Future Computing Systems Volume 1. Number 2. 1986

Interactive Elicitation of Knowledge from Experts

MILDRED L. G. SHAW and BRIAN R. GAINES

Department of Computer Science, University of Calgary, Alberta, Canada T2N 1N4

Abstract

The knowledge engineering necessary to set up an expert system in a new domain is currently labour-intensive. It requires skilled knowledge engineers to elicit and codify the knowledge of an expert. In this paper a description is given of a methodology for knowledge elicitation based on personal construct psychology. It has been implemented in computer programs that interact with experts to enable them to express the constructs underlying their knowledge. The programs use interactive elicitation of repertory grids to elicit knowledge from experts, and entailment analysis of the grids to derive inference rules with which to prime expert system shells. Examples are given of the knowledge elicitation and analysis programs in action of the automatic generation of rules and their application in expert systems, of experiments to validate the methodology, and of recent developments in graphic interaction for knowledge elicitation.

Knowledge engineering

The initial success of expert system 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, and Lenat 1983) has opened up the possibility of widespread usage of expert systems. In particular, the Japanese Fifth Generation Computing System 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 expert systems in new domains. We need to understand more about the nature of expertise in itself (Hawkins 1983) and to be able to apply this knowledge to the elicitation of expertise in specific domains.

The problems of knowledge engineering have been stated clearly by Hayes-Roth *et al.* (1983):

> 'Knowledge acquisition is a bottleneck in the construction of expert systems. The knowledge engineer's job is to act as a go-

© 1986 Oxford University Press & Maruzen Company Limited

> 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 diffifficult aspects of the knowledge engineer's task is helping the expert to structure the domain knowledge, to identify and formalize the domain concepts.'

The problem of knowledge elicitation from a skilled person is well known in the literature of psychology (Nisbett and Wilson 1977). Bainbridge (1979, 1986) has reviewed the difficulties of verbal debriefing and notes that there is no necessary correlation between verbal reports and mental behavior, and that many psychologists feel strongly that verbal data are useless. 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 bypass cognitive defences, including those resulting from automization of skilled behavior. Welbank (1983) has reviewed many of these techniques in the context of knowledge engineering.

This paper is concerned with interactive computer aids to the knowledge engineer based on Kelly's (1955) *Personal Construct Psychology*, which uses a model of the person that has strong systemic and psychological foundations (Gaines and Shaw 1981a, Shaw and Gaines, 1981a, 1981c, 1986, Mancuso and Adams-Webber 1982). Personal construct psychology provides a model of human knowledge representation, acquisition, and processing that has been made operational through computer programs for interactive knowledge elicitation (Shaw 1980, 1982). These may be used in developing the expert's vocabulary and in encoding aspects of his reasoning for a rule-based system. The expert interacts with a program (Shaw and Gaines 1981b, 1984) that encourages him to make clear the distinctions he uses in applying his expertise. This helps him to structure his knowledge and identify and formalize his concepts.

The methodology may also be used in a teaching mode in order to enable others to come to use the expert's vocabulary in the same way as he does. Some experiments have been de-

scribed (Shaw and Gaines 1983a) to validate the use of the methodology in this way by reconstructing the basic distinctions used in the Business Information Analysis and Integration Technique (BIAIT) used to determine the accounting needs of a company (Carlson 1979, Sowa 1984).

Personal construct psychology

Kelly developed a systemic theory of human cognition based on the single primitive of a *construct*, or dichotomous distinction. For an individual, constructs are (Kelly 1955):

> 'transparent templets which he creates and then attempts to fit over the realities of which the world is composed.'

He proposes that all of human activity can be seen as a process of anticipating the future by construing the replication of events (Kelly 1955):

> '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.'

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

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

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

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

What is a repertory grid?

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

The *elements* are the things which are used to define the area of the topic, and can be concrete or abstract entities. For example, in the context of: inter-personal relations the elements might be people; attitudes to life they might be critical events; job change they might be careers; expertise about metal joining they might be types of rivet; expertise about medical diagnosis they might be symptoms. Before choosing the set of elements, the user must think carefully about the area of the topic and relate the elements to his purpose. The elements should be of the same type and level of complexity, and span the topic as fully as possible. It is usual to start with about six to twelve elements.

The universe of discourse is determined by the elements. The elements originally suggested by Kelly in his work as a psychotherapist were role titles such as: *Self*, *Mother*, *Father*, *Best Friend*, *Threatening Person*, *Rejected Teacher*. This has been carried over into industry with such role titles as: *Myself*, *My Boss*, *My Boss's Boss*, *Subordinate*, *Person Likely to Get On*, *Person Not likely to Get On*, and so on. The subject in both cases is required to supply names of individuals well known to her/him to fit these and other roles as closely as possible. When choosing elements, care must be taken to ensure that each one is well known and personally meaningful to the subject. Each element must be significant to the person in the context of the particular problem.

The *constructs* are the terms in which the elements are similar to, or different from each other. Each construct therefore has two *poles*, each of which has a meaning with respect to its opposite. Any construct or dimension of thinking which is important to the subject is a valid construct. For example, to distinguish between people by saying that x and y are 'blue eyed' whereas b and c are 'brown eyed' may be trivial, and not concerned with the important qualities of x, y, b, and c. However, if you are an eye specialist concerned with prescribing tinted contact lenses, this may be a significant construct. Thoughts

and feelings, objective and subjective descriptions, attitudes and rules-of-thumb all constitute valid constructs. The verbal description of the construct and the labelling of the poles need not be a publicly agreed meaning in the outside world, but only a memory aid to the thinking process. The mapping of the elements onto the constructs produces the two-dimensional grid of relationships.

Eliciting constructs—triadic elicitation

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—this being the least number which will produce both a similarity and a difference. 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 the three Artificial Intelligence books *Handbook of AI*, *Winston's AI*, and *AI Applications for Business* in what way are two alike and thereby different from the other one? We might first of all say that *Handbook of AI* and *AI Applications for Business* are alike since they are *multi-authored*, whereas *Winston's AI* is *single-authored*. 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 (1) being *multi-authored*, or (2) being *single-authored*, shown in Figure 1. This also shows the significance of the term *personal* in *personal construct* since it would not obviously be a publicly agreed description that *Winston & Horn LISP* is single-authored whereas Davis and Lenat's *Knowledge-Based Systems in AI* is multi-authored. In this case if reflects a single concerted effort as opposed to more than one topic.

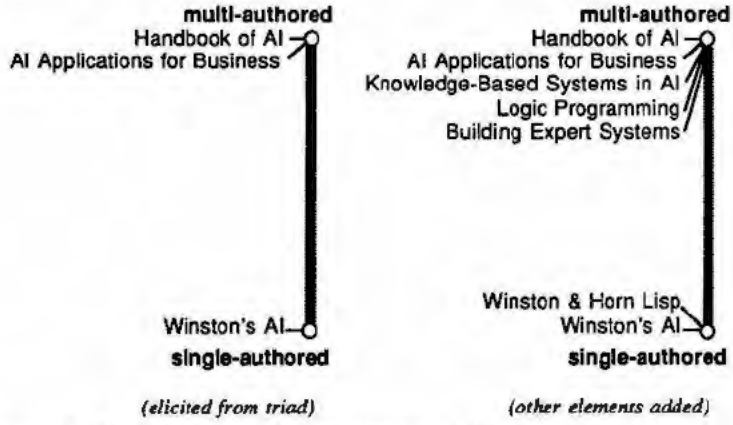

Figure 1. Triadic elicitation of a construct about AI books.

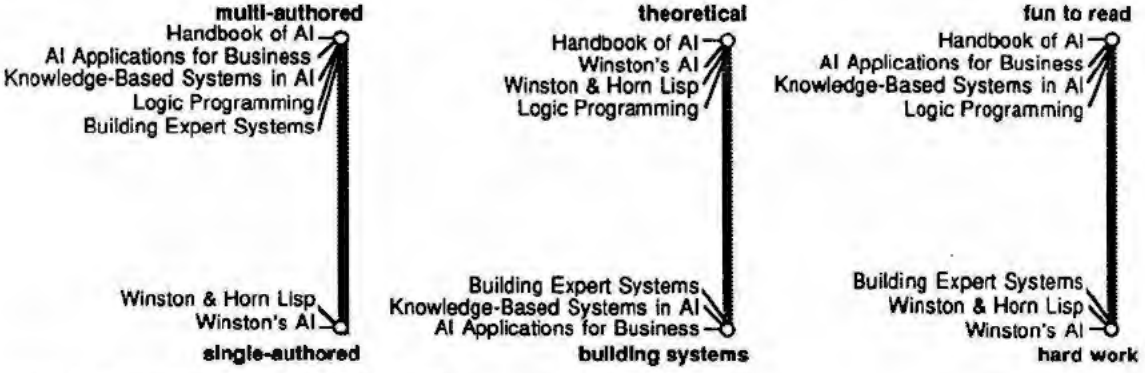

Figure 2. Three constructs applied to AI books.

Then the second and subsequent constructs are elicited in exactly the same way using different triads each time. Figure 2 shows the grid of Figure 1 with a further two constructs elicited. The third construct shown here illustrates that constructs can be factual, imaginary, emotional, or whatever is important to the person generating the grid.

Eliciting constructs—rating scales

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 3. Thus, in this case, the third construct means that *Logic Programming* and *Knowledge-Based Systems in AI* are considered the most *fun to read* books, *Winston's AI* is both *fun to read* and *hard work*, and *Building Expert Systems* and *Winston & Horn LISP* are the most *hard work*.

Note how the use of a multi-point scale with an odd number of values allows for a central rating 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

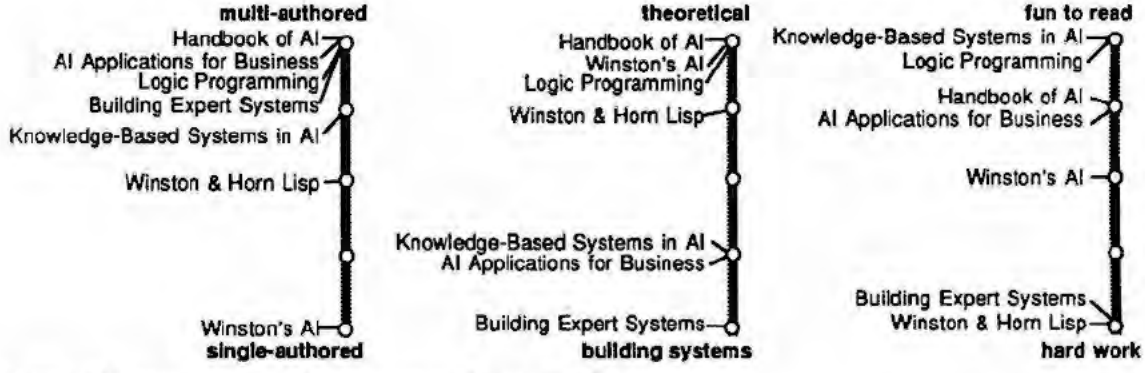

Figure 3. Three multi-point scale constructs applied to AI books.

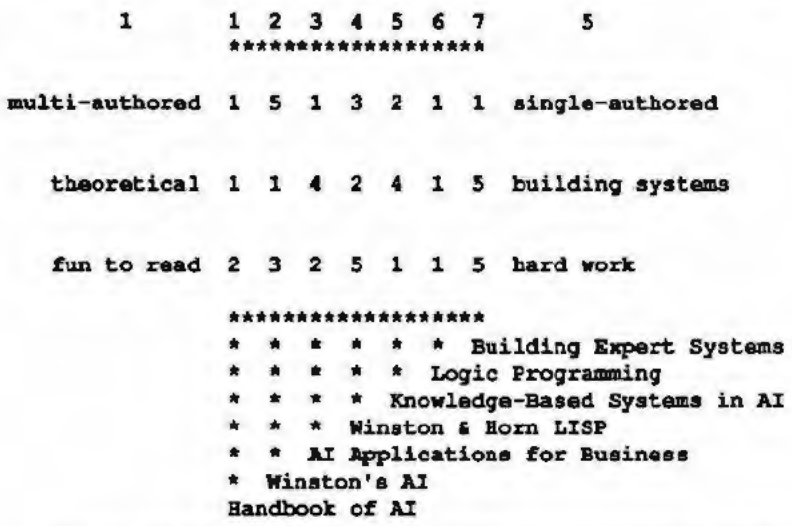

Figure 4. Repertory grid data structure for three multi-point constructs applied to AI books.

