Kacprzyk/Yager (Ed.): Management decision support systems using fuzzy sets and possibility theory
ISR 83 © Verlag TÜV Rheinland GmbH, Köln

FOUNDATIONS OF EXPERT SYSTEMS

Brian R. Gaines & Mildred L. G. Shaw

Department of Industrial Engineering
University of Toronto, Ontario, Canada M5S 1A4
&
Department of Computer Science, York University,
4700 Keele Street, Downsview, Ontario, Canada M3J 1P3

The use of fifth generation computer-based expert systems as management decision support tools will be a major trend during the next decade. However, the knowledge technology underlying such systems is still largely pragmatic and lacks coherent foundations. In this paper we illustrate the operation and application of expert systems, and survey the foundations of expertise from both psychological and system theoretic points of view and propose an overall model of expertise. We relate this to the underlying knowledge structures and show how these are based on a calculus of distinctions that itself underlies the theory of fuzzy sets and possibilistic systems. We show how the personal constructs forming distinctions generate a knowledge hierarchy and how this may be analysed through soft systems analysis and conversation theory. Finally, we describe how the theory has been made operational through computer programs.

Key-words: Management decision systems, expert systems, knowledge-based systems, fifth generation computers, fuzzy sets, calculus of distinctions, possibilistic systems, epistemological hierarchy, personal construct psychology, soft systems analysis, conversation theory.

1. KNOWLEDGE AND POWER

"Human knowledge and human power meet in one; for where the cause is not known the effect cannot be produced. Nature to be commanded must be obeyed; and that which in contemplation is as the cause is in operation as the rule" (Bacon, Novum Organum Book 1, III)

Knowledge is power, but what is knowledge? Management decision systems are often taken to be synonymous with management information systems. This is not so. Information is not knowledge and a surfeit of information is apt to weaken our decision processes not support them. This is particularly so for the types of information that computers are currently programmed to carry. It is primarily the precise quantitative data that is most closely analogous to the mathematical structures underlying current computer systems. We have taken that fragment of our knowledge that can be most readily distorted to fit in computers and have recorded it; trimming away any nuances that do not fit. We can then replay it and have instant access to vast quantities of distorted fragments of our knowledge.

This is not power, nor any path to it. Even the numeric data that does fit naturally on the computer is poor in content compared with its sources. This bottom line and these ratios look attractive but do we believe in them? Underlying the figures in the spread sheet is a community of people, production staff, salesmen and customers, whose dynamics are not captured in these simplistic figures. In March we will sell 6 machines, but what does that "6" mean? It was agreed at a Board meeting which authorized this budget. The Executive Committee had a figure of 8 from the VP of that Division but reduced it to 6 as a more conservative estimate. The Divisional VP had a figure of 12 from his Sales VP who obtained this by adding up four estimates from his Regional Sales Managers that totalled 17 machines. The Western Sales Manager had 9 prospects identified by his salesmen as probably closing in time for March shipment. For each potential customer there is a story to be told of requirements, funds, competition, timing, relationships, and so on.

How much do these dynamics underlying the ultimate simple figures matter? We have become so used to regarding financial statements as the operating parameters of a business that this may be a difficult question to answer. The budget is real and the chains of guesses, uncertainties, possibilities, and so on underlying it are but smoke drifting around and sometimes obscuring its hard reality. Yet that haze of uncertainty, of human foible, of changing circumstances, is what managers actually manage. That is the day to day reality of business. Peering through that haze, accommodating its drifting nature, manipulating the remote causes of which it is the perceived effect, these are the skills of management.

A manager uses his knowledge of reality to control it and hence make the ficticious figures of the budget become real. If we are not able to capture his knowledge in the computer then the machine has little part to play in either reproducing or enhancing the manager's expertise. This paper is concerned with recent developments in **expert systems** and **fifth-generation computers** that are a step toward more realistic management decision systems, and the theoretical and practical foundations of **knowledge engineering** for these systems.

