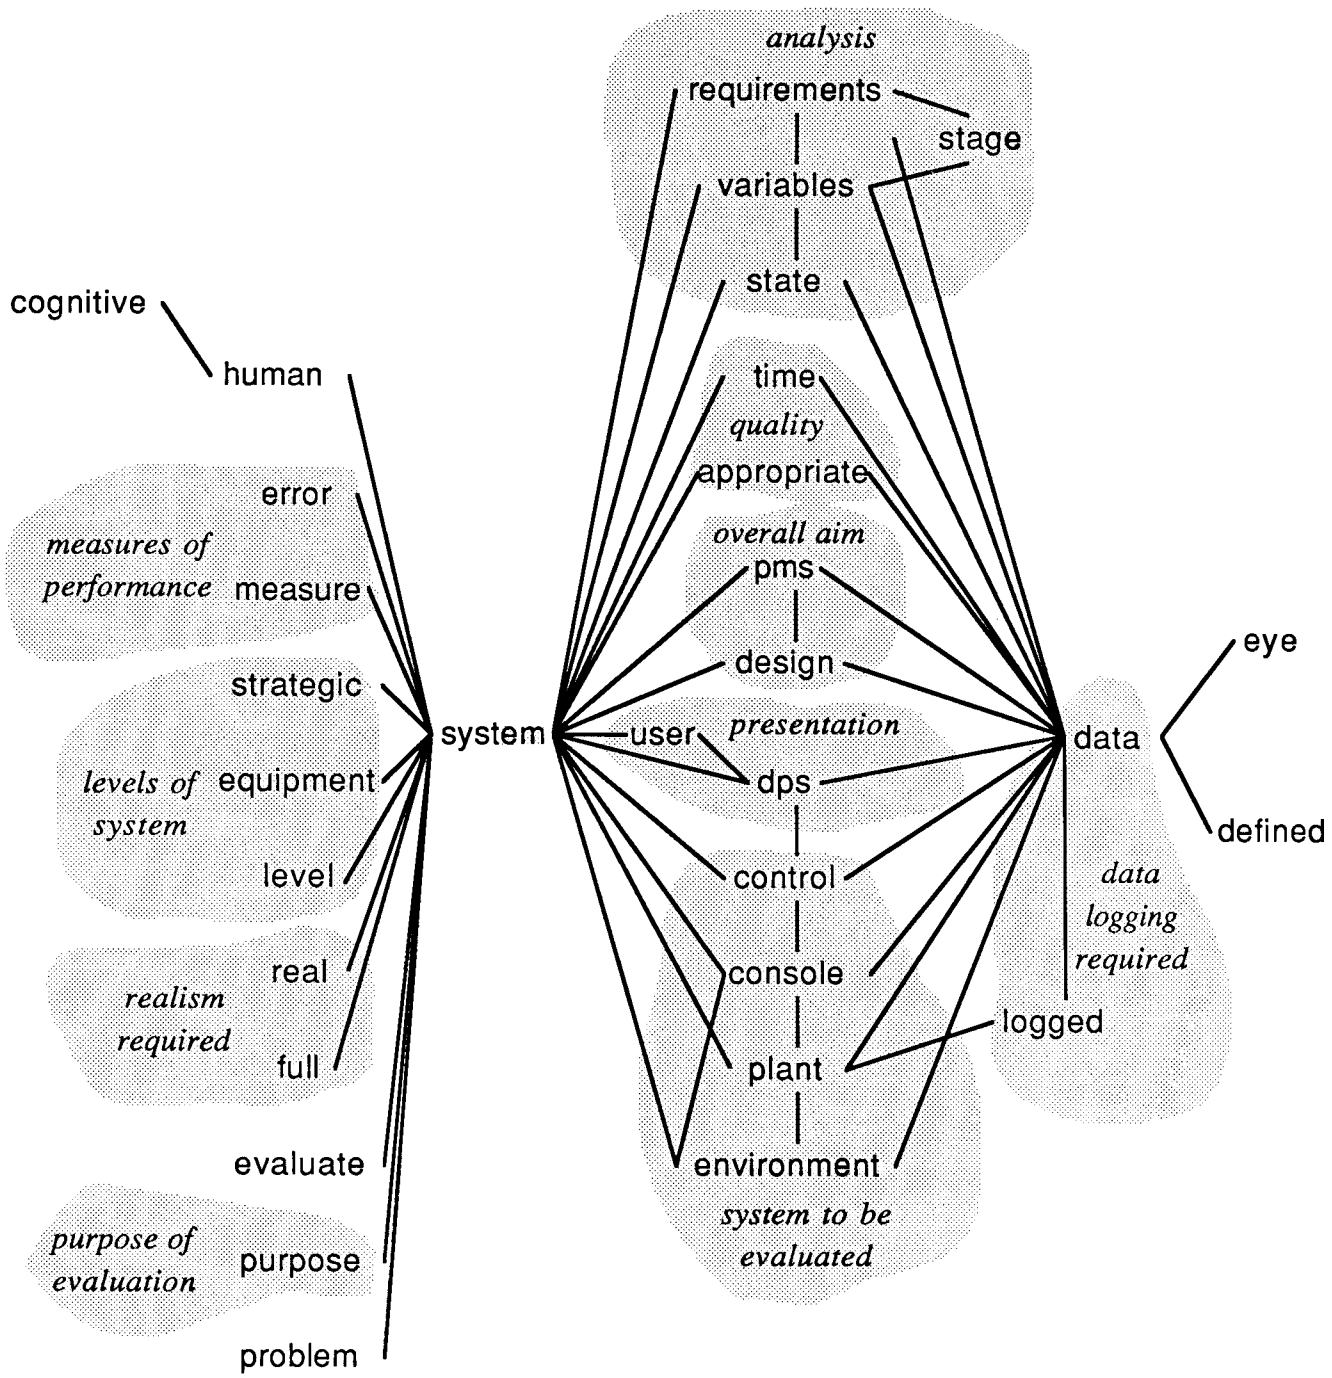

Fig.11 TEXAN clustering of word associations from text with annotation showing knowledge islands