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 ratings on each pole (Shaw and Gaines 1980, Gaines and Shaw 1981b).

The mapping of the elements onto the constructs produces the two-dimensional grid of relationships which can be represented as a numeric data structure as shown in Figure 4. This structure may be viewed as a component of a database in entity–attribute form (Chen 1980): a repertory grid has elements as *entities*, constructs as *attributes*, and allocations of elements to poles of constructs as *values*.

Eliciting constructs—other techniques

In contrast to triadic elicitation, the full context form involves the subject considering the whole set of elements, then thinking of an important way they divide into two (or more) groups. Other methods are splitting the two most similar elements to put one on each pole of a new construct, and naming a construct to describe the difference between the two most dissimilar elements. Variants on these methods include: presenting triads as in the minimum context from but just changing one element of the triad for each construct; including the element 'self' in each triad, if relevant; and asking questions about the circumstance in which the subject might be in the company of the group of elements, if appropriate.

Another useful form of construct elicitation is a technique developed by Hinkle (1965) called *laddering*. Initially, constructs are elicited using a standard method, then one con-

struct is selected and the question asked: 'Which pole is the one you would prefer to be rated on?', then 'Why would you prefer that pole?' and 'What is the opposite of that?' This then provides a new construct which is more superordinate or generalized than the original one. As this process is repeated, a hierarchy or heterarchy of superordinate constructs is elicited.

There are five criteria that should be applied in all techniques for repertory grid elicitation if the results are to be meaningful, as follows.

(1) Constructs should be permeable. That is, the construct should be applicable to all the elements in the set.
(2) All pre-existing constructs should be elicited as well as ones just thought of in the process of the elicitation.
(3) The verbal labels given to constructs should be able to explain to someone else the ideas behind the construct.
(4) The constructs should look at relations between the elements as well as the subject's relationship to the elements.
(5) Constructs should be bipolar. That is, both poles of the construct should be elicited.

Multiple perspectives

Knowledge engineering for expert systems may be seen as a process of making overt the construct system of the expert. Some of the problems noted in the first section arise because of the essentially individualistic nature of construct systems. A person may be able to use his construct system effectively whilst having no basis for communicating it to others. Two people may use exactly the same construct yet refer to its poles by different names. Two people may use the same names for the poles of their constructs and yet use them in different ways. Two people may use similar constructs at the lower level of the hierarchy and yet have them organized in different systems such that their reactions to the same event are quite different. Two people may have similar constructs at nearly all levels of the hierarchy and yet construe a novel event completely differently.

Other problems arise because 'the expert' is not one person but a group of people operating as a team. The constructs underlying the team's expertise may be partly held in common but, in general, they will be distributed among different individuals. In order to elicit the expertise the various perspectives of the different team members must be elicited and combined. The emphasis on differing perspectives here is fundamental. Personal construct psychology is often seen as concerned with individuals, and knowledge engineering is often presented as concerned with *the* expert. However, particularly in complex real-life situations, it is rare to find all the expertise concentrated in one individual, and we are most often concerned with skills that reside in a group or team.

The theory underlying personal construct psychology makes no separation necessary between the individual and the group (Shaw and Gaines 1981a). We have shown elsewhere how it accounts for the way in which communities are formed by individuals and the way in which social norms in communities also act to form individuals (Shaw 1985). The individual can be regarded as a cross-section of the epistemological hierarchy of the community in which he is operating (Shaw 1986a). These considerations are those of Pask's *conversation theory* (Pask 1976), particularly the notion of expertise distributed across a multi-person *P-Individual*, and it is notable that this theory is being used as a foundation for the development of expert systems (Coombs and Alty 1984).

The next subsection describes a useful framework for taking into account both the variability of individual perspectives and the role of multiple perspectives.

Techniques of soft systems analysis

Checkland (1981) has developed a methodology for system definition in ill-defined situations that he terms *soft systems analysis*, and we have found this useful as a framework for applications of PLANET to knowledge engineering. Soft systems analysis proceeds through a sequence of stages as shown in Figure 5. Systems analysis techniques usually assume that the system is sufficiently well defined to commence building a conceptual model, and commence at Stage 4 in terms of this analysis. This is equivalent in knowledge engineering to assuming that the expert has an existing framework for his expertise, but unless the expert is a proficient and explicit teacher this may not be so. Furthermore Checkland empha-

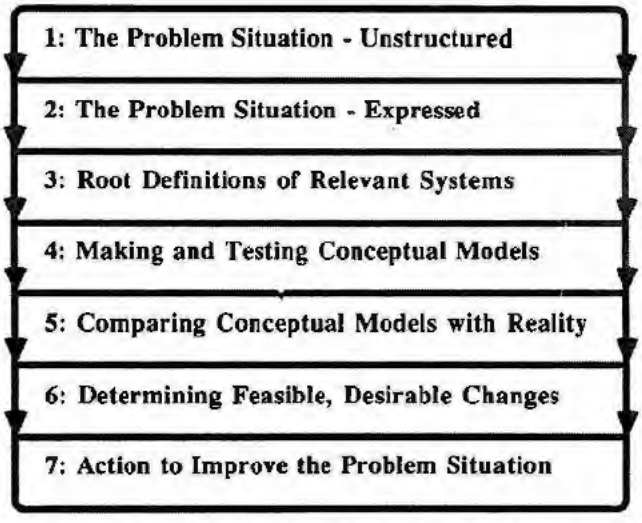

Figure 5. Stages in Checkland's soft systems methodology.

sizes in the context of other domains that the existing framework may be the source of the problems. In terms of knowledge engineering, even if the expert is proficient in explicating his expertise, it may well be that the framework he describes does not correspond very well with the processes actually underlying his expertise.

Hence, soft systems methodology is based on commencing systems analysis several stages back from Stage 4 of conceptualization. Stage 1 is a normative recommendation to attempt to make explicit and put to one side all existing frameworks before commencing the analysis. In terms of knowledge engineering, the knowledge engineer should not commence with a preconceived framework based on his own experience or that of the experts being interviewed. Thus, the importance of Stage 1 is to get back to it from Stage 2 or beyond. The knowledge engineer should not inject preconceptions about solutions into the problem definition, and should treat any situation as unstructured no matter what structure he is given. Otherwise his framework for analysing the expert's knowledge may be such as to prevent that analysis. Kelly (1955) places a similar emphasis in applications of personal construct psychology on the possibility that problems may be created by the use of rigid constructs that paint us into corners. A common source of these in knowledge engineering is the 'text-book' knowledge imparted to the expert in his training. This may be at variance with the skills he has developed and yet still be the model on which he draws to describe his own behavior.

The expression at Stage 2 is intended to be as broad a framework as possible so that no aspects of the problem definition or sources of expertise are left out even if the resulting conceptual framework appears highly unstructured. The structuring of this open-ended framework takes place at Stage 3 where Checkland introduces the notion of *root definitions* of relevant systems, and it is this that provides the formal techniques of soft systems methodology. Stage 3 emphasizes the pluralism of systems analysis, what Kelly had termed *constructive alternativism*, that it is important to examine the problem from a number of viewpoints. These viewpoints may be seen as forming a nested set to which Checkland gives the mnemonic CATWOE as shown in Figure 6. The system is defined through a *Transformation* carried out by people who are the *Actors* within it; it affects beneficially or adversely other people who are its *Customers*, and there is some agency with power of existence over it who is its *Owner*; it has to exist within outside constraints forming its *Environment*, and the whole activity of system definition takes place within an ethos or *Weltanschauung* that affects our views of it. Basden (1983) has noted the utility of Checkland's analysis in developing expert sys-

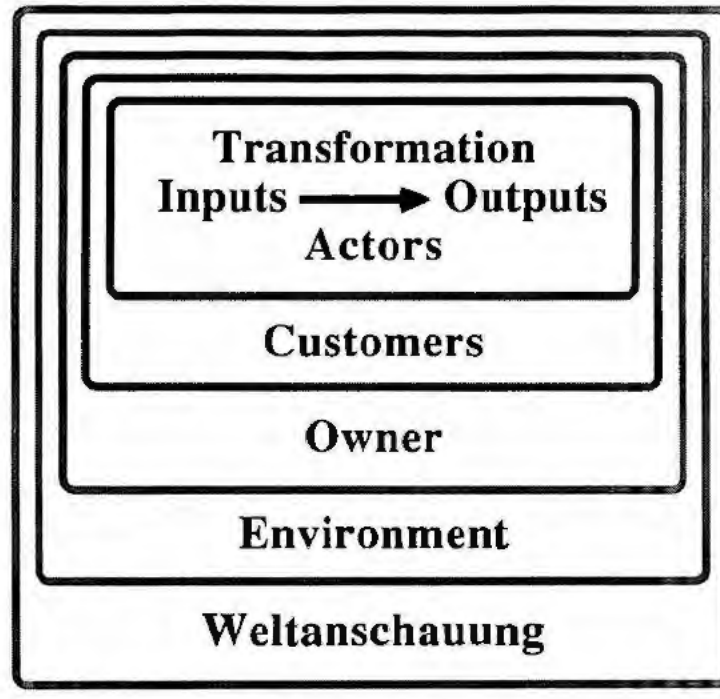

Figure 6. Sources of root definitions relevant systems.

tems, and we have shown how the techniques outlined in this paper may be applied to operationalizing it (Shaw and Gaines 1983b).

Figure 7 shows how the analysis may be applied to knowledge engineering for an expert system concerned with oil exploration. Consider the *transformation* performed by an advisory system on oil exploration to which the *inputs* are exploratory survey data such as seismic surveys and the *outputs* are advice on where to drill. The *actors* are oil exploration data interpretation experts concerned with making these recommendations. Their *customers* are the drilling divisions of oil companies. The *owner* of the problem represents the value system and, as usual in soft systems analysis, there are several choices: the managers, directors, and owners of the oil company are possibilities, as are its customers and the governments that tax it. They have differing perceptions of the problem risk/reward ratios, and so on. The *environment* is the physical, geographic, and economic world within which extraction takes place, and the *weltanshauung* comprises energy policies, pollution concerns, political relationships, and so on. These are the considerations which knowledge engineering for an equivalent automated advisory system has to take into account.

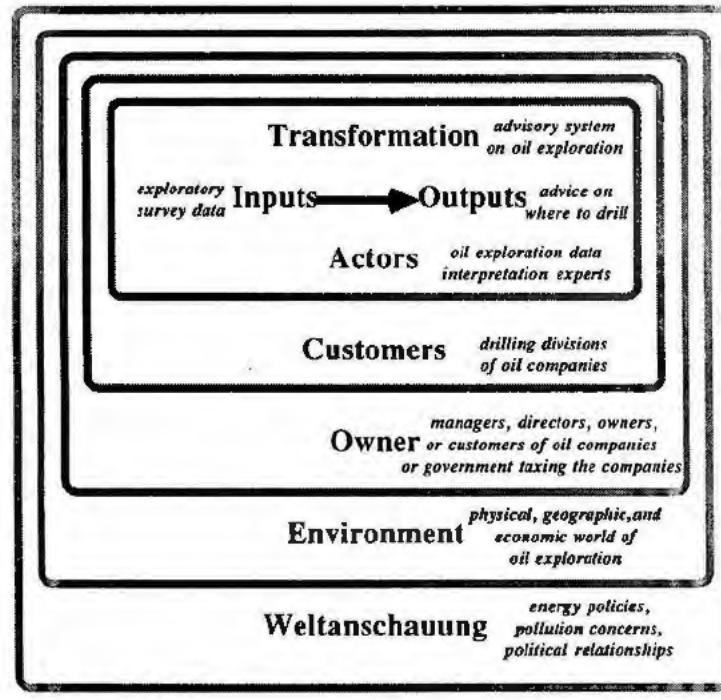

Figure 7. Sources of root definitions of relevant systems for oil exploration advisory system.

