Interactive Knowledge Elicitation

Mildred L. G. Shaw

Department of Computer Science, York University,
4700 Keele Street, Downsview, Ontario, Canada M3J 1P3

This paper outlines a methodology for knowledge elicitation that has been implemented in a computer program that interacts with an expert to enable him to express the constructs underlying his knowledge. Experiments on validating the system by using it to re-construct the Business Information Analysis and Integration Technique (BIAIT) are reported.

Cet article expose une méthodologie d'extraction de connaissance qui fut implémentee dans un programme informatique qui entre en interaction avec un expert en vue de lui permettre d'exprimer les concepts sous-jacent à ses connaissances. Des experiences visant à valider le système en l'utilisant pour reconstruire le "Business Information Analysis and Integration Technique (BIAIT)" sont raportées.

Knowledge Engineering

The initial success of expert system developments (Michie, 1979) 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 development program (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 expert systems in new domains. We need to understand more about the nature of expertise in itself (Hawkins, 1983) and to able to apply this knowledge to the elicitation of expertise in specific domains.

The underlying problem for knowledge engineering shows up well in the "conceptual diagram" of fifth generation systems with its 3 rings of hardware, software and human applications (Moto-oka, 1982, p.29). One notable feature of this diagram is that the hardware and software rings contain detailed structures, but the human ring is empty. It should contain a psychological architecture of the person, but as Norman notes, in summarizing the current state of *cognitive science*, there are major gaps in the underlying theory (Norman, 1980). These gaps and the empty ring indicate problems for the knowledge engineer and limitations to the applications of expert systems:

> *"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 have developed techniques for making use of verbal interaction, for example through interviewing techniques that attempt to by-pass 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 paper 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 (Mancuso and Adams-Webber, 1982). PCP provides a model of human knowledge representation, acquisition and processing, that can be used to fill in the third ring, and has been made operational through computer programs for interactive knowledge elicitation (Shaw, 1980). 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, PEGASUS (Shaw and Gaines, 1981b; Shaw and Gaines, 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 program 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 are described to validate the use of PEGASUS in this way by re-constructing 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 (1955) 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.*

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."*

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 by-pass 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 Keen, 1981); and to study decision making by individuals and groups in management. I have developed an integrated range of programs, PLANET (Shaw, 1982), which 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 paper describes their application to knowledge engineering for expert systems.

Kelly's PCP is important because it develops a complete psychology of both the normal and abnormal, which has strong systemic foundations (Gaines and Shaw, 1981). In the long term these may be more important to knowledge engineering than the techniques currently based on them. The next section briefly reviews the systems analysis aspects of these foundations.

Epistemological Hierarchy

In 1955 when Kelly developed his theory there was neither the mathematics to formalize it nor the computer technology to operationalize it. He sees man's main characteristic being that he *models* the world (the *personal scientist* (Shaw, 1980)) and that the basis of this model can be revized (*constructive alternativism* (Mancuso and Adams-Webber, 1982)). His explication of man's modeling process is extraordinarily simple (Kelly, 1955); personal constructs are dichotomous, or bipolar, distinctions; constructs themselves are construed and hence form a hierarchy or *construct system*; at the upper levels of the hierarchy there tend to be evaluative constructs that form the basis of appraisal of events, manifesting themselves in emotions and the goals for action; at the lower levels there tend to be ways of classifying experience. This hierarchy lies at the heart of modern approaches to modeling in system theory such as Klir's *epistemological hierarchy* (Klir, 1976; Klir, 1985), shown in a psychological variant (Gaines and Shaw, 1981) in Figure 1.

In this figure the lowest level is one of the *constructs* or the distinctions made in interacting with the world. The next level is one of the *experiences* of events which happen to us, and we make happen, in terms of the distinctions already made. Levels above these are *hypotheses* which are rationalizations of experience, *analogies* between these rationalizations and *transcendencies* which are preconceptions underlying rationality. Interaction with the world is, therefore, mediated through the construct system to produce experience which is modeled through the higher levels and leads to predictions, decisions and actions again mediated through the construct system.