2. FIFTH GENERATION COMPUTERS AND EXPERT SYSTEMS

The first four generations of computers were **information processors.** The transition between each generation has been different in nature (Gaines 1984). The first generation (1948-56) used vacuum tubes and delay lines and this **hardware** became transistors and core stores in the second generation (1956-64). The transition to the third generation (1964-72) was dominated by the availability of **software,** operating systems and compilers. The transition to the fourth generation (1972-1980) was one of **person-computer interaction** in which interactive usage and personal computing became common. The transition to the fifth generation (1980-88) has been one to **knowledge processing** with the development of expert systems and knowledge-based systems.

Each generation of computers has made its own contribution to management decision support systems: numeric processing, modeling, databases, information retrieval, and communications. However, all of these aids have been restricted to supporting the manager in handling increasing amounts of complex information, its processing and communication. They have only indirectly supported his decision making because they have not been able to treat information as **knowledge,** that is data embedded in a rich context of associations relating to the real world, the thought processes of people within it and the possibility of actions upon it.

The primary objectives of the Fifth Generation Computer development program initiated by the Japanese (Moto-oka 1982) are two-fold:

a) to provide knowledge processing rather than information processing;
b) to offer natural human-computer interfaces.

These objectives have been accepted as realistic because of the successful development of **expert systems** (Michie 1979) that encode human knowledge and cognitive skills for the computer and make them accessible to others. The significance of such systems is best seen through an example. MYCIN (Shortliffe 1976) is a medical diagnosis expert system which aids a clinician to act as a consultant on infectious diseases. TEIRESIAS (Davis & Lenat 1982) is an interface to MYCIN which aids a clinician in setting up and modifying the knowledge rules for MYCIN. Figure 1 shows the relation between the two systems.

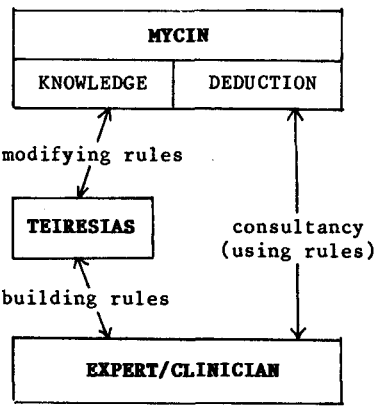

Fig. 1 MYCIN and TEIRESIAS, an advanced expert system

Figure 2 shows how MYCIN operates in an interactive dialog mode to acquire information from the clinician about a patient (clinician's input in **BOLD**).

```
--------PATIENT-1--------
(1) Patient's name: (first-last)
**ALBERT HALL
(2) Sex:
**MALE
(3) Age:
**55
(4) Have you been able to obtain positive cultures from a site at which
    Fred Smith has an infection?
**YES
  --------INFECTION-1--------
  (5) What is the infection?
  **PRIMARY-BACTEREMIA
  (6) Please give the date and approximate time when signs or symptoms
      of the primary-bacteremia (INFECTION-1) first appeared. (mo/da/yr)
  **MAY 5, 1975
  The most recent culture associated with the primary-bacteremia
  (INFECTION-1) will be referred to as:
    --------CULTURE-1--------
    (7) From what site was the specimen for CULTURE-1 taken?
    **BLOOD
    (8) Please give the date and time when this blood culture (CULTURE-1)
```

```
    was obtained. (mo/da/yr)
**MAY 9, 1975
The first significant organism from this blood culture (CULTURE-1)
will be referred to as:
  --------ORGANISM-1--------
  (9) Enter the identity of ORGANISM-1:
  **UNKNOWN
  (10) Is ORGANISM-1 a rod or coccus (etc.):
  **ROD
  (11) The gram stain of ORGANISM-1:
  **GRAMNEG
  (12) Have there been POSITIVE cultures from sites which are NOT
       associated with infections about which you are seeking advice?
  **NO
  (13) Did ORGANISM-1 grow aerobically?
  **YES (8)
  (14) Did ORGANISM-1 grow anaerobically?
  **YES
  (15) What is the suspected portal of entry of ORGANISM-1
       into the sterile site?
```

Fig. 2 MYCIN gathering information

MYCIN eventually makes a diagnosis based on a set of knowledge rules of the form shown in Fig. 3.

```
RULE50
  If  1) the infection is primary-bacteremia, and
      2) the site of the culture is one of the sterile sites, and
      3) the suspected portal of entry of the organism is the
         gastro-intestinal tract
  Then there is suggestive evidence (.7) that the identity of the
       organism is bacteroides.
```