Checkland's methodology structures the task of knowledge engineering and provides techniques that are widely applicable. However, it leaves the burden of analysis to a person, and has no operational form as a computer program. The next sections show how the personal construct psychology approach and the use of computer-based repertory grid techniques can automate the analysis in major part.

PLANET: a computer-based knowledge engineering system

PLANET (Personal Learning, Analysis, Negotiation and Elicitation Techniques, Shaw 1982) is a system for eliciting personal construct systems from a variety of perspectives and then comparing and contrasting the systems elicited. It runs as a menu-driven suite of interactive programs on a variety of computers. The programs concerned with repertory grid elicitation and analysis divides naturally into four groups.

Construct elicitation

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 and Bannister 1977). The coming of powerful, low-cost per-

sonal 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 to the techniques used by people in a form where they can be administered through a conversational interactive computer program. PLANET contains two such programs: PEGASUS, which elicits a single grid from a person and feeds back comments and suggestions during the elicitation; and ARGUS, which elicits a set of grids from a person using a number of different perspectives.

PEGASUS is a highly interactive program that elicits constructs through a conversational dialog with an expert, generated by feedback of its ongoing analysis of his construct system. PEGASUS offers many options one of which is to receive feedback from a stored bank of expert constructs, thus enabling a student to learn to use an expert's construct system. An example of PEGASUS in action is given later in this paper.

ARGUS is a variant of PEGASUS that puts one individual in the position of several others and elicits his view of their viewpoints as well as his own (Shaw and McKnight 1981). Thus it elicits a set of grids simultaneously from one person holding several roles or points of view. Firstly, the elements (roles) are elicited followed by three constructs from fixed triads. These construct labels are then used for a new set of ratings to be entered for each roll name in turn, and at each stage another construct which is felt to be important for that role is added. The final set of grids, all with the same element and construct labels, can be processed with the structure analysis programs SOCIOGRIDS, CORE, or MINUS described below.

Repertory grid analysis

Repertory grids in themselves encode information about a person's way of looking at the world. This information can be used in its own right for some purposes since it is an aid to remembering the basis for decisions and actions. It can also be analysed in a variety of ways to bring out possible underlying structures, or *construct systems*, in a person's world view and its relationship to those of other. There are a number of forms of analysis that are widely used for different purposes, and PLANET offers all the commonly used techniques plus new developments in recent years. What form of analysis should be used in a particular case is partly a matter of personal preference and partly a matter of purpose. Comparisons have been made in the literature of different analyses with the same data (Shaw 1981). In exploring the use of repertory grids for knowledge engineering in a specific domain it is worth repeating such comparisons with familiar data to determine what are the most applicable analyses and presentations.

It is convenient to divide analysis techniques into three groups: those concerned with the analysis of a single grid; those concerned with a pair of grids; and those concerned with a group of grids.

Single-construct-system analysis

For any given construct we may regard the numbers in the grid as a *vector of values* giving the assignment of each element in turn to one or other of the poles of the construct. From this point of view each construct becomes represented as a point in a multi-dimensional space whose dimension is the number of elements involved. A natural relation to examine between constructs is then the *distance* between them in this space. Two constructs which are zero distance apart are such that all elements are construed in the same way in relation to them and hence we might infer that they are being used in the same way—in some sense they are *equivalent* constructs. For constructs which are not equivalent we may analyse the entire constellation in space to determine a set of axes such that the projection of each construct onto the first axis accounts for most of the distance between them, the projection onto the second axis accounts for most of the remaining distance, and so on. This is a *principal components analysis* of the construct space, and it is related to the *factor analysis* of semantic space used in the study of semantic differentials (Osgood, Suci and Tannenbaum 1957). We may also group constructs together that are close together in space, using some form of *cluster analysis.*

The FOCUS algorithm is a distance-based hierarchical cluster analysis technique that sorts the constructs into a linear order such that constructs closest together in the space are also closest together in the order. It has the advantage in presentation that the sorting is used only to re-present the original grid re-organized by the *neighbourness* of constructs and elements. It is left to the user to construe his own personal meaning into the result and confirm this directly in terms of the original data. Element and construct matching scores are computed by subtracting the rating values in each position of the two elements of constructs, summing them, then mapping the result into 0%–100% similarity for the elements, and −100%–100% similarity for the constructs. Two matrices are thus formed: one for elements, and one for constructs. Clusters are computed by selecting the highest numbers from these matrices—that is, the most similar in terms of the ratings given—then continuing this process for the next cluster, until all elements and constructs are incorporated. The tree diagrams show these clusters and are imposed on the re-sorted original grid data (Shaw 1980). The program FOCUS provides a hierarchical clustering of an expert's construct system

that preserves the data elicited from him so that the sources of the analysis are evident and can be discussed.

The PRINGRID algorithm is another distance-based cluster analysis using standard principal components analysis techniques and giving the same results as Slater's INGRID (Slater 1976, 1977). This analysis has been used widely in clinical studies with the repertory grid because it gives a visually meaningful map of some of the relations between elements and constructs. The PRINGRID program gives a non-hierarchical cluster analysis based on principal components that can be used to gauge the major dimensions along which an expert is making distinctions.

It is possible to give a logical analysis of repertory grids by regarding constructs as predicates applying to elements. The ENTAIL algorithm is a logical analysis of the construct system, taking the expert's distinctions to be fuzzy predicates (Zadeh 1972, Gaines 1976). The ENTAIL program derives asymmetric implications between the construct poles so that one can infer how a new element might be placed on one construct given how it is placed on others (Gaines and Shaw 1981b, Shaw and Gaines 1982). It can also provide rules for input to an expert system shell (Boose 1985, Gaines and Shaw 1986a). This is discussed in more detail with the example given in a later section.

Analysis of a pair of construct systems

There are often times when it is desirable to compare several repertory grids. Shaw (1980) discusses *exchange grids* for the measurement of understanding and agreement between either two people or two occasions. To do this, two people, possibly experts with differing points of view, each elicit a grid in an area of common knowledge or experience. Each may choose his own elements independently of the other, and elicit and rate his constructs quite separately. Each then makes two copies of his grid, leaving out the rating values. Both these copies are filled in by the other person, one as he himself uses those constructs on those elements and the other as he thinks the original was completed. By comparing pairs of these grids it is possible to map the extent of overlap of the agreement and understanding between the two people.

There are three ways by which two grids can be compared. The first is to concatenate the grids having the same elements and use the algorithms of FOCUS, ENTAIL, or PRINGRID as if they were one large grid. One can then explore the interaction of ideas by examining mixed clusters of constructs from the two grids.

A second way is to use the MINUS algorithm, which requires the two grids to have the same element and construct names, and shows the differences between them by subtracting

ratings in equivalent positions. The result shows where there was agreement, indicated by areas of blanks, and grades of difference, indicated by increasing numbers. The MINUS program compares two grids based on the same elements and constructs to highlight the differences between them.

A third way is to use the CORE algorithm, again starting with two grids having the same element and construct names. CORE examines the most changed element and construct alternately, identifies it, and allows it to be deleted from the grids. In this way the core elements and constructs which show agreement or understanding in the two grids are identified and recorded. The CORE program provides an interactive comparison of a pair of grids to determine where there are differences and to extract invariancies, for example over time.

Analysis of a Group of Construct Systems

The SOCIOGRIDS algorithm analyses a set of repertory grids elicited from a group who share the same elements. It compares every pair using the FOCUS algorithm and produces a set of *socionets* which indicates the links of similar construing within the group. A *mode grid* is also produced, showing the constructs which are readily understood by the majority of the group, and this is filed for future processing on one of the structure analysis programs. Every construct which has been used is listed in order of how highly matched it is across all grids, and indicating where it originated.

This technique is implemented in the SOCIOGRIDS program, which can be used for a wide variety of purposes concerned with group structures, cultures, and relationships. It has been used for investigating the relative positions of members of a small group, and the content of the shared terms and values. It has been used to investigate management hierarchies defined operationally through the interrelationships of construct systems, and in conjunction with the Delphi technique (Linstone and Turoff 1975) to promote common values and understanding in a group.

Database Administration

DATA, INPUT, OUTPUT, and other programs within PLANET provide the facilities needed to manage the databases of construct data generated from field studies of expertise.

The Overall PLANET System

PLANET provides a set of tools for the knowledge engineer, enabling him to perform Stage 3 of Checkland's soft systems analysis semi-automatically. Figure 8 shows the applications of the main programs at each stage of knowledge engineering. PEGASUS enables the construct systems of all those involved in the application of expertise to be elicited. The experts involved can all be similar in nature or can have very different

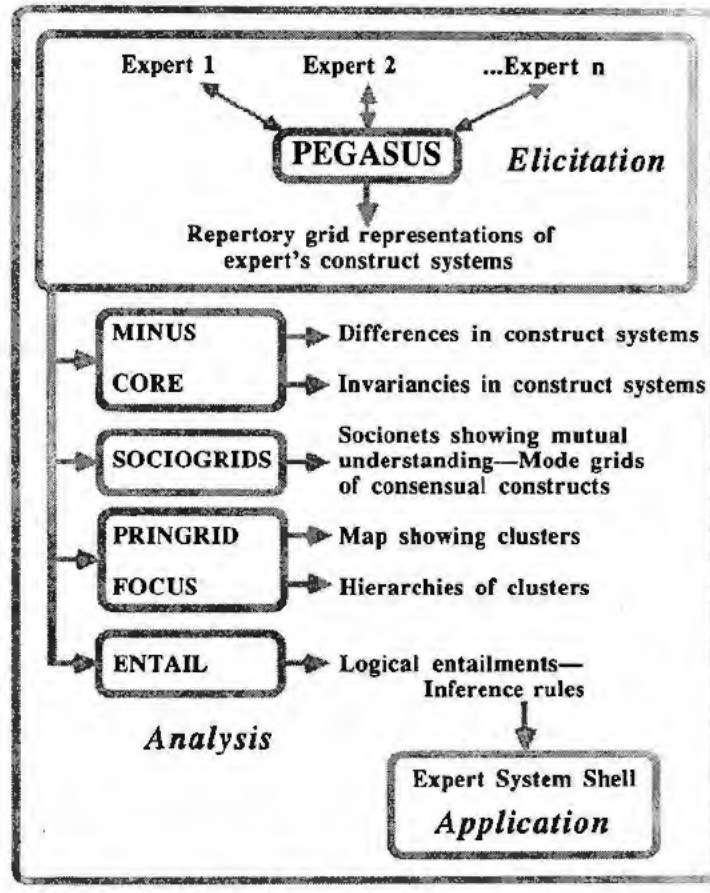

Figure 8. Knowledge engineering facilities available in PLANET.

perspectives, such as *actor*, *customer*, and *owner*. MINUS and CORE enable differences and similarities between expert perspectives to be clarified. SOCIOGRIDS enables expert's understanding of one another's perspectives to be derived, and a consensual grid to be generated. PRINGRID and FOCUS exhibit the relations between constructs and elements through different forms of clustering. ENTAIL induces the logical entailment relations between constructs, displays these, and outputs them as inference rules. These rules can be turned into a decision tree or used to prime an expert system shell.

The following sections illustrate the main programs in PLANET in use in a knowledge engineering context.