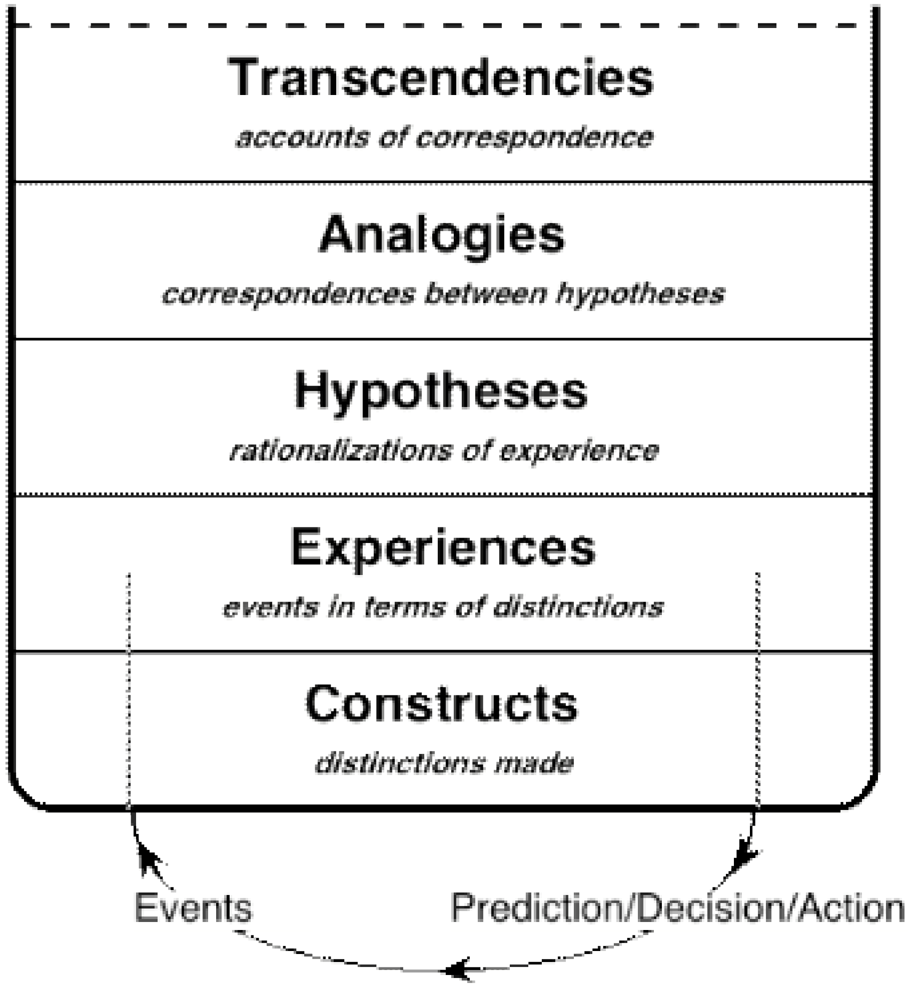

Figure 1 Epistemological hierarchy of a personal scientist

Klir (1985) has used this hierarchy as a framework for a general theory of knowledge acquisition through modeling and Gaines (1977) has shown that the processes represented may be formalized algebraically and computationally. In particular, it may be shown (Gaines and Shaw, 1981) that the calculus of distinctions developed by Brown (1969) is a formal foundation for system theory (Gaines, 1980) leading naturally (Varela, 1975) to the tools of fuzzy reasoning (Zadeh, 1972), and category theory (Goguen, 1974) applied to modeling (Ralescu, 1979) and programming (Goguen, Meseguer and Plaisted, 1983). The *standard uncertainty logic* arising from this analysis encompasses probabilistic and modal logics (Gaines, 1983), and Quinlan (Quinlan, 1983) has shown that the constraint systems of such non-truth-functional systems may be used as a basis for improved inference in expert systems. Thus, the epistemological hierarchy provides a natural framework for knowledge engineering, bridging from the psychological principles of PCP to its systemic and logical foundations.

