Shaw, M. L. G. and Gaines, B. R. (1987). An interactive knowledge-elicitation technique using personal construct technology. In Knowledge Elicitation for Expert Systems: A Practical Handbook Kidd, A. L., ed, pp. 109-136. Plenum Press.

6

An Interactive Knowledge-Elicitation Technique Using Personal Construct Technology

MILDRED L. G. SHAW and BRIAN R. GAINES

1. KNOWLEDGE ENGINEERING

The initial success of expert system (ES) developments (Gevarter, 1983; Michie, 1979; Reitman, 1984) and the development of a number of reasonably domain-independent software support systems for the encoding and application of knowledge (Hayes-Roth *et al.*, 1983) have opened up the possibility of widespread usage of ESs. In particular, the Japanese Fifth Generation Computing System (FGCS) development program (Gaines, 1984; Moto-oka, 1982) 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 be able to apply this knowledge to the elicitation of expertise in specific domains.

MILDRED L. G. SHAW and BRIAN R. GAINES • Department of Computer Science, University of Calgary, Calgary, Alberta, Canada T2N 1N4.

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 *et al.*, 1983).

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 defenses 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 defenses, including those resulting from automization of skilled behavior. Welbank (1983) has reviewed many of these techniques in the context of knowledge engineering.

This chapter is concerned with interactive computer aids to the knowledge engineer based on Kelly's (1955) Personal Construct Psychology (PCP), which uses a model of the person that has strong systemic and psychological foundations (Gaines and Shaw, 1981a; Mancuso and Adams-Webber, 1982; Shaw and Gaines, 1981a, 1981c, 1985). PCP 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 described (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).

2. PERSONAL CONSTRUCT PSYCHOLOGY

Kelly developed a systemic theory of human cognition based on the single primitive of a *construct*, or dichotomous distinction. For an individual, constructs are "transparent templets which he creates and then attempts to fit over the realities of which the world is composed" (Kelly, 1955). He proposes that all of human activity can be seen as a process of anticipating the future by construing the replication of events: "Constructs are used for predictions of things to come, and the world keeps rolling on and revealing these predictions to be either correct or misleading. This fact provides a basis for the revision of constructs and, eventually, of whole construct systems" (Kelly, 1955). Hence, his psychological model of man is strongly epistemological and concerned with the way in which man models his experience and uses this model to anticipate the future. The anticipation may be passive, as in prediction, or active, as in action.

Kelly developed his theory in the context of clinical psychology and hence was concerned to have techniques that 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 ESs 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 the present chapter describes their application to knowledge engineering for ESs.

Kelly's PCP is important because it develops a complete psychology of both the normal and the 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 chapter concentrates on the repertory grid as a technique for eliciting information from an expert.

3. 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 that are used to define the area of the topic, and they can be concrete or abstract entities. For example, in the context of interpersonal relations, the elements might be people; in attitudes to life, they might be critical events; in job change, they might be careers; in expertise about metal joining, they might be types of rivet; in expertise about medical diagnosis, they might be symptoms. Before choosing the set of elements, the user must think carefully about the area of the topic and relate the elements to his purpose. The elements should be of the same type and level of complexity, and should span the topic as fully as possible. It is usual to start with about 6 to 12 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. 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 construct must be central 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 that 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 labeling of the poles need not be a publically agreed meaning in the outside world, but only a memory aid to the thinking process. The mapping of the elements onto the constructs produces the two-dimensional grid of relationships.

3.1. Eliciting Constructs

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 lowest number that 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 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 *multiauthored*, 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 multiauthored or (2) being single-authored, as shown in Figure 1. This also shows the significance of the term *personal* in "personal con-

```
        1       1  2  3  4  5  6  7       2
                ********************

multiauthored   1  2  1  2  1  1  1  single-authored

                ********************
                *  *  *  *  *  *  Building Expert Systems
                *  *  *  *  *  Logic Programming
                *  *  *  *  Knowledge-Based Systems in AI
                *  *  *  Winston & Horn LISP
                *  *  AI Applications for Business
                *  Winston's AI
                Handbook of AI
```

FIGURE 1. A repertory grid with a single construct.

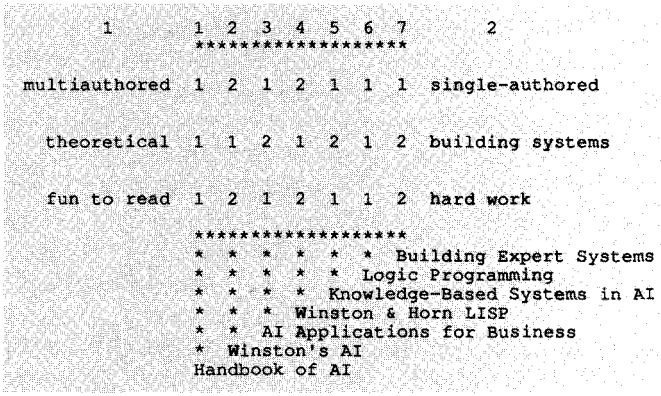

FIGURE 2. A repertory grid with three constructs.

struct'' since it would not obviously be a publicly agreed description that Winston and Horn's LISP is single-authored whereas the Davis and Lenat book is multiauthored. In this case it reflects a single concerted effort as opposed to more than one topic.