Fig. 3 A MYCIN rule

MYCIN works very well when the rules are complete and correct, but it is difficult to set up the rules initially. It is also difficult to trace the rules causing errors in the deductions of MYCIN if the clinician feels that something has gone wrong. TEIRESIAS was added as an auxiliary system with expertise about MYCIN. It has the facility to explain MYCIN's decisions, and help the clinician amend the rules when they lead to incorrect conclusions. TEIRESIAS uses a similar rule-based approach to reasoning as does MYCIN but the rules are now rules about the forms of rules and the use of rules. Figure 4 shows one of these **metarules.**

```
METARULE003:
If  1) there are rules which do not mention the current goal in
       their premise
    2) there are rules which mention the current goal in their
       premise
Then it is definite that the former should be done before the
latter.
```

Fig. 4 A TEIRESIAS metarule

Figure 5 shows the clinician, having completed an interaction with MYCIN and being unsatisfied with the diagnosis, using TEIRESIAS to change the rules.

```
My therapy recommendations will be based on the following probable infection(s)
and potential causative organism(s):
INFECTION-1 is PRIMARY-BACTEREMIA
    <Item 1>  E.COLI [ORGANISM-1]
    <Item 2>  PSEUDOMONAS-AERUGINOSA [ORGANISM-1]
Are all these identities correct?
++++ N
Are there any identities missing?
++++ Y
Which one(s)?
++++ KLEBSIELLA-PNEUMONIAE PROTEUS-NON-MIRABILIS
Are there any identities that shouldn't have appeared?
++++ Y
Which one(s)?
++++ PSEUDOMONAS-AERUGINOSA
Do you want to
      1] track down the problem yourself, or
      2] allow me to guide the debugging, or
      3] ignore it for now and fix it another time?
Type 1, 2 or 3
++++ 2
We'll start with the PSEUDOMONAS-AERUGINOSA first. The following rules were used
in deducing that the identity of ORGANISM-1 is pseudomonas-aeruginosa
RULE184
Since [1.1] the category of ORGANISM-1 is not known
      [1.2] the gram strain of ORGANISM-1 is gramneg
      [1.3] the morphology of ORGANISM-1 is rod
      [1.4] the aerobicity of ORGANISM-1 is facultative
  There is weakly suggestive evidence (.3) that the  identity of ORGANISM-1 is
  pseudomonas-aeruginosa
Is this rule correct?
++++ Y
Should its premise have failed for this case?
++++ Y
Which clause of the premise should have been false? [give #]
++++ 1
No conclusions were reached concerning the category of ORGANISM-1.
Should it have been possible to conclude a category?
++++ Y
Which one?
++++ ENTEROBACTERIACEAE
There were no rules successfully invoked in the attempt to deduce that the
category of ORGANISM-1 is enterobacteriaceae.
The following (potentially applicable) rules were NOT executed because the
starred [**] clause failed:
RULE116
If    1) the identity of ORGANISM-1 is not known
   ** 2) the gram strain of ORGANISM-1 is not known **
      3) the morphology of ORGANISM-1 is not known
      4) the site of CULTURE-1 is csf
      5) the infection is meningitis
      6) the age (in years) of Albert Hall is less than
          or equal to .17
then  There is weakly suggestive evidence (.33) that
      the category of ORGANISM-1 is enterobacteriaceae
RULE050
If    1) the morphology of ORGANISM-1 is rod
      2) the gram strain of ORGANISM-1 is gramneg
```

```
      3) the aerobicity of ORGANISM-1 is facultative
   ** 4) the infection with ORGANISM-1 was acquired
         while the patient was hospitalized **
then  There is suggestive evidence (.7) that the
      category of ORGANISM-1 is enterobacteriaceae
Are all these rules correct?
++++ Y
Should any of them have been successfully invoked?
++++ N
Is there a rule missing?
++++ Y
Ok, then I need a rule that would allow me to deduce that the category of
ORGANISM-1 is enterobacteriaceae.
Would you like to give me such a rule now?
++++ Y
```