It should be noted that the hierarchy is strongly idiosynchratic and formalization should not be taken as implying uniformity. 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 this is necessary. Knowledge engineering for expert systems 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 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.

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. The next section is concerned with a framework for systems analysis that can cope with this situation and structure it for practical applications.

Soft Systems Analysis

The difficulty with eliciting the hierarchy of distinctions, or constructs, of an expert is the lack of a conceptual framework that can be expressed and used in a dialog between expert and knowledge engineer. Checkland (1981) has termed this a problem of *soft systems analysis* and emphasized the importance of the analyst not bringing his own preconceptions to the problem domain. He has developed a number of techniques for soft systems analysis, starting with a structuring of the stages involved:

1: The Problem Situation - Unstructured
2: The Problem Situation - Expressed
3: Root Definitions of Relevant Systems
4: Making and Testing Conceptual Models
5: Comparing Conceptual Models with Reality
6: Determining Feasible, Desirable Changes
7: Action to Improve the Problem Situation

Figure 2. Stages in Soft Systems methodology

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 what structure we are given. Our framework for analysing 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. His 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:

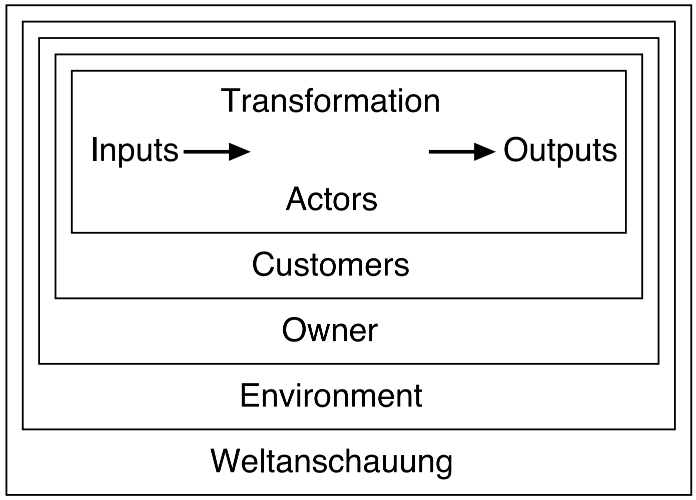

Figure 3. Root Definitions of Relevant Systems

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

Applying the analysis to an expert system 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 which 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 an individual and we are most often concerned with skills that reside in a group or team. In fact the theory underlying PCP, the hierarchy of distinctions underlying modeling, makes no separation necessary between the individual and the group (Shaw and Gaines, 1981c). I 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 and Gaines, 1981a). 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 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).

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 give the analysis these features.

PLANET: A Knowledge Engineering System

The program system that I 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 & Elicitation Techniques) (Shaw, 1982; Gaines and Shaw, 1984). 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 divide naturally into four groups:

Construct elicitation. *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. *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 non-hierarchical 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. The entailment structure it generates can be used as a decision tree expressing the relationship between an expert's data and his conclusions, both personally expressed.

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 or over a group of people. *SOCIOGRIDS* analyses the total construct structure generated by a number of people construing the same elements from their individual viewpoints. It produces *socionets* showing the relationships between 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.