PEGASUS in action

This section shows a sample of the PEGASUS dialog in eliciting expertise from a manager relating to the appraisal of his subordinates (Shaw 1980). PEGASUS commences its interaction by asking for the purpose of the elicitation and then for

```
Type in your purpose for doing this grid?staff appraisal

Name six elements.

You must choose a set of six elements keeping in mind why you
want to do this grid. They could be people, events, pieces of
music, pictures, books, or what you want but whatever you
choose they must be of the same type and each must be well
known to you. Try to choose specific things. Now type each
one after each question mark.

Do not forget to press the RETURN key after each.

  Element 1?S1
  Element 2?S2
  Element 3?S3
  Element 4?S4
  Element 5?S5
  Element 6?S6

Here is a list of your elements.

   1  S1
   2  S2
   3  S3
   4  S4
   5  S5
   6  S6

Do you want to change an element?no

Can you think of another element that belongs with the 6 that
you have got so far?y

  Element 7?S7
```

Figure 9. PEGASUS: initial element selection.

a set of elements relevant to this purpose, as shown in Fig. 9. It is important to establish an initial context such that the expert selects a set of elements which are of the same type and relevant to the purpose of the elicitation. Note that in the dialog the computer output is shown in normal face and the human input in *italic*.

One technique used by PEGASUS to elicit constructs is that of triadic comparison, as shown in Fig. 10. A combination of three elements is the minimum necessary to establish both a similarity and a difference. The poles of personal constructs are not just a logical predicate and its negation. Construing the three elements sets up a context, or a perspective, which corresponds to the logical notion of *relevance*. The two poles of the construct anchor the extremal ends of the dimension of construing that defines this relevance.

Another technique used by PEGASUS to elicit constructs is that of matching the ratings given to elements and asking for a

```
Triad for elicitation of construct 1
  1 S1
  2 S2
  3 S3

Can you choose two of this triad of elements which are in
some way alike and different from the other one.

What is the number of the one which is different?3

Now I want you to think about what you have in mind when you
separate the pair from the other one. How can you describe
the two ends or poles of the scale which discriminate S1 and
S2 on the left pole from S3 on the right pole?

Just type one or two words for each pole to remind you what
you are thinking or feeling when you use this construct.

  left pole rated 1 --?intelligent
  right pole rated 5 --?dim

Now assume that S1 and S2 are assigned the value 1 and S3 is
assigned the value 5 according to how you feel about them
please assign to each of the other elements in turn a
provisional value from 1 to 5

  1 S1     1
  2 S2     1
  3 S3     5
  4 S4    ?5
  5 S5    ?3
  6 S6    ?3
  7 S7    ?5
  8 S8    ?2
  9 S9    ?3
```

Figure 10. PEGASUS: triadic construct elicitation.

construct that differentiates between two highly matched elements, as shown in Figure 11.

A technique that allows PEGASUS to elicit additional elements is to match constructs across elements and ask for a new element that differentiates between two highly matched constructs, as shown in Figure 12.

Through a combination of triadic comparison, element matching, and construct matching as shown in the flow chart of Figure 13, a PEGASUS interaction elicits from the user his system of constructs relating to the topic being studied. Figure 14 is the final grid of the manager.

PRINGRID and FOCUS in action

The principal components analysis performed by PRINGRID enables the elements and constructs to be plotted in such a way that some relations between them are apparent. Figure 15 shows a plot of the first and second components produced by the PRINGRID analysis of the grid of Figure 14. The positive

```
The two elements 5 S5 and 6 S6 are matched at the 90 percent
level. This means that so far you have not distinguished
between S5 and S6

Do you want to split these?y

Think of a construct which separates these two elements, with
S5 on the left pole and S6 on the right pole.

  left pole rated 1 --?self starters
  right pole rated 5 --?need a push

According to how you feel about them please assign to each of
the other elements in turn a provisional value from 1 to 5

  5 S5     1
  6 S6     5
  1 S1    ?2
  2 S2    ?1
  3 S3    ?5
  4 S4    ?5
  7 S7    ?5
  8 S8    ?3
  9 S9    ?4
 10 S10   ?5
```

Figure 11. PEGASUS: element elicitation by breaking a construct match.

```
The two constructs you called
  1 intelligent -- dim
  4 little supervision reqd -- need supervision
are matched at the 66 percent level

This means that most of the time you are saying intelligent
you are also saying little supervision reqd and most of the
time you are saying dim you are also saying need supervision.
Think of another element which is either intelligent and need
supervision or little supervision reqd and dim. If you really
cannot do this then just press RETURN after the first
question mark but please try.

Then you must give this element a rating value on each
construct in turn. After each question mark type a value from
1 to 5

What is your element?S10

Type in the ratings for this element on each construct.
Left pole rated 1, right pole rated 5.

  intelligent -- dim?5
  willing -- unwilling?2
  new boy -- old sweats?3
  little supervision reqd -- need supervision?3
```

Figure 12. PEGASUS: construct elicitation by breaking an element match

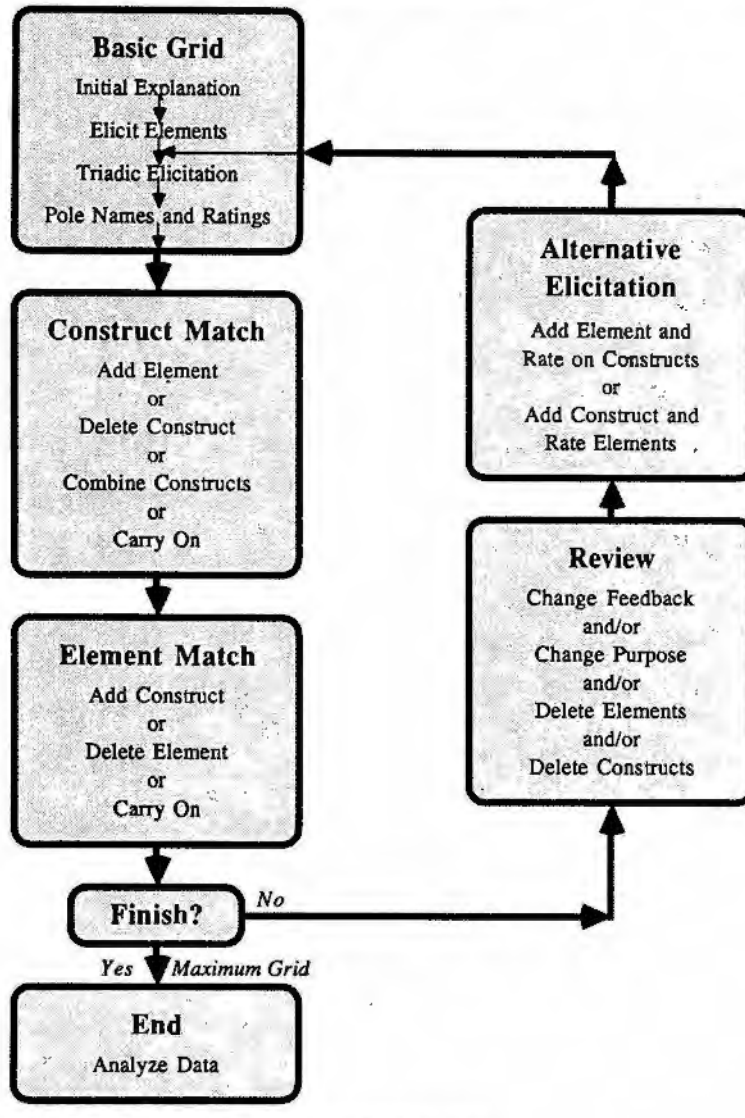

Figure 13. PEGASUS: flow of control in elicitation.

cluster of {*intelligent, motivated, willing, professional*} is apparent together with its neighboring *overall rating high* and *new boy*. The negative cluster of {*dim, less motivated, unwilling, less professional*} is also apparent together with its neighboring *overall rating low* and *old sweat*. The places of individuals such as S5 near the positive cluster and S4 near the negative cluster are also apparent.

A more refined cluster analysis of the grid of Figure 14 is given by FOCUS, as shown in Figure 16. For the constructs, the primary positive cluster linked to *overall rating high* is

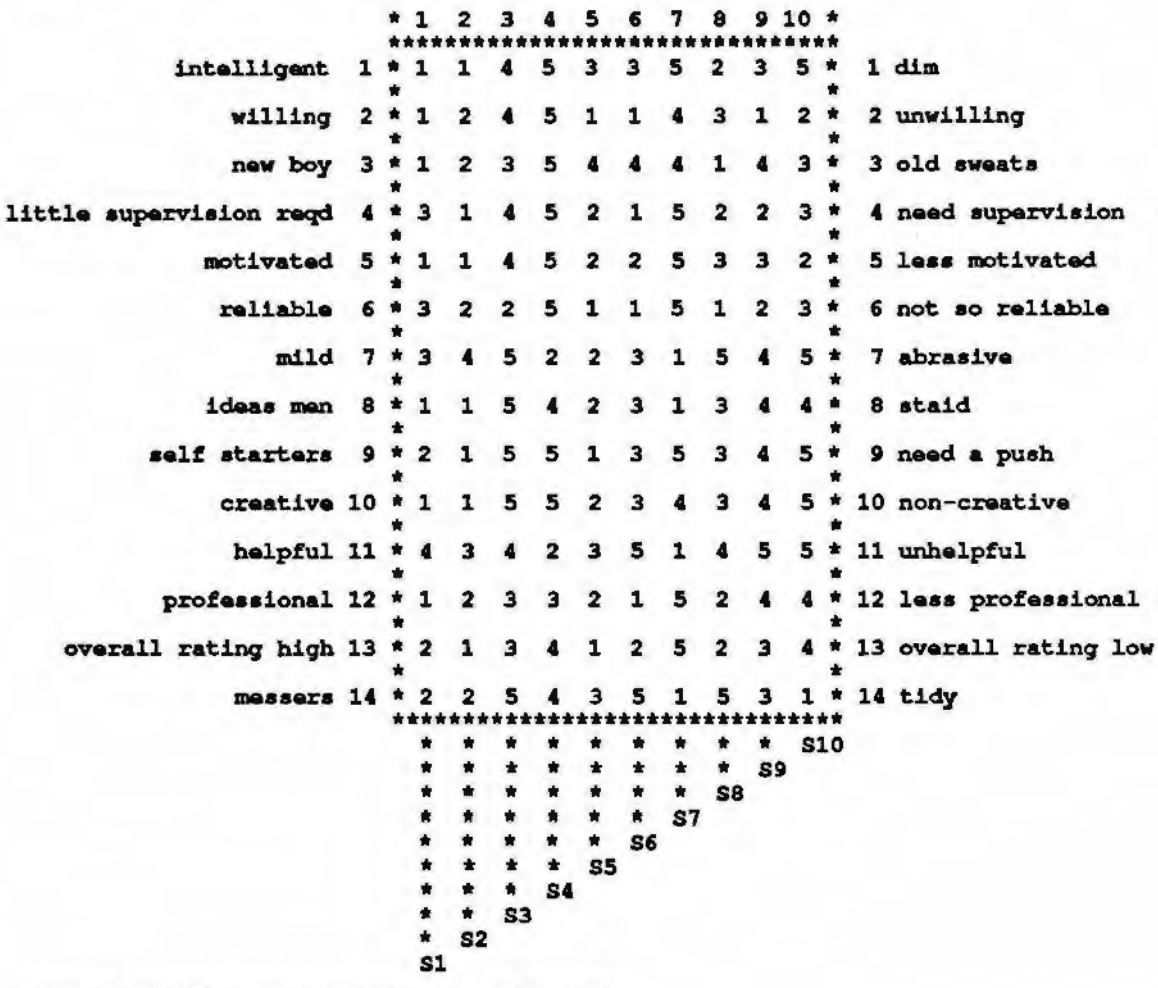

Figure 14. PEGASUS: repertory grid elicited on staff appraisal.

{*professional, reliable, little supervision reqd*}, and this is linked to a secondary positive cluster of {*willing, motivated, intelligent, creative, self-starters, ideas men*} with a lesser link to {*mild, helpful, new boy*}. The negative clusters are the converse of these. For the elements, the primary positive cluster rated *overall rating high* is {S1, S2, S5, S6, S8}, the primary negative cluster is {S4, S7} and this is linked to a secondary negative cluster of {S9, S10, S3}.