Fig. 5 Changing MYCIN's knowledge base by interaction with TEIRESIAS

In this dialog the clinician noticed an identity was missing in those deduced by MYCIN. He chooses to allow TEIRESIAS to guide the debugging, and it uses the metarule of starting with the identity which appeared incorrectly. The program then reconstructs the deductive process, asking the clinician to confirm its actions at each stage, or comment on anything incorrect. It confirms the facts about the particular case in hand and checks that the rules are correct. Before very long the issue has become sharply focused for the clinician and the system. The guideline is that a rule must be found which is relevant to this particular patient, but true in general, and enhancing the set of rules of knowledge. The new rule is added to the set of knowledge rules, with a tag explaining who added it and in what circumstances. This allows the ultimate level of explanation to be given in terms of who added the rule and the reasons given. Whereas MYCIN's rules are specific to microbial infections, those of TEIRESIAS are more general and can be used in other domains. Davis (1983), for example, shows TEIRESIAS being used as an investment decision system for clients of a stockbroker.

The highest level of generality is shown in the rules of Lenat's AM (Davis & Lenat 1982) which searches for "interesting" conjectures in mathematics. What is remarkable about AM is the very notion of what it is doing. This is the highest level of scientific creation, looking for the patterns underlying knowledge. The concept of something being interesting seems peculiarly human and certainly too vague to form the basis of a computer program. However, AM is not only able to generate profound conjectures in mathematics but also meaningful ways of arriving at them. It is judged by the quality of its musings as it considers what it is doing.

```
19. If concept C possesses some very interesting property lacked by one
      of its specializations S,
    Then both C and S become slightly more interesting.
38. If there are no known examples for the interesting concept X,
    Then consider spending some time looking for such examples.
```

Fig. 6 Two rules of high generality from AM

Figure 6 shows two of the few hundred rules used by AM. Most of them are sufficiently general to apply to other situations. In a business context, for example, Rule 19 indicates that a committee which has decision-making powers that do not apply to its sub-committees below a certain size has an interesting property (that of requiring a quorum). Its sub-committees also have an interesting property (that of being quorate).

Expert systems based on the MYCIN/TEIRESIAS concepts have been developed for a wide range of scientific, industrial and commercial applications (Gevarter 1983). There are many situations in commercial management, government policy and personal decision-making where there are no absolute rules but only a variety of guidelines coming from different sources with different priorities. It is difficult for people to keep all these rules in mind, relate them one to another, use them in decision-making, note their effects, and modify them when they have undesired consequences. The computer is coming to provide an effective medium for conversation with the multiple rule sources on whom we rely, consciously or subconsciously, in arriving at our decisions and actions. If we can encode the expertise expressed in these rules then we are able to use the computer not so much to make our decisions for us, but rather to discuss and explain the basis for our decisions with us so that we are able to understand our own rationale (Gaines & Shaw 1984).

What appears remarkable about expert system developments is that they are concerned with recording what had previously been regarded as very high level human expertise, difficult to explain to another person let alone program for a computer. However, there is an alternative point of view that notes that computers have always been used to encode human expertise. An accountancy program for a business that keeps track of purchases and sales and prepares invoices, purchase and sales ledgers, and so on, may be seen as the recording of the expertise of an accountant for use by an ordinary businessman. An auditor evaluating such a program may expect to be able to ask it exactly those questions he would ask the accountant, for example, from what original information did you calculate this figure? If the program does not have built-in "audit trail" facilities which enable it to answer such questions then it is inadequate in exactly the same way that an accountant would be who could not answer that question; it is a simulation of an incompetent accountant. It merely manipulates information without using knowledge.

Thus, expert systems may be seen as bridging the gap between information processing and knowledge processing. In the first four generations of computers we could encode only that part of human knowledge concerned with facts, arithmetic and precise logical deduction. This did not enable us to emulate a person's capability to reason with imprecise, possibly contradictory, information that requires knowledge of context, social norms, and the expectations of others, and the ability to use rules which are heuristic in that they are generally useful but may sometimes be wrong. Fifth generation systems using knowledge processing and improved person-computer interaction through natural language are providing facilities to emulate the capabilities of a skilled colleague and advisor.

1. When the burning zone temperature is drastically low:
   (a) reduce kiln speed;
   (b) reduce fuel.
2. When the burning zone temperature is slightly low:
   (a) increase I.D. fan speed;
   (b) increase fuel rate.