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.

A scale allowing more distinctions than the pair 1 and 2 may be

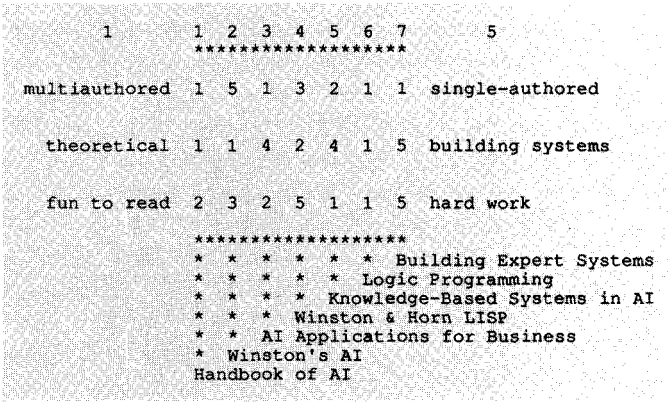

FIGURE 3. A repertory grid with three constructs rated on a 5-point scale.

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 and Horn LISP* are the most *hard work*.

Note how the use of a multipoint 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 (Gaines and Shaw, 1981b; Shaw and Gaines, 1980).

4. TECHNIQUES FOR REPERTORY GRID ELICITATION AND ANALYSIS

In the past, the elicitation of repertory grids has been a task requiring a skilled psychologist who can draw out a person's constructs while 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 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. 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.

An example of PEGASUS in action is given later in this paper. ARGUS elicits a set of grids simultaneously from one person holding several roles or points of view. First, 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 role name in turn, and at each stage another construct that 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.

4.1. 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 analyzed 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 others. 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), and if someone is using repertory grids for the first time, it is worth exploring the different techniques on their own data.

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.

4.2. Analysis of a Single Grid

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 another of the poles of the construct. From this point of view each construct becomes represented as a point in a multidimensional 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 that 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 that are not equivalent, we may analyze 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 on 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. 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 represent the original grid reorganized by the "neighborness" 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 or constructs, summing them, then mapping the result into 0 to 100% similarity for the elements, and −100 to 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 resorted original grid data (Shaw, 1980).

The PRINGRID algorithm is another distance-based cluster analysis using standard principal component 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. However, both INGRID and FOCUS are less applicable to knowledge engineering for expert systems where a logical analysis is more relevant.

The alternative to distance-based methods of grid analysis is a logical analysis, looking at constructs as predicates applying to elements. The ENTAIL algorithm 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. This is discussed in more detail with the example given in a later section.

4.3. Analysis of a Pair of Grids

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 may 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 methods that can be used to compare two grids. 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 highlights 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.

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 that show agreement or understanding in the two grids are identified and recorded.

4.4. Analysis of a Group of Grids

The SOCIOGRIDS algorithm analyzes a set of repertory grids elicited from a group who share the same elements. It compares every pair of grids 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 that 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 that has been used is listed in order of how highly matched it is across all grids, and indicating where it originated.

This technique 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 to promote common values and understanding in a group.

5. SOFT SYSTEMS ANALYSIS

Personal construct psychology provides cognitive foundations for knowledge elicitation, and repertory grid techniques provide a corre-

sponding methodology. However, neither in itself gives a framework for carrying out knowledge engineering. This section considers such a framework, Checkland's (1981) *soft systems analysis*, which is consistent with PCP and can be operationalized using repertory grid elicitation and analysis.

5.1. The Significance of Different Perspectives

Kelly emphasizes the *dynamic* and *personal* nature of the construct system. It is subject to change according to feedback from failure to anticipate events. It is individualistic in the constructs used, in the vocabulary used to name the construct poles, in the relations between constructs in the hierarchy, and in those constructs most likely to change when some change takes place.

Knowledge engineering for ESs may be seen as a process of making overt the construct system of the expert. 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 construction system effectively while 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.

This essentially *soft* aspect of the systems analysis of expert knowledge is contrary to the normal expectation that there is a "right" answer, valid for all conditions. Checkland (1981) has termed this a problem of soft systems analysis and emphasized the importance of the analyst's not bringing his own preconceptions to the problem domain. He has also shown the significance to taking into account a number of different perspectives in knowledge engineering.

5.2. Techniques of Soft Systems Analysis

Checkland (1981) has developed a number of techniques for soft systems analysis, starting with a structuring of the stages involved as shown in Figure 4.