ENTAIL In action

ENTAIL gives a dependency analysis of the constructs in a grid by deriving significant logical entailments consistent with the data. From a clustering point of view this analysis is significant because it can show up one-way relations that may be missed by distance-based clustering techniques. However, in

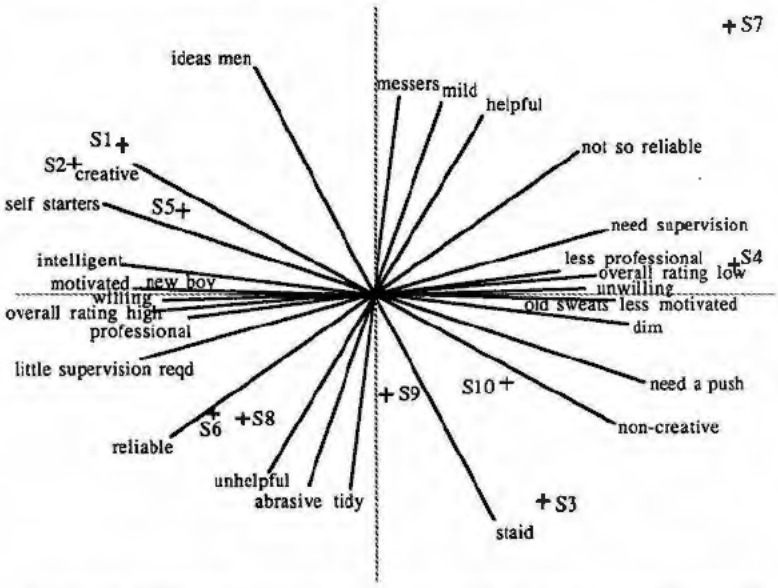

Figure 15. PRINGRID principal components analysis of repertory grid on staff appraisal—first component horizontal, second vertical.

the present context it is more significant that the entailments may be treated as inference rules to be used in an expert system. Figure 17 shows the ENTAIL analysis of the grid of Figure 14. 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 (Gaines and Shaw 1986a).

The data of Figure 14 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 17 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*. The rules form a decision tree expressing the relationship between an expert's data and his conclusions, both personally expressed. For example, the staff appraiser might use these results to determine whether a new

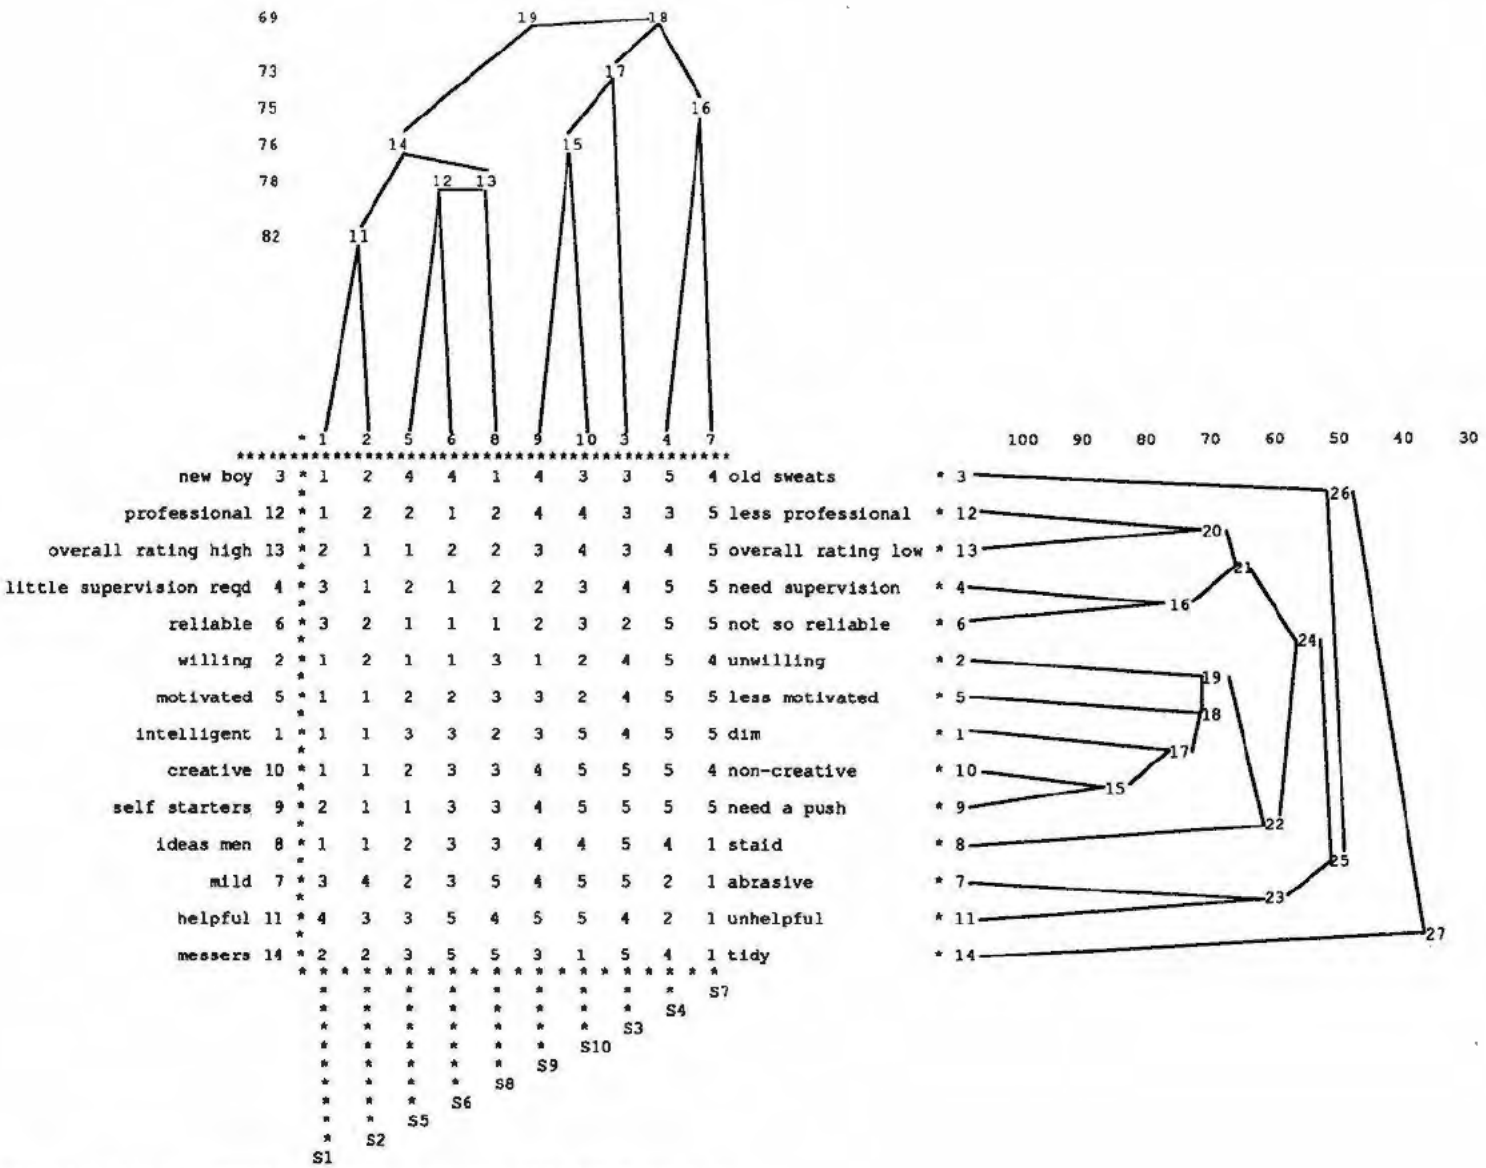

Figure 16. FOCUS hierarchical clustering of repertory grid on staff appraisal.

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	profeasional→overall rating high
L13 → L12	0.80	1.00	0.17	overall rating high→professional

Figure 17. ENTAIL logical analysis of repertory grid on staff appraisal.

subordinate is likely to achieve a high rating based on his evaluation of intelligence, creativity, reliability, professionalism, and the ability to self-start. These criteria might also be used in the assessment procedures for candidates for positions in his division.

The entailments derived by ENTAIL can be converted directly to rules that may be loaded into a standard expert system shell. The truth value of the rule codes into the truth value representation used by the shell. The rules may be checked by the expert for their content, but their prior evaluation for uncertainty reduction makes it unlikely that spurious rules will be generated. Figure 18 shows the Nexpert (Roy 1986) in operation loaded with the entailments of Figure 17. Interaction with Nexpert enables the expert to see the derived rules in action. He can determine their consequences with test date, analyze new hypothetical cases, and see the interrelations 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 revised system using Nexpert's facilities.

Thus PEGASUS and ENTAIL have acted as tools for rapid prototyping of an expert system by priming the shell with a meaningful system that exemplifies what is required to the expert in a context he can readily understand. Without such

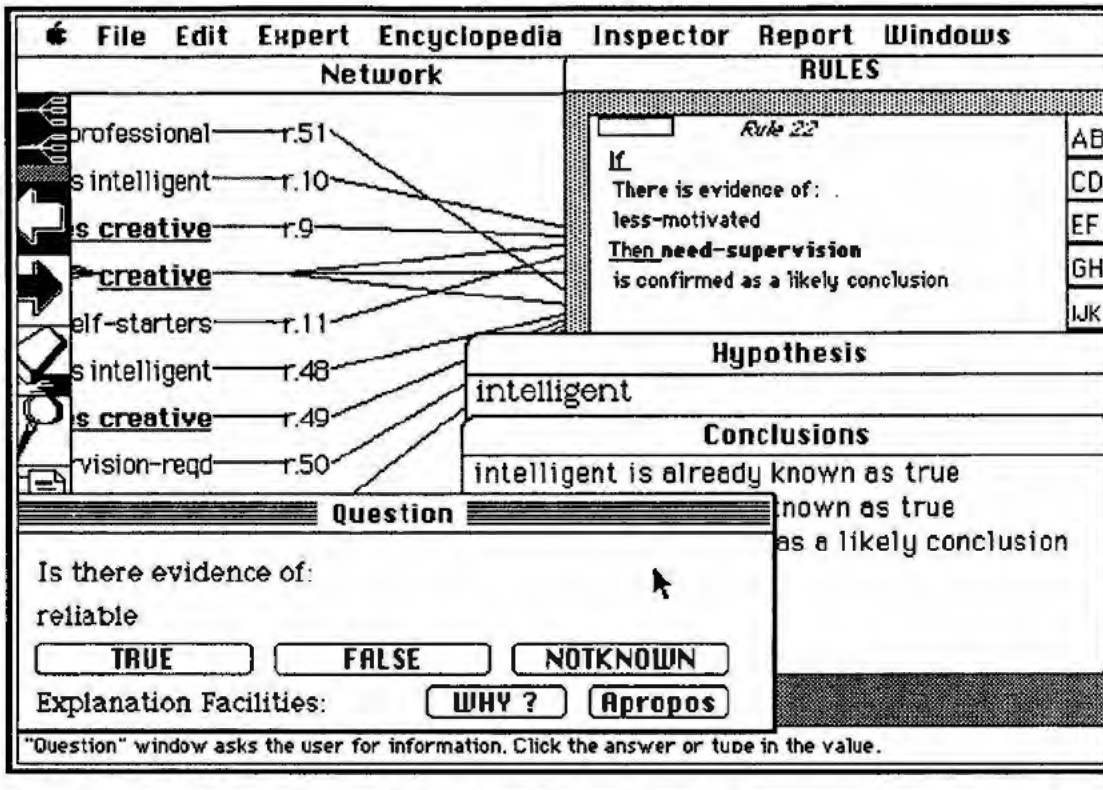

Figure 18. Inference rules, derived by ENTAIL from repertory grid on staff appraisal, in use in Nexpert expert system shell.

prototyping tools the initial priming of an expert system shell is a daunting task and the knowledge engineer has a major role to play as an intermediary between the expert and the system. With such prototyping the expert himself is the dominant partner in the knowledge elicitation process and the knowledge engineer can concentrate on supporting the expert in his use of the tools to express, and further his own understanding of, his knowledge.