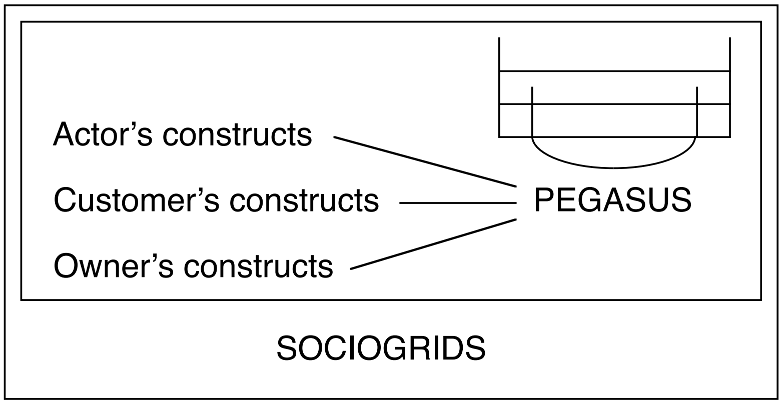

Figure 4 PLANET automating soft system analysis

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. PEGASUS enables the construct systems of all those involved in the application of expertise to be elicited, SOCIOGRIDS enables the consensual structure to be derived from these, and ENTAIL enables this to be turned into a decision tree when appropriate.

The next section is concerned with validating the methodology described applied to knowledge engineering.

BIAIT: A Test Case for Validation

One of the problems of evaluating any methodology for knowledge engineering is to have some 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:

1. Does the supplier bill the customer, or does the customer pay cash?
2. Does the supplier deliver the product at some time in the future, or does the customer take the order with him?
3. Does the supplier keep a profile of the customer, or is every transaction a surprise?
4. Is the price negotiated or fixed?
5. Is the product rented or purchased?
6. Does the supplier keep track of the product after it is sold or not?
7. Is the product made to order or provided from stock?

Sowa notes that these features "seem almost obvious once they are presented, but finding the right features and categories may take months or years of analysis. The proper set of categories provides a structural framework that helps us organize the detailed facts and general principles of a system. Without them, nothing seems to fit, and the resulting system is far too complex or unwieldy."

The BIAIT features are examples of constructs applied to orders placed between a customer and supplier. They provide a useful test case with which to evaluate the personal construct elicitation methodology provided in PLANET. An empirical study is being carried out to determine whether PEGASUS can be used to elicit the BIAIT distinctions from those with some accounting or business knowledge. It is particularly interesting to determine how open such a study may be:

can the knowledge engineer start at Checkland's Stage 1 and operate in a completely unstructured way or does he have to "seed" the solution by a process of selection. We usually assume that even if we cannot comprehend the expertise we can recognize an expert, and hence a key strategy for knowledge engineering is the selection of experts with whom to collaborate. This has strong prima facie validity but it would also be interesting to demonstrate it empirically. To do so through the PLANET analysis would itself be an indication of the validity of the methodology.

Hence, the initial 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.

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. In the dialog the computer output is shown in normal face and the human input in **bold**:

`Type in your purpose for doing this grid?` **`examining business record keeping`**

`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.`

`Element 1?` **`selling newspapers from a stand`**
`Element 2?` **`selling antiques in a market`**
`Element 3?` **`selling custom drapes`**
`Element 4?` **`selling large appliances`**
`Element 5?` **`hiring out a car to a customer`**
`Element 6?` **`providing TV rental services`**
`Element 7?` **`selling car insurance`**
`Element 8?` **`providing electricity for home`**
`Element 9?` **`providing dental services`**
`Element 10?` **`being agent for home purchase`**

One technique used by PEGASUS to elicit constructs is that of triadic comparison:

```
Triad for elicitation of construct 1
```

```
1 selling newspapers from a stand
2 selling antiques in a market
3 selling custom drapes
```

```
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?` **`1`**

```
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 selling antiques in a market and selling custom drapes on
the left pole from selling newspapers from a stand 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 --? **no. of different objects**
right pole rated 5 --? **same object in quantity**

```
Now assume that selling antiques in a market and selling custom drapes are
assigned the value 1 and selling newspapers from a stand 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

 2 selling antiques in a market      1
 3 selling custom drapes             1
 1 selling newspapers from a stand   5
 4 selling large appliances         ?3
 5 hiring out a car to a customer   ?1
 6 providing TV rental services     ?2
 7 selling car insurance            ?4
 8 providing electricity for home   ?5
 9 providing dental services        ?2
10 being agent for home purchase    ?3
```

Another technique is that of matching the ratings given to elements and asking for a construct that differentiates between two highly matched elements:

```
The two elements 1 selling newspapers from a stand and 8 providing electricity
for home are matched at the 87 percent level. This means that so far you have
not distinguished between selling newspapers from a stand and providing
electricity for home.
```