The importance of Stage 1 is to get back to it from Stage 2 or beyond. We should not inject preconceptions about solutions into the problem definition and need to treat any situation as unstructured no matter

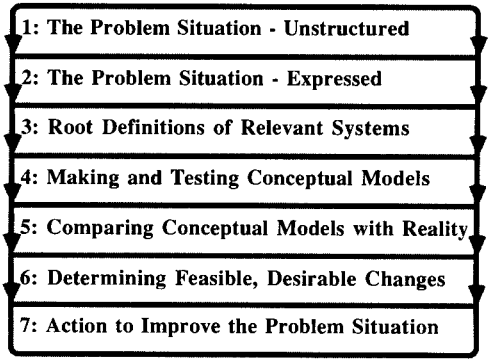

FIGURE 4. Stages in soft systems methodology.

what structure we are given. Our framework for analyzing the expert's knowledge may be such as to prevent that analysis. The expression at Stage 2 is still unstructured, and it is Checkland's study of Stage 3, *root definitions* of relevant systems, that provides the formal framework for a structure.

Checkland's approach is particularly interesting because it emphasizes the pluralism of systems analysis, what Kelly has termed *constructive alternativism,* that it is important to examine the problem from a number of viewpoints. These may be seen as forming a nested set to which Checkland gives the mnemonic CATWOE, as shown in Figure 5.

The system is defined through a Transformation carried out by people who are the Actors within it; it affects beneficially or adversely

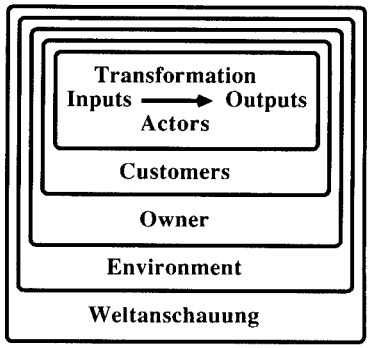

FIGURE 5. Root definitions of relevant systems.

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

Applying the analysis to an ES for oil exploration, we consider a *transformation* to which the *inputs* are survey data and the *outputs* are decisions to drill. The *actors* are experts concerned with making these decisions. Their *customers* are 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 government that taxes 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, and so on. These are the considerations that knowledge engineering has to take into account.

The emphasis on differing viewpoints here is fundamental. PCP is often seen as concerned with the different perspectives of 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. In fact, the theory underlying PCP 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, 1983). These considerations are those of Pask's *conversation theory* (Pask, 1975), particularly the notion of expertise distributed across a multiperson *P-Individual*, and it is notable that this theory is being used as a foundation for the development of ESs (Coombs and Alty, 1984).

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 PCP approach can automate the analysis in major part.

6. PLANET: A COMPUTER-BASED KNOWLEDGE-ENGINEERING SYSTEM

The program system that we have developed for doing soft systems analysis from a variety of viewpoints and then comparing and contrasting the systems elicited is called PLANET (Personal Learning, Analysis, Negotiation, and Elicitation Techniques; Shaw, 1982). It runs as a menu-driven suite of interactive programs on a variety of computers, and, while preliminary versions were developed on the PDP10, there proved to be no difficulty in issuing PLANET for the Apple II, enabling its widespread use in field studies. The programs are all concerned with repertory grid elicitation and analysis, and they divide naturally into the following four groups.

Construct Elicitation. PEGASUS is a highly interactive program that elicits constructs through a conversational dialogue 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. 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.

Single Construct System Analysis. 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. PRINGRID is a nonhierarchical cluster analysis based on principal components that can be used to gauge the major dimensions along which distinctions are being made. ENTAIL is a logical analysis of the construct system taking the expert's distinctions to be fuzzy predicates (Gaines, 1976; Zadeh, 1972). The entailment structure it generates can be used as a decision tree expressing the relationship between an expert's data and his conclusions. It can also provide rules for input to an ES shell (Boose, 1985; Gaines and Shaw, 1985).

Multiple Construct System Analysis. MINUS compares two grids based on the same elements and constructs to highlight the differences between them. CORE provides an interactive comparison of such grids to determine where there are differences and extract invariancies, e.g., over time. SOCIOGRIDS analyzes the total construct structure generated by a number of people construing the same elements from their individual viewpoints. It produces *socionets* showing the relationships be-

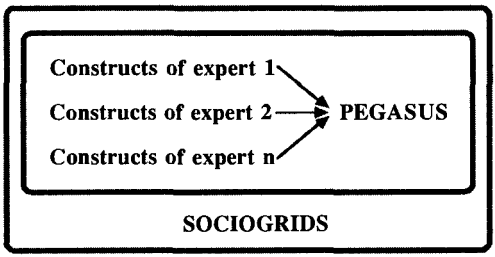

FIGURE 6. PLANET automating integration of multiple viewpoints.

tween individual construct systems within the group and a *mode grid* reflecting the constructs shared across the group.