Validation—SOCIOGRIDS in action

The methodology described is very easy to implement and run. Experts enjoy using the system, and need far less attention from a knowledge engineer. It is important initially that the expert gains some familiarity with the construing and rating process. It is essential that the expert becomes familiar with the kinds of questions being asked, and finishes his grid with a representative set of elements all at the same level of abstraction, and a representative set of constructs which cover all cases. The best way to achieve this is not through discussion but by experience in the informal use of PLANET, construing books, movies, friends, as an exercise before undertaking the main knowledge elicitation.

We put great emphasis on the need to be able to validate any technique for knowledge engineering. It is not sufficient to show that reasonable expert systems can be developed. One must attempt to evaluate the accuracy and completeness of the knowledge transfer. This is not an easy task because there are few well-established domains of expertise where the accuracy of the elicitation can be tested. One needs clear-cut test cases where the conceptual framework of experts has been both elicited and validated. Sowa (1984) describes one such case where the record-keeping needs of a business can be prescribed through the BIAIT analysis technique (Carlson 1979) based on seven features such as 'Does the supplier bill the customer, or does the customer pay cash?' and 'Does the supplier deliver the product at some time in the future, or does the customer take the order with him?'

A study has been carried out using PLANET to determine whether PEGASUS can be used to elicit the BIAIT distinctions from those with some accounting or business knowledge (Shaw and Gaines 1983a, Shaw 1984). The experiments have been based on taking a spectrum of those with assumed expertise in business record keeping, from businessmen and accountants, to accounting students at various stages of their education, and using PEGASUS to elicit their constructs relating to business transactions. Again in the interests of an initial open-ended approach, the construing has not been tightly constrained to relate to the record-keeping aspects of the transactions. It has been interesting to determine what distinctions

are made that do not relate to record-keeping, or relate to it in ways different from BIAIT.

Various forms of analysis have been applied to the resulting grids. The constructs elicited have been compared with the BIAIT constructs by analysing with ENTAIL the compound grids obtained by merging the constructs elicited from each person with those corresponding to the BIAIT constructs. This has shown that all the BIAIT constructs exist within those elicited. The set of grids obtained has also been analysed with SOCIOGRIDS to see whether the variations of expertise across the group of subjects used shows up clearly. This is a good test of the basic validity of the technique and its foundations in tne notion of construing. For example, in an earlier study (Shaw 1980) of quality control for a garment company the SOCIOGRIDS analysis reconstructed the management hierarchy of the firm from the construct systems of staff involved in construing garment faults.

The operation of SOCIOGRIDS is based on what Kelly terms the *commonality corollary* (Kelly 1955) that: to the extent that one person employs a construction of experience that subsumes that of another in a given domain, that person can understand the cognitive processes of the other person in that domain. Thus, an expert system for a domain can be defined as one which subsumes the relevant constructions of experience of experts in that domain. Figure 19 shows one stage of a socionet sequence produced by SOCIOGRIDS, in which an arrow from one person to another shows that the first person *out-construes* the other. The results are mainly what one would expect, with the businessmen and the economics students closely related to BIAIT. The only disconcerting result is the

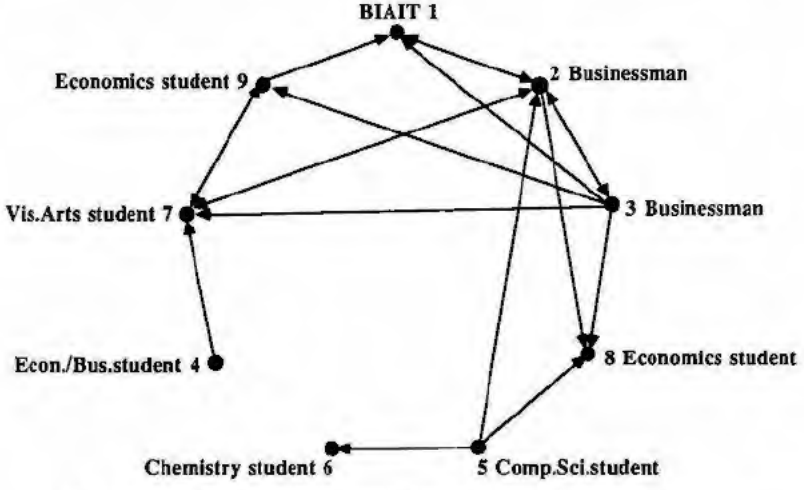

Figure 19. SOCIOGRIDS socionets from BIAIT validation experiment data.

position of the Visual Arts student who is clearly a central part of the cluster. On enquiring into the background of this student, it transpired that she had only recently become a student, having spent the previous 25 years as an administrative secretary in an insurance organization.

Gammack and Young (1985) have reported a further validation of construct elicitation through repertory grids, in the context of statistical expertise. They used fifteen forms of statistical distribution as elements and report such constructs as: *continuous—discrete* and *useful-in-modeling—common-test-statistic*. They report that

> 'hierarchical clustering applied to the data yielded known families of distributions such as the closely related F, gamma and log gamma distributions.'

Validation—expert system developments

The capability of PLANET and related systems to generate useful expert systems rapidly is clearly another important dimension of validation. We have reported its use in developing an expert advisory system for managers responsible for configuring flexible manufacturing systems (Gaines and Shaw 1986b). Boose (1984) has developed an Expertise Transfer System (ETS) based on similar techniques to PLANET and operating on a Xerox 1108 LISP Machine. Figure 20 shows the screen of the machine running ETS at the stage of generating inference rules from the grid ratings. At Boeing Computer Services, Boose, and other users of the system have prototyped several hundred expert systems for a wide range of problems using ETS.

Boose notes 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 (Boose 1985). He assesses its value for rapid prototyping in these terms:

> 'Knowledge engineers feel that ETS can save at least two months of project calendar time. This estimate is made by comparing the time taken for knowledge acquisition for similar projects. The reduction in time results in savings of both the expert's time and the knowledge engineer's time.'

Boeing's success with ETS has attracted widespread industrial interest and, for example, Wahl (1986) reports that Lockheed is also making effective use of repertory grid techniques in knowledge engineering.

Further developments

Knowledge engineering as the process of making explicit the knowledge acquired by human society and held by its experts is an immense problem that may be regarded as a fundamental

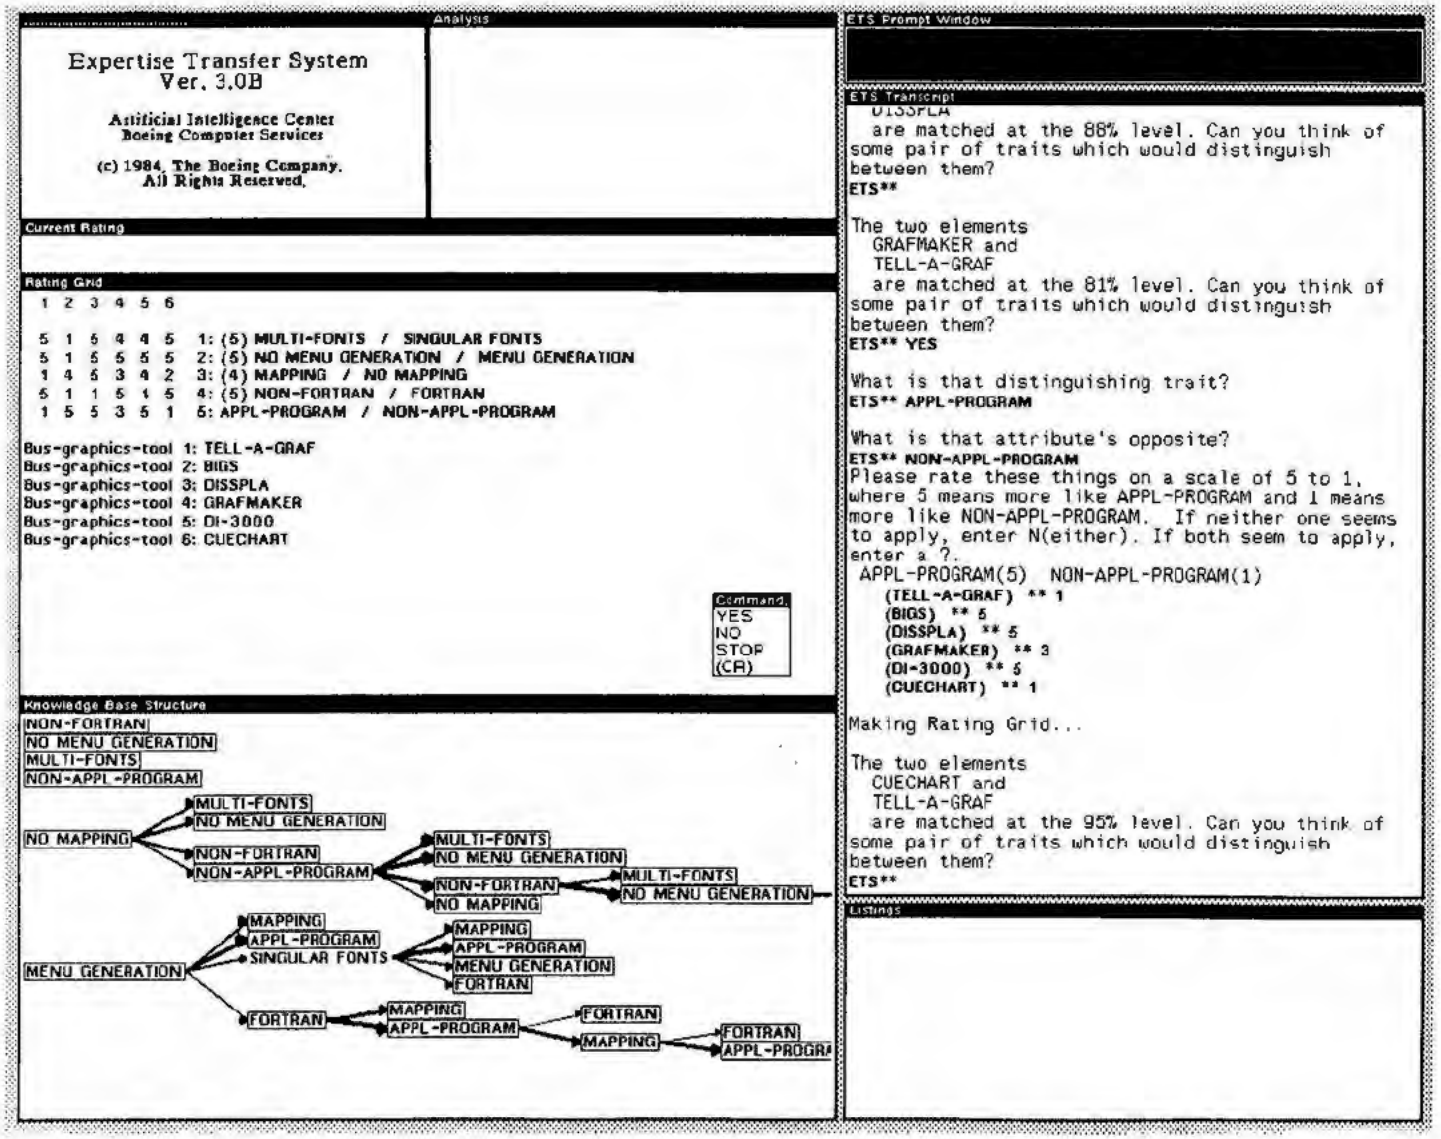

Figure 20. ETS running on a LISP machine with display of dialog, grid, and derived entailments (Boose 1985).

dynamic underlying the development of our society (Shaw and Gaines 1986). Developments in expert systems have not so much generated the problem as focussed attention on its industrial and commercial significance in this information age. Thus, we see developments such as PLANET and ETS as only the first steps towards increasingly advanced techniques for knowledge elicitation, representation, application, and, ultimately, creation.