Fig. 7 Rules from a lime kiln operator's manual

The rules on which a particular expert system is based may be seen as ways of representing knowledge. This is not a new concept as such rules have been used in the past quite independently of computer systems. Figure 7 is an excerpt

from a training manual for lime kiln operators (Peray & Waddell 1972). It seems that the encoding of expertise into a set of fairly vague but highly applicable rules is a standard way of making it available to others. Mamdani and Assilian (1981) have shown that rules such as this may be interpreted algorithmically to produce computer control programs that are more effective than those derived from classical control theory. In 1974 they developed the first expert system that took human knowledge expressed verbally and encoded directly as computer programs which acted on the environment. The system they controlled was a small steam engine and the rules were of the form:

> If the pressure error is positive and big and
> the change in pressure error is not negative medium or big
> then make the heat change negative and big

Fig. 8 Rule from fuzzy logic controller

It will be noted that the form of this rule is imprecise, using terms such as "medium" and "big." Similarly the MYCIN rule of Fig. 3 notes that there is "suggestive evidence" for its inference rather than that it is true. Mamdani and Assilian were able to encode these imprecise rules and make inference from them using Zadeh's (1965) fuzzy sets theory and its application to decision processes (Zadeh 1973). Shortliffe uses a similar method in MYCIN where the ".7" in the rule is a truth value in a multivalued logic, and Lenat uses such a measure of "interest" in AM. Our scientific training has taught us to regard imprecision as undesirable and it was not easy for the pioneers of expert systems to determine that it was a necessary component of any system adequate to encode human knowledge. However, a major weakness of past information processing systems, such as databases, has been that they are not able to cope with the vague expression of information in everyday life (Gaines 1981).

As applications move out of the numeric realms of the accountant and into the less quantitative areas of the manager so also does a requirement for high precision in specifying retrievals and updates of a database become increasingly unnatural. As an example of the value of being able to store and process imprecise information consider a personnel database which keeps track of an employee's age. Current database systems require the age in years and a value "young" cannot be stored. Yet this information is adequate to exclude the employee in an enquiry such as "list all employees within two years of pensionable age." This is an example where we are imprecise because we lack information. It is still possible to argue that precision in itself is always a virtue. A reply to a request for a delivery date that says "soon" would be rated less satisfactory than one that says "in 3 days time"; the latter is more "businesslike."

However, there are other examples where the imprecision is intrinsic to the situation and being precise would falsely represent it. For example:

1) we will deliver 7 parcels each weighing 15.2 Kilograms at the rear entrance of building 6A on 15th February at 9.03 p.m.;
2) we will deliver some heavy equipment to your site Saturday evening;
3) see you with the goods over the weekend;

may each refer to the same event but are clearly not equivalent; each conveys an exact meaning that (presumably) properly represents what is to occur. If we prefer the precision of the first statement it is not for its own sake but because the tighter tolerances it implies on the actual situation allow us to plan ahead with greater accuracy and less use of resources. However, if the third statement really represents all that can be said it would be ridiculous to replace it with either of the previous ones. It would be equally ridiculous to say nothing, however, since even the least precise of the three statements does

provide a basis for planning and action. A key aspect of executive action is planning under uncertainty and normal language provides a means for imprecision to be clearly and exactly expressed.

In retrieving information from a database the requirement for artificial precision is at least irritating and at worst highly misleading. For example the request:

a) list the young salesmen who have a good selling record for household goods in the north;

is perfectly comprehensible to a person. Translating it into:

b) list salesmen under 25 years old who have sold more than $20,000 of goods in the categories.... to shops in the regions....;

generates unnecessary work and makes no allowances for the whole spectrum of tradeoffs possible. It will not list the salesman of 26 who has made a real killing, or the one of 19 who sold $19,000, etc. The second request is more precise but the first represents far more accurately the actual **meaning** of the retrieval required. It is this comprehension of meaning that we term **knowledge** and see as essential to expert systems.