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

PLANET provides a set of tools for the knowledge engineer enabling him to operate at Stage 3 of Checkland's soft systems analysis semi-automatically. As shown in Figure 6, 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 could have very different perspectives, such as "actor," "customer," and "owner". SOCIOGRIDS enables the consensual structure to be derived from the PEGASUS results, and ENTAIL enables this to be turned into a decision tree, or rules to prime an expert system shell, when appropriate.

The next section shows a sample of the PEGASUS dialogue in eliciting expertise from a manager in the appraisal of his subordinates (Shaw, 1980).

7. PEGASUS IN ACTION

PEGASUS commences its interaction by asking for the purpose of the elicitation and then for a set of elements relevant to this purpose, as shown in Figure 7. It is important to establish an initial context such that the expert selects a set of elements that are of the same type and relevant to the purpose of the elicitation. Note that in the dialogue the computer output is shown in normal face and the human input in italics.

One technique used by PEGASUS to elicit constructs is that of triadic comparison, as shown in Figure 8. A combination of three elements is

```
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 7. PEGASUS initial element selection.

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 construct that differentiates between two highly matched elements, as shown in Figure 9.

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 10.

Through a combination of triadic comparison, element matching, and construct matching, as shown in the flow chart of Figure 11, a PEGASUS interaction elicits from the user his system of constructs relating to the topic being studied. Figure 12 is the final grid of the manager.