Do you want to split these? **Y**

```
Think of a construct which separates these two elements, with selling
newspapers from a stand on the left pole and providing electricity for home on
the right pole.
```

left pole rated 1 --? **supplied one at a time**
right pole rated 5 --? **direct delivery**

```
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 selling newspapers from a stand   1
 8 providing electricity for home    5
 2 selling antiques in a market     ?1
 3 selling custom drapes            ?2
 4 selling large appliances         ?1
 5 hiring out a car to a customer   ?1
 6 providing TV rental services     ?4
 7 selling car insurance            ?3
 9 providing dental services        ?1
10 being agent for home purchase    ?1
```

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:

```
The two constructs you called

5 direct delivery -- supplied one at a time
9 'hands-on' delivery -- nonpersonal delivery

are matched at the 75 percent level
```

```
This means that most of the time you are saying direct delivery you are also
saying 'hands-on' delivery and most of the time you are saying supplied one at
a time you are also saying nonpersonal delivery. Think of another element
which is either direct delivery and nonpersonal delivery or 'hands-on'
delivery and supplied one at a time. 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?
```

Through a combination of triadic comparison, element matching, and construct matching, a PEGASUS interaction draws out from the user his conceptual framework relating to the purpose of the elicitation.

SOCIOGRIDS in Action

The construct elicitation described in the previous section is being carried out with a number of subjects from those experienced in business record-keeping to students with some knowledge of accounting. It is clearly possible to analyse for each subject whether the BIAIT constructs are among those elicited, how closely they are mimiced, and which ones are missing, and this analysis is being done. However, as noted earlier, one can also attempt to validate the methodology by seeing whether it is possible to reconstruct the ranking in expertise of the subjects through a SOCIOGRIDS analysis. For example, in an earlier study (Shaw, 1980) of quality control for a garment company the SOCIOGRIDS analysis re-constructed 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 (1955) terms the *commonality corollary* 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.

SOCIOGRIDS derives linkages between people (socionets) by comparing grids in pairs and calculating the average over all constructs in grid A of the similarity of the best match for each construct in grid B. This measures to what extent the system of construction modeled in grid B subsumes that modeled in grid A. These subsumption averages are ordered and used to produce socionets showing the degree to which one person's construction system subsumes that of another. SOCIOGRIDS derives a consensual construction system (mode grid) by comparing each construct in turn with those in every grid and calculating the average over all grids of the best match for the construct in each grid. This measures to what extent a given construct is represented in the construction system of every person, and hence is consensual.

Link 1 in Figure 5 shows the most highly matched pair which is BIAIT and a businessman; link 2 shows another matched pair which is an economics student and a visual arts student. As the links are formed these initial pairs are joined into one group at link 4. A new subgroup is formed at link 6, then gradually the other participants are brought in and joined into one group. The nodes forming the central part of the network are BIAIT, the two businessmen, an economics student and the visual arts student. This can be seen at link 14 which is shown more completely in Figure 6.