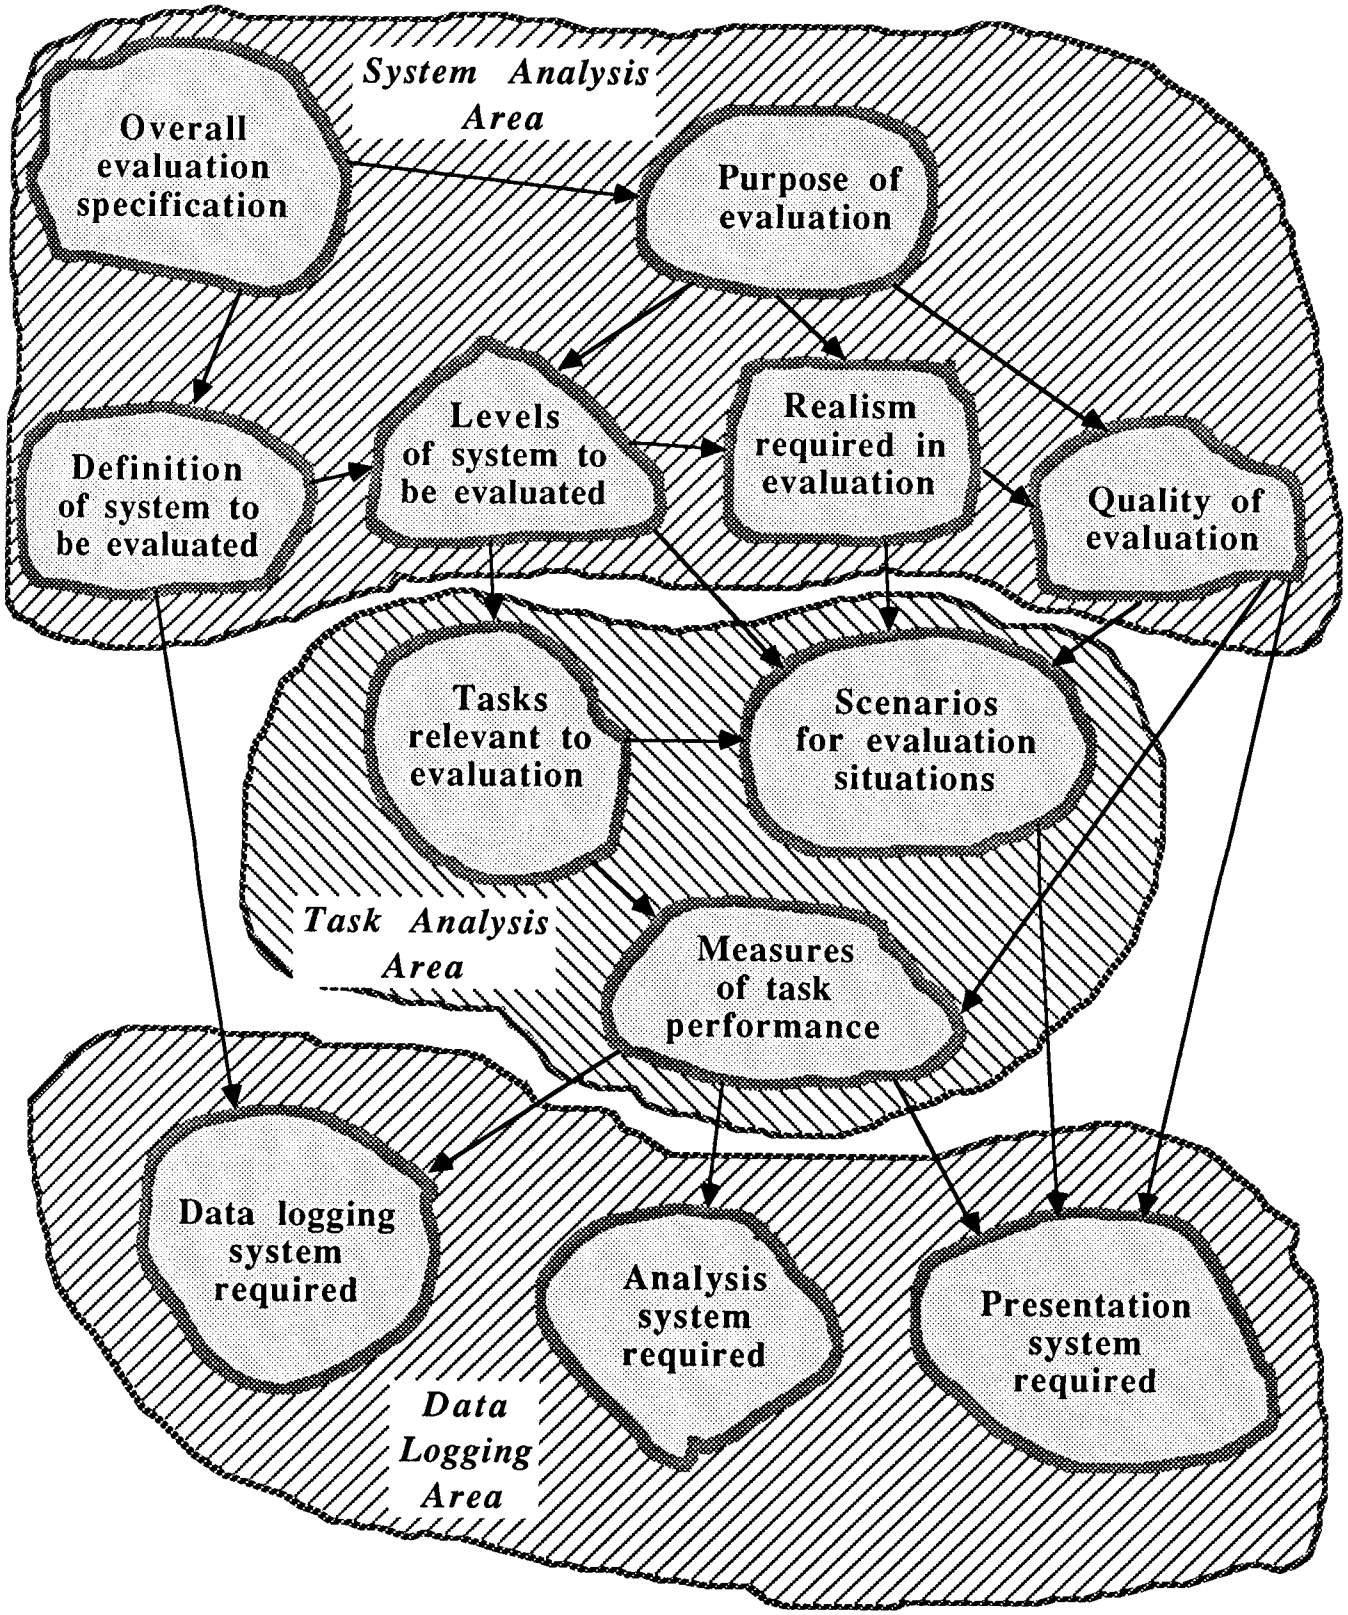

Fig.12 Knowledge islands specified by expert for performance measuring system

ANALOG: SCHEMA ANALYSIS

The groupings of Figures 11 and 12 when combined with the construct classifications of repertory grids as in Figure 8 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 (Gaines & Shaw 1982). 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 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). We have incorporated a version of ATOM 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

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. Figure 13 shows the presentation of agreement when one participant has made his own assignment of elements to constructs elicited from another participant.

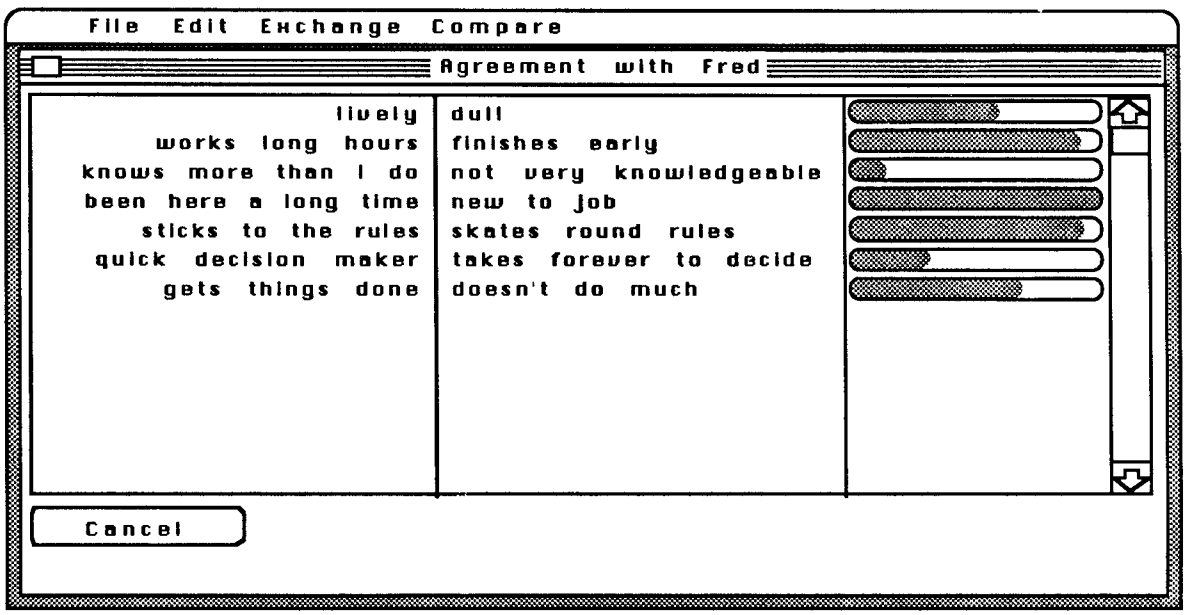

Fig.13 Presentation of degrees of agreement when one participant has rated elements using another's constructs

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, Shaw & Chang 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.

PCS may be seen as the first step towards knowledge support systems that operate interactively with a group of inter-related experts.

STEPS TOWARD AN INTEGRATED KNOWLEDGE SUPPORT SYSTEM

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 we have preserved the integrity of each of these tools, enabling each to be utilized effectively in a stand-alone mode. However, we have also made the first steps 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. We see the next generation of knowledge support tools as being 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.

CONCLUSIONS

Expert system technology is now widely recognized as having a wide range of significant industrial applications. However, the emphasis in the literature is on the software and hardware technologies of expert system shells and Lisp machines. Recognition that the primary proprietary resource and product that is involved is *knowledge* itself is becoming widespread among those seriously involved in expert system projects. With this comes recognition also that the main bottleneck in development is the *knowledge engineering* involved in acquiring, structuring, encoding and validating that knowledge. This has led to emphasis on knowledge engineering as a new profession involved in managing these knowledge processes.