```
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 8. PEGASUS triadic construct elicitation.

```
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 9. PEGASUS construct elicitation by breaking element 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 10. PEGASUS element elicitation by breaking construct match.

8. ENTAIL IN ACTION

ENTAIL (Gaines and Shaw, 1981b, 1985; Shaw and Gaines, 1982) gives a dependency analysis of the grid by listing logical entailments consistent with the data. The dependency analysis is also being extended to be interactive so that it may lead to the elicitation of further elements if the system suggests dependencies with which the expert disagrees, or to the modification of ratings if the expert suggests dependencies that are not possible according to the data he has entered. The overall objective is to allow the expert to explore the dimensions of his conceptual structure, testing it against specific cases, and discussing its implications, until he is comfortable that the representation elicited reasonably represents his knowledge.

Figure 13 shows the ENTAIL analysis of the grid of Figure 12. The entailments are shown with three values, the first a truth value from 0 to 100 (e.g., L 9 → L 13 truth 100), the second the probability of the hypothesis being true, and the third the uncertainty reduction generated by asserting this entailment. All three are in the range from 0 to 100. The uncertainty reduction measures the information carried by the assertion and is used to ensure that trivial entailments consistent with the data are pruned (Gaines and Shaw, 1985).

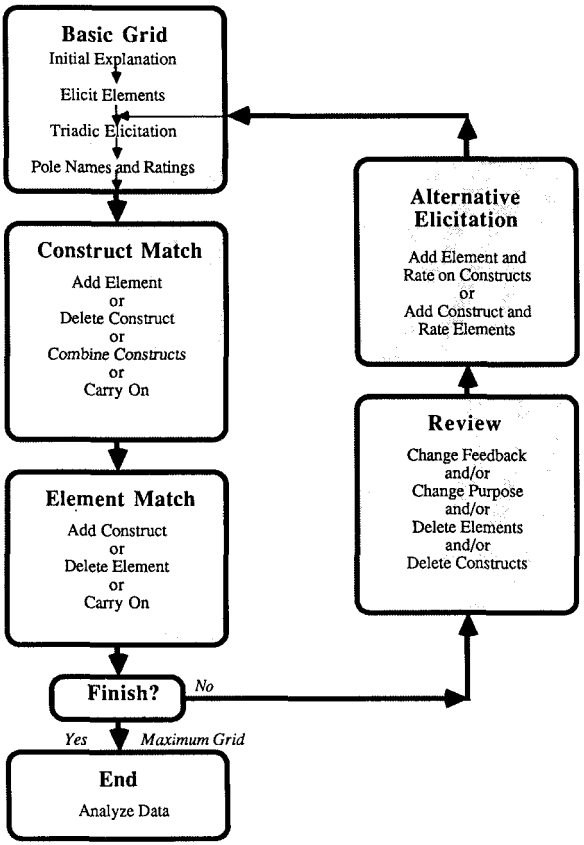

FIGURE 11. Flow of control in PEGASUS elicitation.

In our initial analysis of entailment we introduced the notion of entailments being true "usually" and interpreted this as "in all except one instance." The original version of ENTAIL generates entailments that are usually true, and these have proved useful with some data where there is a common pattern upset only by an isolated case. The concept of "usually" can be generalized and formalized within the framework already established, and it leads to an interesting two-dimensional analysis of fuzzy data.

We quantify the notion of a hypothesis being true "usually" by ascribing a probability to the hypothesis measured by the relative frequency with which it is true. For example, using the data of L9 and L13, L9 → L13 has truth value 1 with probability 1; it is always absolutely

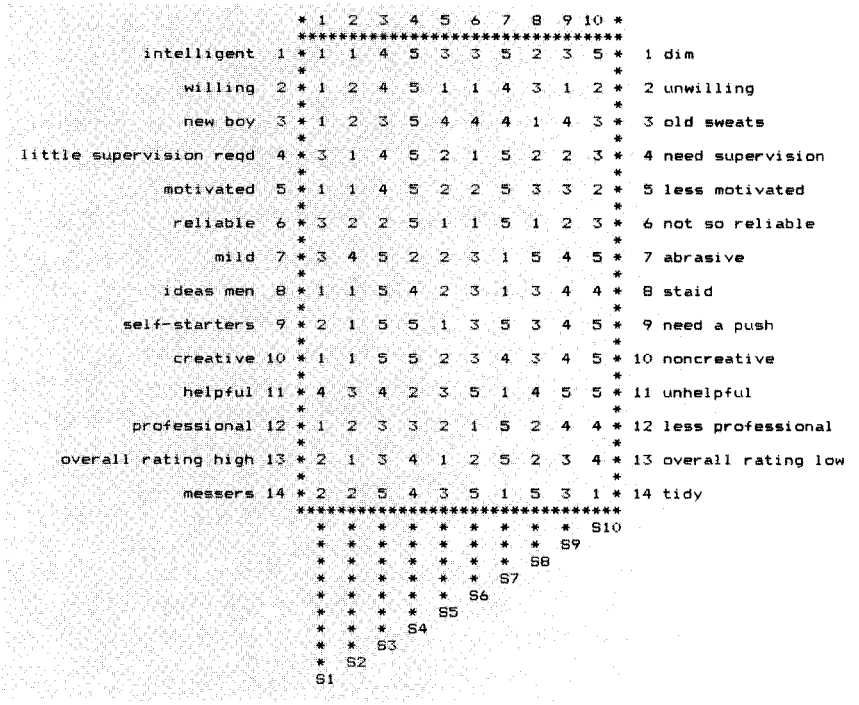

FIGURE 12. PEGASUS repertory grid elicited on staff appraisal.

true. However, L13 → L9 is more interesting because it has truth value 0.5 with probability 1, truth value 0.75 with probability 0.9 (fails in 1 case out of 10), and truth value 1.0 with probability 0.4 (fails in 6 cases out of 10).

This type of trade-off distribution between the truth value of an assertion and its probability of being true occurs in inductive inference with fuzzy data, for example, in inferring structure from fuzzy sequential data (Gaines, 1979). We can evaluate the possible trade-offs by using the uncertainty reduction analysis based on expected surprise already described by Gaines and Shaw (1985b).

ENTAIL checks for equivalence separately from the check for one-way entailment. This is because the uncertainty reduction from the hypothesis of equivalence, $x \equiv y$, cannot be inferred from the separate uncertainty reductions of the logically equivalent pair of hypotheses, $x \rightarrow y$ and $y \rightarrow x$, although it will be at least the sum of these. In order to

rank-order the hypotheses, the uncertainty reduction for one-sided entailment is doubled for comparison with that of two-sided equivalence.

Note how a hypothesis of high truth value, such as R13 → R9, can give less information reduction than one of lower truth value, such as L1 → L9. Note also how a one-sided hypothesis, such as L1 → L10, may give greater uncertainty reduction than its two-sided equivalent, L1 ≡ L10, and hence have a higher ranking. Similarly, the two one-sided hypotheses, such as L9 → L10 and L10 → L9, can also occur in preference to the two-sided hypotheses, L9 ≡ L10. The list of hypotheses in Figure 13 has been terminated when the uncertainty reduction falls below 35% for a two-sided hypothesis. It can be truncated at any level and the entailments plotted in graphical form. However, since approximate inference is involved, the graph will not necessarily be transitive, and it

```
Entail    Truth Prob. Information (Cutoff 35)