4. KNOWLEDGE ENGINEERING

The transition from information processing to knowledge processing is not proving an easy one. The Japanese project has been presented in the conceptual diagram of Fig. 9 (Moto-oka 1982 p.29) where the rings for hardware and software are filled with detail but that for the human is empty:

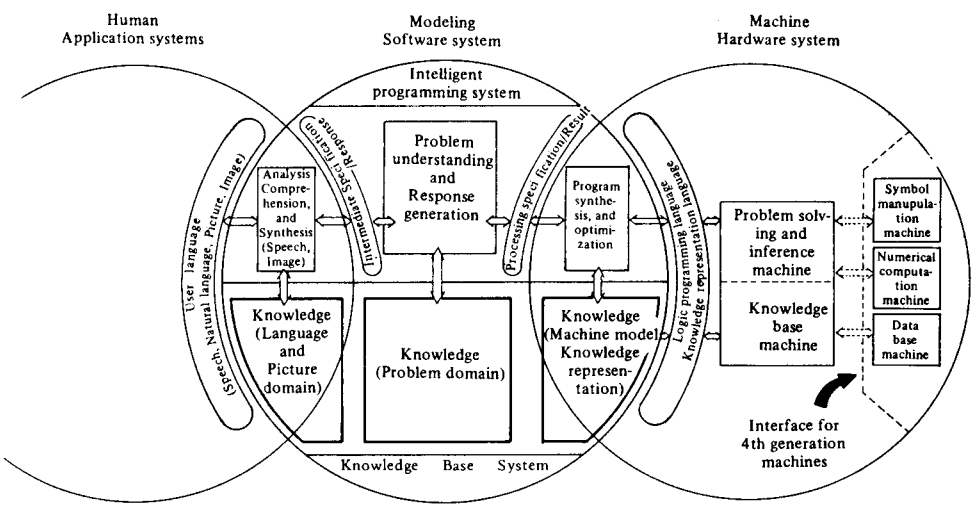

Fig. 9 Conceptual diagram of a fifth generation computer system

The third ring should contain an account of human knowledge acquisition and processing, and its lack has created problems in what Feigenbaum (1980) terms **knowledge engineering**, the reduction of a large body of knowledge to a precise set of facts and rules. As Hayes-Roth, Waterman and Lenat (1983) note:

"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."

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 de-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.

We also have the problems not only of acquiring knowledge but also of representing it and using it for inferences and as a basis for action. The next section introduces philosophical, psychological and systemic foundations for knowledge engineering that can be made operational in computer programs.

5. DISTINCTION MAKING, PERSONAL CONSTRUCTS AND SYSTEMS

Clinical psychology is the prime source of literature on techniques for by-passing cognitive defenses. Kelly's **Personal Construct Psychology** (PCP) is particularly significant because it develops a complete psychology of both the normal and abnormal, which has strong systemic foundations (Gaines & Shaw 1981a) and has been operationalized through computer programs (Shaw 1980). PCP models a person as using **personal constructs** as distinctions made about experience:

"transparent templets which he creates and then attempts to fit over the realities of which the world is composed" (Kelly 1955)

and emphasizes the use of constructs as a basis for knowledge acquisition:

"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)

When Kelly developed his theory there was neither the mathematics to formalize it nor the computer technology to operationalize it. Despite this, PCP has developed to become a major psychological theory (Mancuso & Adams-Webber 1982) and to provide a variety of techniques (Shaw 1981) used in clinical (Shepherd & Watson 1982) educational (Pope & Keen 1981) and managerial (Shaw 1980) psychology. In recent years strong links with system theory have developed that give formal foundations for PCP. Brown bases his logical system in **Laws of Form** on the notion underlying personal constructs of **making a distinction:**

"The theme of this book is that a universe comes into being when a space is severed or taken apart...By tracing the way we represent such a severance, we can begin to reconstruct, with an accuracy and coverage that appear almost uncanny, the basic forms underlying linguistic, mathematical, physical and biological science, and can begin to see how the familiar laws of our own experience follow inexorably from the original act of severance." (Brown 1969)

Kelly's psychology and Brown's logic may both be seen as based on taking as primitive an axiom of unrestricted predication. In set-theoretic terms they take the axiom of comprehension (Fraenkel, Bar-Hillel & Levy 1973) as **a priori** and derive all other concepts from this. This is not the standard and familiar approach, and this makes it very misleading to translate their discussions into logic and set theory; they are defining a basis of knowledge which is prior to both and from which logic and set theory must be derived. The key questions are: to what does a construct apply; what is it that is distinguished? We fall into a trap if we have any answer based on other knowledge, particularly the obvious ones of logic and set theory. The proper answer is that what is distinguished is defined only by being distinguished; it requires an auxiliary ontological postulate to reify this by assuming some entity is thus defined.