This paper has argued that while knowledge engineering is certainly the area to emphasize in the commercial exploitation of expert systems technology, it is misleading to assume that it necessarily involves the detailed involvement of knowledge engineers. Knowledge engineering can, and should, be automated through knowledge support systems that are used directly be the subject domain experts. The knowledge engineer should operate only at the meta-level in helping the expert to select and apply the tools within these knowledge support systems. This has a number of significant advantages:

— efficient transfer of expertise without the overhead of intermediaries
— accurate transfer of expertise without distortion by intermediaries
— increased awareness by experts of the nature of their knowledge and expertise.

A number of practical knowledge engineering tools have been described which illustrate various possibilities in automating knowledge engineering. Some of these tools have already been used widely for rapid prototyping of industrial expert systems, and others are experimental illustrating the frontiers of

current research. The integration of such tools into versatile knowledge support systems is a major research thrust at a number of organizations concerned with expert systems development. Any project involved in expert systems applications should monitor the development and availability of such tools, and regard them as a basic technology of at least equal importance to those of expert system shells.

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.

BIBLIOGRAPHY

Adler, J. & Matheson, J. (1986). Expert knowledge elicitation using personal construct theory. Thesis, Department of Industrial Engineering, University of Toronto (April).

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

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.

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

Fransella, F. & Bannister, D. (1977). **A Manual for Repertory Grid Technique**. London: Academic Press.

Freud, S. (1914). **Psychopathology of Everyday Life**. London: Benn.

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. (1984). Perspectives on fifth generation computing. **Oxford Surveys in Information Technology, 1**, 1-53.

Gevarter, W.B. (1983). Expert systems: limited but powerful. **IEEE Spectrum, 18**, 39-45.

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. (1981a). 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. (1981b). 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. (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. (1984). Hierarchies of distinctions as generators of system theories. Smith, A.W., Ed. **Proceedings of the Society for General Systems Research International Conference**. pp. 559-566. California: Intersystems.

Gaines, B.R. & Shaw, M.L.G. (1985). Three world views and system philosophies. Banathy, B.H., Ed. **Proceedings of the Society for General Systems Research International Conference**. pp. 387-396. California: Intersystems.

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. J.E. (Ed.). **Proceedings of Second International Conference on Simulation in Manufacturing**, 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.

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

Hayes-Roth, F., Waterman, D.A. & Lenat, D.B., Eds. (1983). **Building Expert Systems**. Reading, Massachusetts: Addison-Wesley.

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.

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.

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.

Mancuso, J.R. & Adams-Webber, J.R., Eds. (1982). **The Construing Person**. New York: Praeger.

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.

Michie, D., Ed. (1979). **Expert Systems in the Micro Electronic Age**. Edinburgh: Edinburgh University Press.

Moto-oka, T., Ed. (1982). **Fifth Generation Computer Systems**. Amsterdam: North-Holland.

Mulhall, D.J. (1977). The representation of personal relationships: an automated system. **International Journal of Man-Machine Studies, 9**(3), 315-335 (May).

Reitman, W., Ed. (1984). **Artificial Intelligence Applications for Business**. Norwood, New Jersey: Ablex.

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

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. (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. (1981a). 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).

Shaw, M.L.G. & Gaines, B.R. (1981b). The personal computer and the personal scientist. **BCS'81: Information Technology for the Eighties**. pp. 235-252. London: Heyden Press for British Computer Society (July).

Shaw, M.L.G. & Gaines, B.R. (1981c). Exploring personal semantic space. Rieger, B., Ed. **Empirical Semantics II**. pp. 712-791. Bochum, West Germany: Studienverlag Brockmeyer.

Shaw, M.L.G. & Gaines, B.R. (1983a). 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. (1983b). Eliciting the real problem. Wedde, H., Ed. **International Working Conference on Model Realism,**. pp. 100-111. Oxford: Pergamon Press.

Shaw, M.L.G. & Gaines, B.R. (1984). Deriving the Constructs Underlying Decision. Zimmerman, H.-J., Gaines, B.R. & Zadeh, L.A., Eds. **Fuzzy Sets and Decision Analysis. TIMS Series Studies in the Management Sciences**. Amsterdam: North-Holland.

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.

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