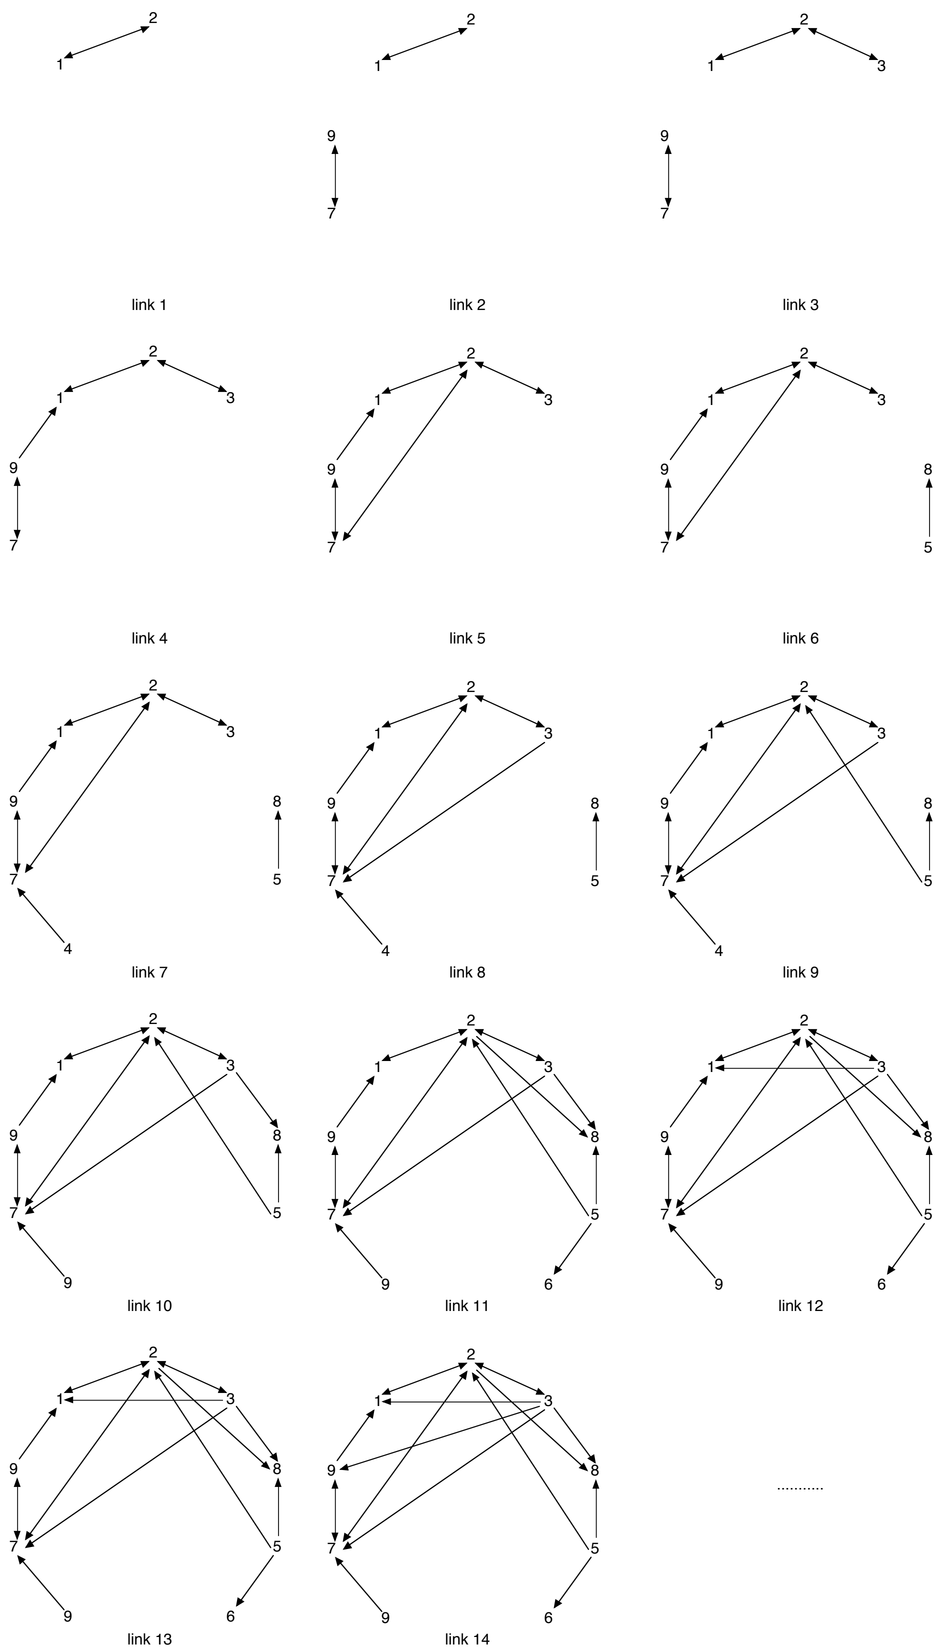

Figure 5 Series of socionets produced by SOCIOGRIDS

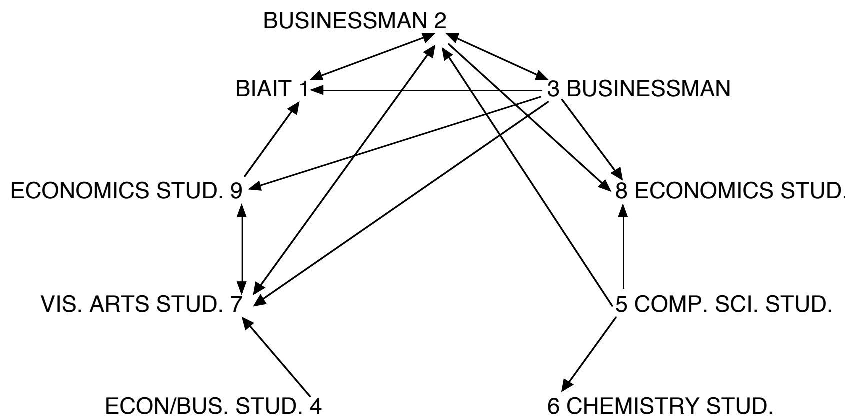

Figure 6 Socionet produced by SOCIOGRIDS

This is mainly what one 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. 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.

Other Developments

ENTAIL (Gaines and Shaw, 1980; Shaw and Gaines, 1982) gives a dependency analysis of the grid by listing logical entailments not falsified by the data and it was expected that this would aid removal of unnecessary constructs. However, this is only partially successful because the grid as elicited contains mixed cases which show up as non-extreme ratings (2,3,or 4). For example, some customers of Hertz pay cash and others by credit card. When these customers are rated on other constructs this information is lost. The interaction is being extended to query the rating of elements that are not at extremes. The user will be given the opportunity to describe the elements in more detail as short case histories (i.e. instantiated case frames) and to add qualifications to these as he construes them in relation to new constructs.

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 which 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.

The system we have described is concerned primarily with the expert's conceptual structure for describing problems. It can also be applied to his conceptual structure for describing solutions or actions, in this case the different forms of record keeping appropriate to different types of business. Rule elicitation then consists of determining the relationship between these two conceptual spaces. This enables the techniques described in this paper to be used to construct complete systems using decision trees or other means to reproduce expertise.

Summary and Conclusions