One important area of advance is to tap the knowledge already explicitly encoded in books, papers, technical reports, and equipment manuals. We have implemented a text analysis that clusters words based on their association as a front end to new grid elicitation programs. The textual analysis will be used to aid the partitioning of the knowledge domain into relatively independent *knowledge islands*, and to suggest elements and constructs to the expert within these knowledge islands.

Kelly (1955) stresses the hierarchical nature of construct systems, and emphasizes the significance of *superordinate* construct systems at the upper levels of the hierarchy and of *core* constructs which we are reluctant to change because they subsume important experience. It is apparent that there are such significant constructs underlying all others and that making them explicit in a given domain has great power as a knowledge represenation technique: for example, the *case grammars* (Fillmore 1968) and *conceptual dependencies* (Schank 1975) used in language processing. We have derived a set of such constructs based on Popper's (1968) three worlds theory and shown that they may be used to generate the key constructs underlying Pask's (1980) conversation theory and Checkland's (1981) soft systems analysis, both of which have already proved useful in knowledge engineering (Gaines and Shaw 1985). We have successfully applied these constructs to analyse the expertise underlying evaluative methodologies in human factors engineering (Gaines and Shaw 1986c), and are developing techniques for incorporating them in the automatic elicitation process.

We emphasize, above, 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. PLANET provides facilities for the analysis of group data. However, PEGASUS elicits grids from one person at a time with no interaction involved and, while it can be used in a group elicitation, the resultant grid reflects the consensus only and does not highlight different perspectives and disagreements. A group elicitation technique has been developed which involves multiple interacting users communicating through terminals on a network. The system elicits the con-

struct systems of a number of users separately but then allows them to exchange their grids 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 system has been called a *Participant Construct System* (PCS, Shaw 1986b) because it follows the paradigm suggested by Chang (1986) of a *participant system* (PS). This captures the notion of computer-based communication between multiple users that is *essential* to the performance of their tasks, rather than merely incidental to their use of a time-shared computer or computer network. One may distinguish a *weak* notion of a PS in which the performance of an individual task involves essential interaction with a system with which others must also interact, such as an airline reservation system. A *strong* notion of a PS is one in which the individual task is perceived by the individual as being only part of a mutual task essentially involving a number of individuals interacting through the PS. It is the strong form of PS that is of interest in this project, where the participants are using the system to further a mutual goal, and the interaction of others with the system is seen by one user as being esential to the fulfilment of his individual objectives:

- 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.

A preliminary version of PCS has been developed on the Macintosh computer, and Figure 21 shows a construct being elicited using the natural click-and-drag techniques of the window/mouse interaction rather than the numeric rating scale of Figure 10. Figure 22 shows the natural representation of the construct match screen giving element elicitation similar to that of the dialog of Figure 12. Figure 23 shows the presentation of agreement when one participant has made his own assignment of elements to constructs elicited from another participant. PCS has been developed as part of a teleconferencing system for group problem solving rather than as a knowledge engineering tool explicitly targeted on the development of expert systems. However, it provides an interesting environment

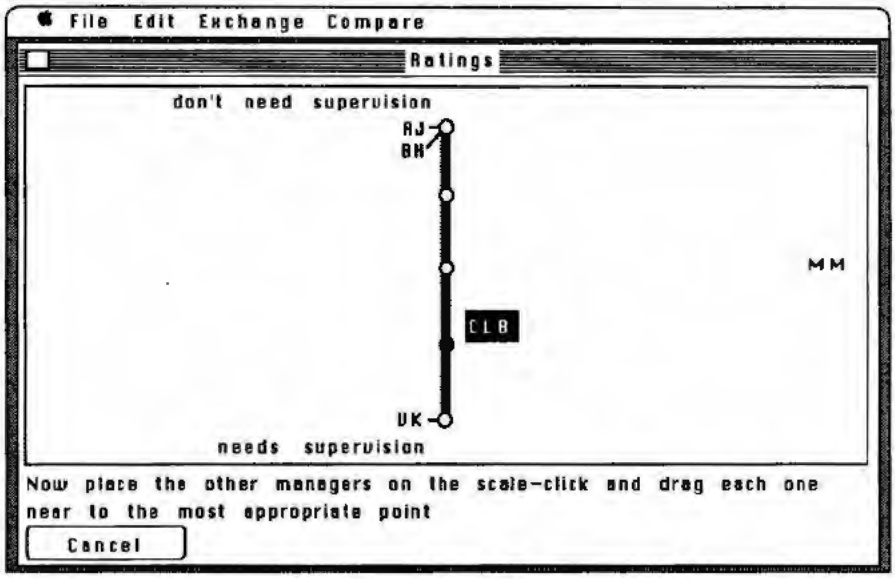

Figure 21. Placing elements on constructs using click and drag in PCS running on a Macintosh.

with which to experiment with group construct elicitation and suggests approaches to group knowledge engineering.

There are many other extensions to existing construct elicitation and analysis techniques that will widen the range of applicability. The rating scale copes with basic constructs that can be placed on an ordinal scale. Boose has noted the desirability of allowing categorial data to be input directly: for

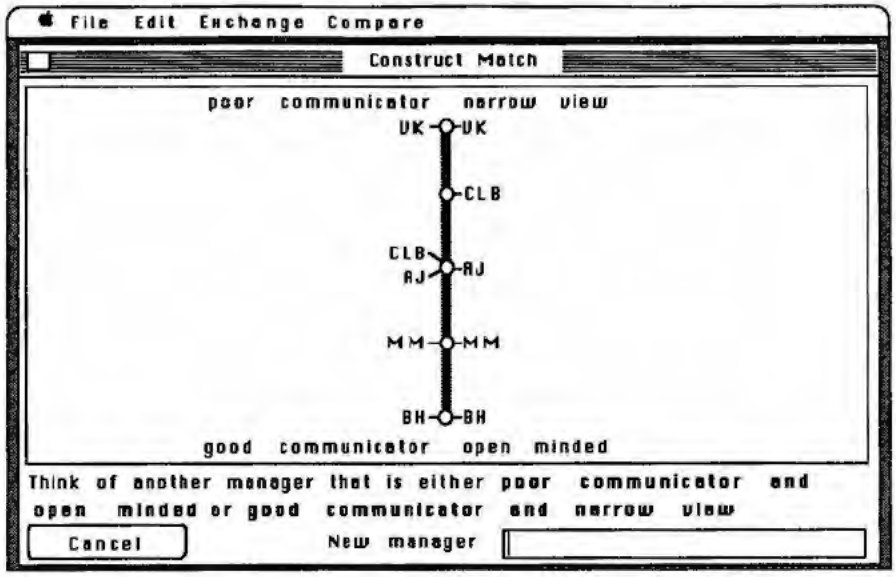

Figure 22. Elements elicitation through construct match in PCS.

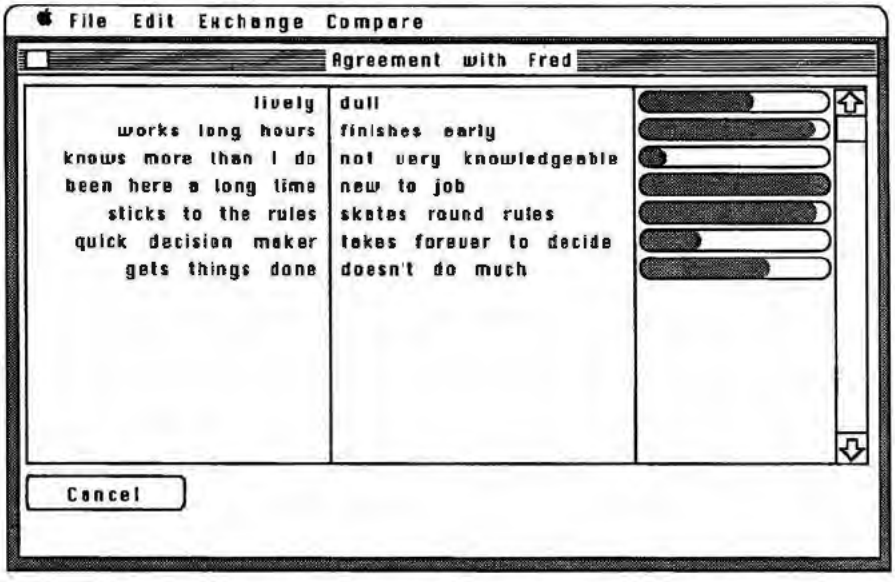

Figure 23. Degrees of agreement when one participant has placed elements on the constructs of another in PCS.

example, colors. For some significant constructs such as space and time there are more refined scales that have additive and ratio properties that may be regarded as precisiations of the underlying informal constructs (Gaines 1976), and it is important that such properties can be used in the analysis. The inductive reasoning techniques used in ENTAIL can be extended to such data using the same techniques as for the variable-precision logics developed by Michalski and Winston (1986).

Categoricity and other mathematical properties may be treated within the framework of personal construct theory as the application of higher-level constructs to more basic constructs. The resulting structures of distinctions are equivalent to the algebraic theories underlying measurement methodologies (Krantz, Luce, Suppes, and Tversky 1971). The development of techniques for eliciting and analysing construct hierarchies is also required to overcome some of the other limitations, noted by Boose (1985, 1986) in that a single repertory grid represents only the knowledge about a particular domain and cannot account for the relationships between multiple domains, such as those involved in causal reasoning. However, there appear to be no intrinsic limitations to the extension of methodologies based on personal construct psychology to an ever-widening range of knowledge representation and elicitation problems. PLANET and ETS were designed as front ends to expert system shells, priming the knowledge base with initial inference rules and thus enabling TEIRESIAS-like

(Davis and Lenat 1982) facilities to come into play. However, it would be attractive to integrate on-line knowledge elicitation techniques with the rule-matching analogical inferences of TEIRESIAS. This would be particularly valuable in attempts to encode knowledge on a large scale such as the MCC encyclopedic study, CYC (Lenat, Prakash, and Shepherd 1986). It would also be valuable to have the inferencing integrated with the elicitation so that the consequences of knowledge entry could be investigaed at any stage. Boose (1985) has incorporated an inferencing decision tool in ETS as part of the validation feedback to the expert. The mixed possibilistic/probabilistic hypotheses produced by ENTAIL arise in other inductive problems (Gaines 1979) and are consistent with a general logic of uncertainty (Gaines 1983) that will eventually be incorporated as an inferencing extension to PLANET.

Conclusions

Personal construct psychology indentifies the system of construction being used by an expert as characteristic of his thought processes, and the elicitation of his constructs as a fundamental methodology of knowledge engineering. Soft systems analysis provides a framework for applying this concept to the ill-defined problems typical of expert system applications. These concepts can be used as a basis for structuring interviews to elicit knowledge from experts and to combine their multiple perspectives. However, manual elicitation is tedious, and in this paper a description is given of automatic computer-based techniques for on-line elicitation, analysis, and feedback. These lead to improved elicitation and have proved very effective in knowledge elicitation for expert systems.

PEGASUS is a program which elicits from an expert the fundamental distinctions, or constructs, he is using in applying his expertise. MINUS and CORE are programs analyzing differences and similarities between experts. INGRID and FOCUS are cluster-analytic programs showing relations in, or between, the construction systems of experts. ENTAIL analyses the logical dependencies between constructs and can be used to develop rules for an expert system shell. SOCIOGRIDS analyses the relations between experts and can be used to develop consensual construct systems. The programs and methodology have been widely used since 1976 not only for expert system development but also for a wide variety of knowledge elicitation studies in psychology, education, and management. Studies have been carried out to validate the methodology experimentally using the BIAIT business record-keeping constructs as a standard.