L 9 → L13   100   100   59 self-starters → overall rating high
L 1 → L 9    75   100   58 intelligent → self-starters
L10 → L 8   100   100   57 creative → ideas men
R 9 → R 1    75   100   56 need a push → dim
L 1 → L10    75   100   52 intelligent → creative
R 8 → R10   100   100   52 staid → noncreative
R13 → R 6    75   100   52 overall rating low → not so reliable
L10 → L 9    75   100   51 creative → self-starters
R10 → R 1    75   100   49 noncreative → dim
L 9 → L10    75   100   47 self starters → creative
L10 ≡ L 1    75   100   47 creative ≡ intelligent
R 1 ≡ R10    75   100   47 dim ≡ noncreative
R 9 ≡ R10    75   100   47 need a push ≡ noncreative
R13 → R 9   100   100   46 overall rating low → need a push
R 4 → R13    75   100   44 need supervision → overall rating low
R 5 → R 4    75   100   44 less motivated → need supervision
R 5 → R13    75   100   44 less motivated → overall rating low
L 1 → L 3    75   100   42 intelligent → new boy
L 6 → L13    75   100   42 reliable → overall rating high
L 4 ≡ L13    75   100   39 little supervision reqd ≡ overall rating high
L13 ≡ L 6    75   100   39 overall rating high ≡ reliable
R 1 → R 6    50   100   39 dim → not so reliable
R 6 ≡ R13    75   100   39 not so reliable ≡ overall rating low
R13 ≡ R 4    75   100   39 overall rating low ≡ need supervision
R12 → R13    75   100   38 less professional → overall rating low
R13 → R12    75   100   38 overall rating low → less professional
L 8 ≡ L10   100    70   37 ideas men ≡ creative
R 9 → R 4    50   100   37 need a push → need supervision
R 9 → R12    50   100   37 need a push → less professional
R 9 → R13    50   100   37 need a push → overall rating low
R10 ≡ R 8   100    70   37 noncreative ≡ staid
L 4 → L 5    75   100   36 little supervision reqd → motivated
L 4 → L 9    50   100   36 little supervision reqd → self-starters
L12 → L 9    50   100   36 professional → self-starters
L12 ≡ L13    75   100   36 professional ≡ overall rating high
L13 → L 5    75   100   36 overall rating high → motivated
L13 → L 9    50   100   36 overall rating high → self-starters
R 2 → R 4    75   100   35 unwilling → need supervision
R 2 → R 5    75   100   35 unwilling → less motivated
R 2 → R13    75   100   35 unwilling → overall rating low
R 3 → R 1    75   100   35 old sweats → dim
```

FIGURE 13. Entailment analysis of repertory grid on staff appraisal.

is important to preserve the intransitivity. It avoids the paradoxes of approximate reasoning, such as that of iterated approximate inference—for example, that a person with one more hair than a bald man is still bald (Gaines, 1976, 1983).

The entailments derived by ENTAIL can be converted directly to rules that may be loaded into a standard ES shell, as Boose (1985) has shown. 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. The rules form a decision tree expressing the relationship between an expert's data and his conclusions, both personally expressed. For example, the data of Figure 12 may be regarded as those 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 13 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."

9. VALIDATION

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 gain some familiarity with the construing and rating process. It is essential that the expert become familiar with the kinds of questions being asked, and finish his grid with a representative set of elements all at the same level of abstraction, and a representative set of constructs that 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, 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 analyzing 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 analyzed 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 the 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 that subsumes the relevant constructions of experience of experts in that domain. Figure 14 shows a socionet produced by SOCIOGRIDS in which an arrow from one person to another shows that the first person "outconstrues" the other. The results are mainly what one

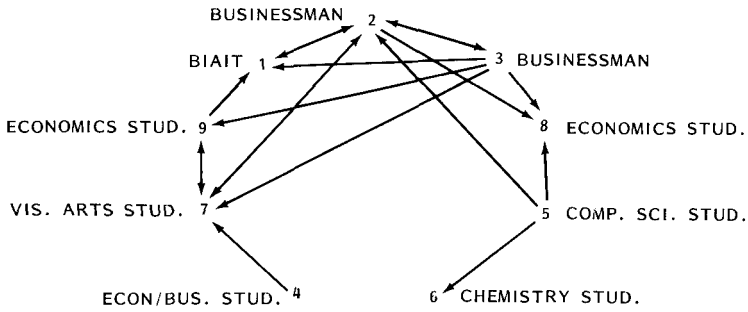

FIGURE 14. Socionet produced by SOCIOGRIDS from BIAIT validation experiment.

would expect, with the businessmen and the economics students closely related to BIAIT. The only disconcerting result is the position of the Visual Arts student, who is clearly a central part of the cluster. Enquiry into the background of this student revealed that she had only recently become a student, having spent the previous 25 years as an administrative secretary in an insurance organization.