PLANET, an interactive computer aid to expedite the initial stages of knowledge engineering has been described and illustrated through an example. It is based on the personal contruct model of human knowledge acquisition and reasoning that has deep systemic and psychological foundations. The use of these foundations to operationalize the soft systems analysis at the heart of knowledge engineering has been described. PLANET contains a number of sub-programs to do this: PEGASUS elicits from an expert the fundamental distinctions, or constructs, he is using in applying his expertise; SOCIOGRIDS analyses the relationships between members of a group through their differing construct systems and determines consensual constructs; ENTAIL analyses the logical dependencies between constructs and can present these as rules.

A basis for validating the methodology by eliciting knowledge in an area where an optimum structure has already been developed has been presented. An ongoing empirical trial of the techniques in a test case based on the BIAIT analysis of the record-keeping requirements of business enterprises has been described. The initial results show that the spectrum of expertise in the group of subjects used is being re-constructed in the analysis and hence that the methodology shows promise of being valid.

Acknowledgements

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

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.

Brown, G.S. (1969). **Laws of Form**. London, George Allen & Unwin.

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, UK, Wiley.

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

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

Gaines, B.R. (1977). System identification, approximation and complexity. **International Journal of General Systems 2**(3) 241-258.

Gaines, B.R. (1980). General systems research: quo vadis? Gaines, B.R., Ed. **General Systems 1979**. pp.1-9. Kentucky, Society for General Systems Research.