Boeing Computer Services has implemented a related sys-

tem, ETS, on a LISP machine and produced a wide range of expert systems in an industrial environment. Their experience suggests that current implementations of the methodology are better suited for analysis than for synthesis problems: for example, debugging, diagnosis, interpretation, and classification, rather than design and planning. It is currently difficult to apply the methodology to knowledge engineering involving *deep* causal knowledge or strategic knowledge. However, in situations where it is applicable, Boeing has found that the methodology can save at least two months of project calendar time, and decision support systems can be prototyped in a few hours.

Recent developments of the methodology include: extension of the foundations of the work to include groups of experts and the combination of expert systems and simulation techniques; improved techniques for the derivation of rules from repertory grid data giving a natural knowledge representation for uncertain data, combining fuzzy and probabilistic logics; an information-theoretic measure of the significance of derived rules; improved interactive techniques for on-line knowledge engineering from groups of domain experts interacting through a computer network.

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 Dr John Boose and Boeing Computer Services for permission to use Figure 20.

References

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

BAINBRIDGE L. (1986). 'Asking questions and accessing knowledge', *Future Computing Systems* **1**(2) (this issue).

BASDEN A. (1983). 'On the application of expert systems', *International Journal of Man–Machine Studies* **19**(5) (November) 461–477.

BOOSE J. H. (1984). Personal construct theory and the transfer of human expertise, *Proceedings of AAAI-84* 27–33.

BOOSE J. H. (1985). 'A knowledge acquisition program for expert systems based on personal construct psychology', *International Journal of Man–Machine Studies* **20**(1) (January) 21–43.

BOOSE J. H. (1986). 'Rapid acquisition and combination of knowledge from multiple experts in the same domain', *Future Computing Systems* **1**(2) (this issue).

CARLSON W. M. (1979). 'The new horizon in business information analysis', *Data Base* **10**(4) 3–9.

CHANG E. (1986). 'Participant systems', *Future Computing Systems* **1** (in press).

CHECKLAND P. (1981). *Systems Thinking, Systems Practice* (Wiley, Chichester).

CHEN P. P. (ED.) (1980). *Entity–relationship Approach to Systems Analysis and Design* (North Holland, New York).

COOMBS M. J. and ALTY J. (1984). 'Expert systems: an alternative paradigm', *International Journal of Man–Machine Studies* **19**(5) (November) 461–477.

DAVIS R. and LENAT D. B. (1982). *Knowledge-Based Systems in Artificial Intelligence* (McGraw-Hill, New York).

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

FILLMORE C. (1968). 'The case for case', *Universals in Linguistic Theory* (eds. E. Bach and R. Harms, Holt, Rinehart and Winston, New York) pp. 1–88.

FRANSELLA F. and BANNISTER D. (1977). *A Manual for Repertory Grid Technique* (Academic Press, London).

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

GAINES B. R. (1976). 'Foundations of fuzzy reasoning', *International Journal of Man–Machine Studies* **8**(6) (November) 623–668.

GAINES B. R. (1979). 'Sequential fuzzy system identification', *Fuzzy Sets and Systems* **2**(1) (January) 15–24.

GAINES B. R. (1983). 'Precise past—fuzzy future', *International Journal of Man–Machine Studies* **19**(1) (July) 117–134.

GAINES B. R. (1984). 'Perspectives on fifth generation computing', *Oxford Surveys in Information Technology* **1** 1–53.

GAINES B. R. and SHAW M. L. G. (1981a). 'A programme for the development of a systems methodology of knowledge and action', *General Systems Research and Design: Precursors and Futures* (ed. W. J. Reckmeyer, Society for General Systems Research, Louisville, Kentucky) pp. 255–264.

GAINES B. R. and SHAW M. L. G. (1981b). 'New directions in the analysis and interactive elicitation of personal construct systems', *Recent Advances in Personal Construct Technology* (ed. M. L. G. Shaw, Academic Press, London) pp. 147–182.

GAINES B. R. and SHAW M. L. G. (1985). 'Hierarchies of Distinctions as generators of system theories', *Proceedings of the Society for General Systems Research International Conference* (ed. A. W. Smith, Intersystems, California) pp. 559–566.

GAINES, B. R. and SHAW M. L. G. (1986a). 'Induction of inference rules for expert system', *Fuzzy Sets and Systems* **18**(3) (April) 315–328.

GAINES B. R. and SHAW M. L. G. (1986b). 'Knowledge engineering for an FMS advisory system', *Proceedings of 2nd International Conference on Simulation in Manufacturing*. IFS, Bedford, U.K., pp. 51–61.

GAINES B. R. and SHAW M. L. G. (1986c). 'A systemic analysis of human–computer interaction in complex systems', *International Journal of Man–Machine Studies* **25**(1) (January) 1–27.

GAMMACK J. G. and YOUNG R. M. (1985). 'Psychological techniques for eliciting expert knowledge', *Research and Development in Expert Systems* (ed. M. Bramer, Cambridge University Press, Cambridge) pp. 105–116.

GEVARTER W. B. (1983). 'Expert systems: limited but powerful', *IEEE Spectrum* **18** 39–45.

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

HAYES-ROTH F., WATERMAN D. A. and LENAT D. B. (eds.) (1983). *Building Expert Systems* (Addison-Wesley, Reading, Mass.).

HINKLE D. N. (1965). 'The change of personal constructs from the Viewpoint of a Theory of Implications', *PhD Thesis* (Ohio State University).

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

KRANTZ D. H., LUCE R. D., SUPPES P. and TVERSKY A. (1971). *Foundations of Measurement* Vol. 1 (Academic Press, New York).

LANDFIELD A. (1976). 'A personal construct approach to suicidal behaviour', *Dimensions of Intrapersonal Space* Vol 1. (ed. P. Slater, John Wiley, London) pp. 93–107.

LENAT D. B., PRAKASH M. and SHEPHERD M. (1986). 'CYC: using common sense knowledge to overcome brittleness and knowledge acquisition bottlenecks', *AI Magazine* **6**(4) 65–85.

LINSTONE H. A. and TUROFF M. (eds.) (1975). *The Delphi Method: Technique and Applications* (Addison-Wesley, Reading, Mass.) pp. 93–107.

MANCUSO J. R. and ADAMS-WEBBER J. R. (eds.) (1982). *The Construing Person* (Praeger, New York).

MICHALSKI R. S. and WINSTON P. H. (1986). 'Variable precision logic', *Artificial Intelligence* (to appear).

MICHIE D. (ed.) (1979). *Expert Systems in the Micro Electronic Age* (Edinburgh University Press, Edinburgh).

MOTO-OKA T. (ed.) (1982) *Fifth Generation Computer Systems* (North-Holland, Amsterdam).

NISBETT R. E. and WILSON T. D. (1977). 'Telling more than we can know; verbal reports on mental processes', *Psychological Review* **84** 231–259.

OSGOOD C. E., SUCI G. J. and TANNENBAUM P. H. (1957). *The Measurement of Meaning* (University of Chicago Press, Urbana, Ill.).

PASK G. (1976). *Conversation Theory* (Elsevier, Amsterdam).

PASK G. (1980). 'Developments in conversation theory—Part I', *International Journal of Man–Machine Studies* **13**(4) (November) 357–411.

POPE M. L. and SHAW M. L. G. (1981). 'Personal construct psychology in education and learning', *Recent Advances in Personal Construct Technology* (ed. M. L. G. Shaw, Academic Press, London) pp. 105–114.

POPPER K. R (1968). 'Epistemology without a knowing subject', *Logic, Methodology and Philosophy of Science III* (ed. B. Van Rootselaar North-Holland, Amsterdam) pp. 333–373.

REITMAN W. (ed.) (1984). *Artificial Intelligence Applications for Business* (Ablex, Norwood, N.J.).

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

ROY J. (1986). 'Expert systems in Nexpert', *Mac Tutor* **2**(2) (February) 48–51.

SCHANK R. C. (1975). *Conceptual Information Processing* (North-Holland, New York).

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

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

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

SHAW M. L. G. (1984). 'Interactive knowledge elicitation', *Proceedings of CIPS Session 84* (Canadian Information Processing Society, Calgary) pp. 202–208.

SHAW M. L. G. (1985). 'Communities of Knowledge', *Anticipating Personal Construct Psychology* (eds. F. Epting and A. W. Landfield, University of Nebraska Press, Lincoln) pp. 25–35.

SHAW M. L. G. (1986a). 'Epistemological hierarchies in knowledge and expert systems', *Metamodels and Metasystems* (ed. J. P. Van Gigch, Abacus Press, Cambridge, MA).

SHAW M. L. G. (1986b). 'PCS: a knowledge-based interactive system for group problem solving', *Proceedings of 1986 International Conference on Systems, Man and Cybernetics* (in press). Atlanta, (October).

SHAW M. L. G. and GAINES B. R. (1980). 'Fuzzy semantics for personal construing', *Systems Science and Science* (Society for General Systems Research, Kentucky) pp. 146–154.

SHAW M. L. G. and GAINES B. R. (1981a). 'The personal scientist in the community of science', *General Systems Research and Design: Precursors and Futures* (ed. W. J. Reckmeyer, Society for General Systems Research, Kentucky) pp. 59–68.

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

SHAW M. L. G. and GAINES B. R. (1981c). 'Exploring personal semantic space', *Empirical Semantics II* (ed. B. Rieger, Studienverlag Brockmeyer, Bochum, West Germany) pp. 712–791.

SHAW M. L. G. and GAINES B. R. (1982). 'Eliciting entailment', *Progress in Cybernetics and Systems Research* Vol. IX (eds. R. Trappl, L. Ricciardi, and G. Pask, Hemisphere, Washington) pp. 425–435.

SHAW M. L. G. and GAINES B. R. (1983a). 'A computer aid to knowledge engineering', *Proceedings of British Computer Society Conference on Expert Systems*, British Computer Society, London, pp. 263–271.

SHAW M. L. G. and GAINES B. R. (1983b). 'Eliciting the real problem', *International Working Conference on Model Realism* (ed. H. Wedde, Pergamon Press, Oxford) pp. 100–111.

SHAW M. L. G. and GAINES B. R. (1984). 'Deriving the constructs underlying decision', *Fuzzy Sets Decision Analysis. TIMS Series Studies in the Management Sciences* (eds. H. J. Zimmerman, B. R. Gaines, and L. A. Zadeh, North-Holland, Amsterdam).

SHAW M. L. G. and GAINES B. R. (1986). 'A framework for knowl-

edge-based systems unifying expert systems and simulation', *Intelligent Simulation Environments* (eds. P. A. Luker and H. H. Adelsberger, Society for Computer Simulation, La Jolla, Calif.) pp. 38–43.

SHAW M. L. G. and MCKNIGHT C. (1981). 'ARGUS: a program to explore intra-personal personalities', *Recent Advances in Personal Construct Technology* (ed. M. L. G. Shaw, Academic Press, London) pp. 125–134.

SHEPHERD E. and WATSON J. P. (eds.) (1982). *Personal Meanings* (John Wiley, London).

SLATER P. (ed.) (1976). *Dimensions of Intrapersonal Space* Vol. 1 (John Wiley, London).

SLATER P. (ed.) (1977). *Dimensions of Intrapersonal Space* Vol. 2 (John Wiley, London).

SOWA J. F. (1984). *Conceptual Structures: Information Processing in Mind and Machine* (Addison-Wesley, Reading, Mass.).

WAHL D. (1986). 'An application of declarative modeling to aircraft fault isolation and diagnosis', *Intelligent Simulation Environments* (eds. P. A. Luker and H. H. Adelsberger, Society for Computer Simulation, La Jolla, Calif.).

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

ZADEH L. A. (1972). 'Fuzzy languages and their relation to human intelligence', *Proceedings of International Conference on Man and Computer* (Karger, Basel).