The capability of PLANET and related systems to generate useful expert systems rapidly is clearly another important dimension of validation. Boose (1985) discusses the strengths and weaknesses of his Expertise Transfer System (ETS), which is based on the PLANET system but operates on a Xerox 1108 LISP machine. At Boeing Computer Services he has produced a wide range of expert systems by this methodology. He says that it is better suited for analysis than for synthesis problems—for example, debugging, diagnosis, interpretation, and classification rather than design and planning; that it is difficult to apply to "deep" causal knowledge or strategic knowledge. On the other hand he says: "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" (Boose, 1985).

Figure 15 shows the screen of the LISP machine running (ETS) at the stage of generating conclusion rules from the grid ratings. ETS has been used to supply the initial knowledge base to TEIRESIAS (Davis and Lenat, 1982), which helps the expert and knowledge engineer to debug and refine an existing knowledge base since TEIRESIAS is not capable of eliciting such information on its own.

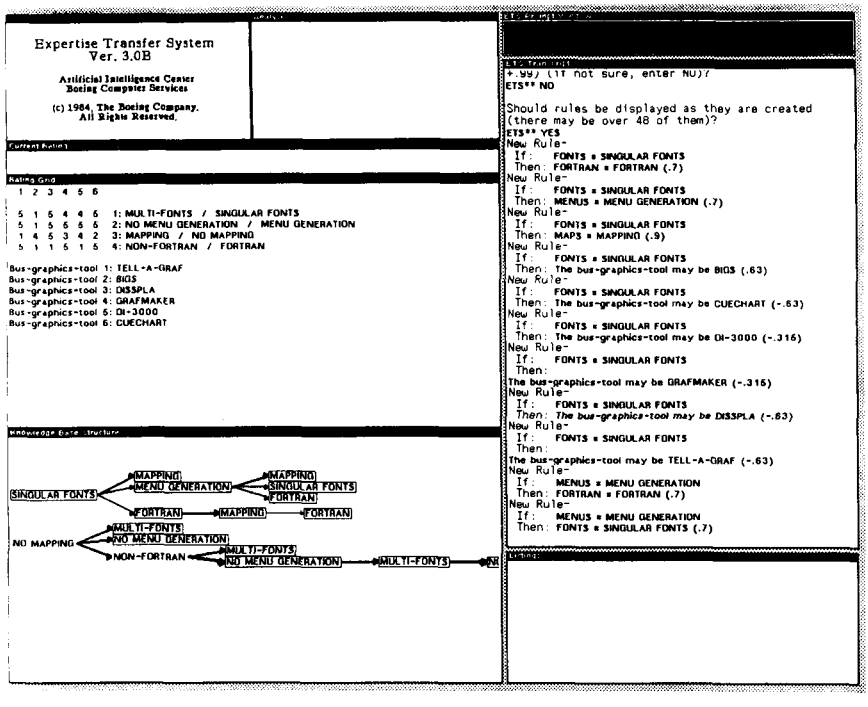

FIGURE 15. ETS elicitation on a LISP machine (Boose, 1985).

10. GUIDELINES SUMMARY

1. Kelly's personal construct psychology identifies the system of construction being used by an expert as fundamental to his thought processes.
2. Checkland's soft system analysis structures the application of these concepts to the ill-defined problems typical of expert system applications.
3. These methodologies can be used to structure interviews eliciting knowledge from experts.
4. However, manual elicitation is tedious, and on-line analysis and feedback lead to improved elicitation.
5. Automatic computer-based techniques have been developed and prove very effective in knowledge elicitation for expert systems.

6. PEGASUS is a program that elicits from an expert the fundamental distinctions, or constructs, he is using in applying his expertise.
7. ENTAIL analyzes the logical dependencies between constructs and can be used to develop rules for an ES shell.
8. SOCIOGRIDS analyzes the relations between experts and can be used to develop consensual construct systems.
9. The programs and methodology have been widely used since 1978 not only for expert system development but also for a wide variety of knowledge-elicitation studies in psychology, education, and management.
10. Studies have been carried out to validate the methodology experimentally using the BIAIT business record-keeping constructs as a standard.
11. Boeing Computer Services has implemented a related system, ETS, on a LISP machine and produced a wide range of expert systems in an industrial environment.
12. The methodology is better suited for analysis than for synthesis problems—for example, debugging, diagnosis, interpretation, and classification rather than design and planning.
13. It is currently difficult to apply to "deep" causal knowledge or strategic knowledge.
14. In situations where it is applicable, the methodology can save at least 2 months of project calendar time, and decision support systems can be prototyped in a few hours.

ACKNOWLEDGMENTS. Financial assistance for this work has been made available by the National Sciences and Engineering Research Council of Canada. We are grateful to Dr. John Boose and Boeing Computer Services for permission to use Figure 15.

11. REFERENCES

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

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

Boose, J. H. (1985). A knowledge acquisition program for expert systems based on personal construct psychology. *International Journal of Man-Machine Studies,* 23(5), 495–525.

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

Checkland, P. (1981). *Systems thinking, systems practice.* Chichester, England: Wiley.

Coombs, M., and Alty, J. (1984). Expert systems: An alternative paradigm. *International Journal of Man-Machine Studies, 20*(1), 21–43.

Davis, R., and Lenat, D. B. (1982). *Knowledge-based systems in artificial intelligence.* New York: McGraw-Hill.

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., and 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). Foundations of fuzzy reasoning. *International Journal of Man-Machine Studies, 8*(6), 623–668.

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

Gaines, B. R. (1983). Precise past—Fuzzy future. *International Journal of Man-Machine Studies, 19*(1), 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. In W. J. Reckmeyer (Ed.), *General systems research and design: Precursors and futures.* (pp. 255–264). Louisville, Ky.: Society for General Systems Research.