We have suggested elsewhere (Gaines 1980) that what is distinguished is a **system** and that this notion is a sufficient definition for a system. A sytem is what is distinguished as a system and in this lies the essence of systems theory. To distinguish some entity as being a system is a necessary and sufficient criterion for its being a system and this is uniquely true for systems. This definition may appear too weak in that it does not involve the expected properties of a system. The Oxford English Dictionary notes that a system is "a group, set or aggregate of things, natural or artificial, forming a connected or complex whole." However, these are properties that derive from our definition: a set of things is treated as distinct from another and it is this which gives it coherence; it is also this which increases its complexity by giving it one more characteristic than it had before, that it has now been distinguished.

6. THE LOGIC OF KNOWLEDGE SYSTEMS

The significance of an axiom of unrestricted predication as a foundation for a theory of knowledge is that it forces us to develop system theories that do not run into the problems of this axiom in conventional mathematics. In set theory such predication leads to the inconsistency that Russell discovered in Frege's (1884) **Foundations of Arithmetic**, and requires the axiom of comprehension to be severely restricted by a complex set of rules. However, it has long been speculated that the inconsistency may be better ascribed to the law of contradiction in standard logic: that the same predicate cannot be both true and false about the same proposition. Maydole (1975) proved that most contenders for an alternative logic led to inconsistencies in set theory, and White (1979) showed that Lukasiewicz infinitely valued logic gives a consistent set theory.

This was a particularly important result in bringing theory and practice together because it was just this logic that Zadeh had taken as a basis for fuzzy sets theory and which Mamdani and Assilian had used in their expert system for linguistic control. The links from Brown's calculus of distinctions to Lukasiewicz logic had already been established. Varela (1975) used the calculus to analyse self-referential systems using a construction based on Russell's paradox which gave a 3-valued logic. Gaines (1976b) used a recursive form of this construction to generate the infinite valued Lukasiewicz logic and Gaines (1983) shows that the axioms encompass modal, probability and fuzzy logics. The significance of this construction is that the axioms are taken first as defining a non-truth-functional **standard uncertainty logic**: fuzzy logic is then derived by adding the constraint that the logic be truth functional; probability logic is derived by the alternative constraint that the law of contradiction apply. Figure 10 shows the relationships between the various logics of uncertainty that may be derived from the Lukasiewicz axioms for a system of distinctions underlying the constructs representing a person's knowledge.

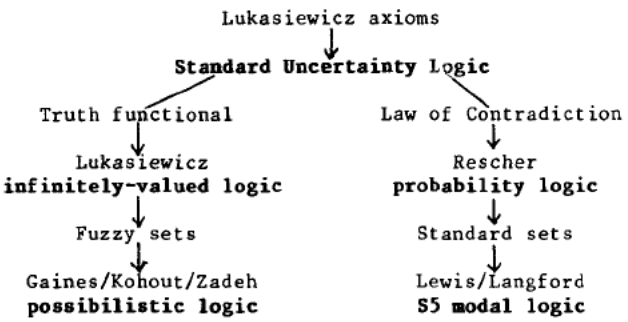

Fig. 10 Logics of uncertainty generated from calculus of distinctions

The need for a system adequate to represent knowledge to encompass both imprecision and non-truth-functionality is obvious in retrospect. Most of the distinctions we make in everyday life do not have hard boundaries. The notion of a red rose is totally unambiguous and acceptable in human communication and reasoning. However, the set of red roses is not a standard set. It has boundary cases which are neither precisely red nor not-red. Yet we can not exclude these borderline cases from being considered as red roses. The laws of contradiction and the excluded middle do not apply. Zadeh (1965) allows for this in fuzzy sets theory by defining a continuous variable as the **degree of membership** to a set, rather than the binary, [0,1], possibility in the standard characteristic function for set membership.