Gaines, B.R. (1983). Precise past - fuzzy future. **International Journal of Man-Machine Studies 19**(1) 117-134.

Gaines, B.R. and Shaw, M.L.G. (1980). New directions in the analysis and interactive elicitation of personal construct systems. **International Journal Man-Machine Studies 13** 81-116.

Gaines, B.R. and Shaw, M.L.G. (1981). A programme for the development of a systems methodology of knowledge and action. Reckmeyer, W.J., Ed. **General Systems Research**

and Design: Precursors and Futures. pp.255-264. Louisville, Kentucky, Society for General Systems Research.

Gaines, B.R. and Shaw, M.L.G. (1984). **The Art of Computer Conversation**. Englewood Cliffs, New Jersey, Prentice-Hall.

Goguen, J.A. (1974). Concept representation in natural and artificial languages: axioms, extensions and applications for fuzzy sets. **International Journal Man-Machine Studies 6** 513-561.

Goguen, J.A., Meseguer, J. and Plaisted, D. (1983). Programming with parametrized abstract objects in OBJ. Ferrari, D., Bolognani, M. and Goguen, J., Ed. **Theory and Practice of Programming Technology**. pp.163-193. Amsterdam, North-Holland.

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., Ed. (1983). **Building Expert Systems**. Reading, Massachusetts, Addison-Wesley.

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

Klir, G.J. (1976). Identification of generative structures in empirical data. **International Journal of General Systems, 3** 89-104.

Klir, G.J. (1985). **Architecture of Systems Problem Solving**. New York, Plenum Press.

Mancuso, J.C. and Adams-Webber, J.R., Ed. (1982). **The Construing Person**. New York, Praeger.

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

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

Norman, D. (1980). Twelve issues for cognitive science. **Cognitive Science 4**(1) 1-32.

Pask, G. (1975). **Conversation, Cognition and Learning**. Amsterdam, Elsevier.

Pope, M.L. and Keen, T.R. (1981). **Personal Construct Psychology and Education**. London, Academic Press.

Quinlan, J.R. (1983). Inferno: a cautious approach to uncertain inference. **Computer Journal 26**(3) 255-269.

Ralescu, D. (1979). A system theoretic view of social identification. Ericson, R.F., Ed. **Improving the Human Condition: Quality and Stability in Social Systems**. pp.531-538. Louisville, Society for General Systems Research.

Shaw, M.L.G. (1980). **On Becoming A Personal Scientist: Interactive Computer Elicitation of Personal Models Of The World**. 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. van Gigch, J.P., Ed. **Metamodels and Metasystems**. pp.185-198. Cambridge, MA, Abacus Press.

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

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). The personal scientist in the community of science. Reckmeyer, W.J., Ed. **General Systems Research and Design: Precursors and Futures**. pp.59-68. Kentucky, Society for General Systems Research.

Shaw, M.L.G. and Gaines, B.R. (1982). Eliciting entailment. Trappl, R., Ricciardi, L. and Pask, G., Ed. **Progress in Cybernetics and Systems Research. Vol. IX**. pp.425-435. Washington, Hemisphere.

Shaw, M.L.G. and Gaines, B.R. (1983). Eliciting the real problem. Wedde, H., Ed. **Adequate Modeling of Systems**. pp.100-111. Berlin, Springer.

Shaw, M.L.G. and Gaines, B.R. (1984). Deriving the constructs underlying decision. Zimmermann, H.J., Zadeh, L.A. and Gaines, B.R., Ed. **Fuzzy Sets and Decision Analysis. TIMS Series Studies in the Management Sciences, 20**. pp.335-355. Amsterdam, North-Holland.

Shepherd, E. and Watson, J.P. (1982). **Personal Meanings**. Chichester, Wiley.

Sowa, J.F. (1984). **Conceptual Structures: Information Processing in Mind and Machine**. Reading, Massachusetts, Adison-Wesley.

Varela, F.J. (1975). A calculus for self-reference. **International Journal of General Systems 2** 5-24.

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

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