Gaines, B. R., and Shaw, M. L. G. (1981b). New directions in the analysis and interactive elicitation of personal construct systems. In M. L. G. Shaw (Ed.), *Recent advances in personal construct technology* (pp. 147–182). London: Academic Press.

Gaines, B. R., and Shaw, M. L. G. (1986). Induction of inference rules for expert systems. *Fuzzy Sets and Systems, 18,* 315–328.

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), 1–47.

Hayes-Roth, F., Waterman, D. A., and Lenat, D. B. (Eds.). (1983). *Building expert systems.* Reading, Mass.: Addison-Wesley.

Kelly, G. A. (1955). *The psychology of personal constructs.* New York: Norton.

Landfield, A. (1976). A personal construct approach to suicidal behaviour. In P. Slater (Ed.), *Dimensions of intrapersonal space* (Vol. 1, pp. 93–107). London: Wiley.

Mancuso, J. R., and Adams-Webber, J. R. (Eds.). (1982). *The construing person.* New York: Praeger.

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.

Pask, G. (1975). *Conversation, cognition and learning.* Amsterdam: Elsevier.

Pope. M. L., and Shaw, M. L. G. (1981). Personal construct psychology in education and learning. In M. L. G. Shaw (Ed.), *Recent advances in personal construct technology* (pp. 105–114). London: Academic Press.

Reitman, W. (Ed.). (1984). *Artificial intelligence applications for business.* Norwood, N.J.: Ablex.

Rogers, C. R. (1967). *On becoming a person: A therapist's view of psychotherapy.* London: Constable.

Shaw, M. L. G. (1980). *On becoming a personal scientist.* London: Academic Press.

Shaw, M. L. G. (Ed.). (1981). *Recent advances in personal construct technology.* 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. (1983). Epistemological hierarchies in knowledge and expert systems. In J. P. van Gigch (Ed.), *Metamodels and metasystems.* Cambridge, Mass.: Abacus Press.

Shaw, M. L. G. (1984, May). Interactive knowledge elicitation. *Proceedings of CIPS SESSION 84* (pp. 202–208). Calgary: Canadian Information Processing Society.

Shaw, M. L. G. (1985). Communities of knowledge. In F. Epting and A. W. Landfield (Eds.), *Anticipating personal construct psychology* (pp. 25–35). Lincoln: University of Nebraska Press.

Shaw, M. L. G., and Gaines, B. R. (1980). Fuzzy semantics for personal construing. *Systems science and science* (pp. 146–154). Louisville, Ky.: Society for General Systems Research.

Shaw, M. L. G., and Gaines, B. R. (1981a). The personal scientist in the community of science. In W. J. Reckmeyer (Ed.), *General systems Research and design: Precursors and futures* (pp. 59–68). Louisville, Ky.: Society for General Systems Research.

Shaw, M. L. G., and 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.

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

Shaw, M. L. G., and Gaines, B. R. (1982). Eliciting entailment. In R. Trappl, L. Ricciardi, and G. Pask (Eds.), *Progress in cybernetics and systems research* (Vol. 9, pp. 425–435). Washington, D.C.: Hemisphere.

Shaw, M. L. G., and Gaines, B. R. (1983a, December). A computer aid to knowledge engineering. *Proceedings of British Computer Society conference on expert systems,* Cambridge, pp. 263–271.

Shaw, M. L. G., and Gaines, B. R. (1983b). Eliciting the real problem. In H. Wedde (Ed.), *International working conference on model realism* (pp. 100–111). Oxford: Pergamon Press.

Shaw, M. L. G., and Gaines, B. R. (1984). Deriving the constructs underlying decision. In H.-J. Zimmerman, B. R. Gaines, and L. A. Zadeh (Eds.), *Fuzzy sets and decision analysis. TIMS series studies in the management sciences.* Amsterdam: North-Holland.

Shaw, M. L. G., and Gaines, B. R. (1985). Some perspectives on knowledge engineering for expert systems and simulation. In *Proceedings of workshop on coupling symbolic and numerical computing in expert systems.* Bellevue, Wash.: Boeing Computer Services.

Shepherd, E., and Watson, J. P. (Eds.). (1982). *Personal meanings.* London: Wiley.

Slater, P. (Ed.). (1976). *Dimensions of intrapersonal space* (Vol. 1). London: Wiley.

Slater, P. (Ed.). (1977). *Dimensions of intrapersonal space* (Vol. 2). London: Wiley.

Sowa, J. F. (1984). *Conceptual structures: Information processing in mind and machine.* Reading, Mass.: Addison-Wesley.

Welbank, M. (1983). *A review of knowledge acquisition techniques for expert systems.* BTRL, Ipswich: Martlesham Consultancy Services.

Zadeh, L. A. (1972). Fuzzy languages and their relation to human intelligence. *Proceedings of the international conference on man and computer.* Basel: Karger.