Non-truth-functionality arises because it is not possible ascribe truth-values to the conjunction and disjunction of uncertain events without further information about their relationship. In probability theory this is well-known: if it is 50% certain you will phone today and 50% certain that I will be out of the office, what is the certainty that you will phone when I am out of the office? The answer can range from 25% if the two events are not causally connected to 100% if I do not want to talk to you. It is possible to compute the constraints upon the result but not a single-valued result. Quinlan (1983) has shown that the constraint systems of such non-truth-functional logics may be used as a basis for improved inference in expert systems with uncertain rules and data. Fuzzy logic can be seen as resulting from the requirement to avoid the calculation of such constraints by imposing truth-functionality. The common assumption of independence of events in probabilistic situations is an alternative way of forcing truth-functionality. Both these approaches are expedients whose utility depends on the problem domain.

7. THE HIERARCHY OF MODELING

Distinctions are not just static partitions of experience. They may be operations: actions in psychological terms; processes in computational terms. This shows in Goguen's (1974) category-theoretic formulation of the logics, Ralescu's (1979) application of these to a general theory of modeling, and Goguen and Meseguer's (1983) development of **OBJ**, a category-theoretic programming language. The role of distinctions at the base level of all models is evident in Klir's (1976, 1984) hierarchy of modeling shown in Fig. 11. The loop from events through distinctions up through the modeling hierarchy and then

down again to predictions and actions characterizes what Shaw (1980) has termed the **personal scientist.** Note that the upper levels of modeling are totally dependent on the system of distinctions, or personal constructs, used to express experience through the source system. Klir developed this hierarchy for work on symbolic modeling systems and Gaines (1977) has shown that it forms a basis for general knowledge acquisition algorithms.

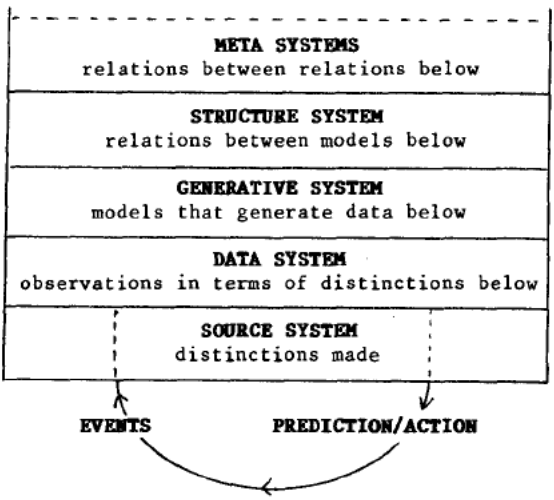

Fig. 11 Epistemological hierarchy of a personal scientist

It is important to note that this hierarchy does not introduce any additional primitives beyond that of making a distinction. The levels of the hierarchy are the results of distinctions that we make. Thus the source system is distinguished as those distinctions that the particular personal scientist makes; it is a distinction about distinctions defining the construct system of an individual. The data system is distinguished as those distinctions that have been made about a particular event; again a distinction about distinctions defining an event. The generative system is distinguished as a set of distinctions that also defines an event; these are model-generated rather than event-generated. It is the match between the model-generated and event-generated distinctions that determines the degree of approximation of the model to the world; this is a distinction about distinctions among distinctions that defines goodness of fit.

A modeling schema results from distinctions about distinctions at each level in the hierarchy. In prediction the key distinction is to what degree a level accounts for the information flowing through it and hence we have termed this distinction one of **surprise** (Gaines 1977), borrowing the term from the economist Shackle (1955) who used it in similar circumstances. Surprise goes in opposition to the degree of membership of a predicted event to an actual event and the expected surprise is a form of entropy. If the degree of membership is m, then taking surprise to be 1-m leads to the expected surprise being a pre-Shannon entropy function, whereas taking it to be $-\log_2 m$ leads to it being the familiar Shannon entropy (Gaines 1976a). Surprise at the lowest level of the hierarchy corresponds to distinctions being inadequate to capture events; surprise at the next level to inadequate variety to experience events; at the next level to inadequate approximation to predict events; at the next level to

inadequate simplicity to explain events; at the next level to inadequate comprehensiveness to account for events. AM's "interest" may be explicated as surprise generated when rules such as property inheritance are violated.

8. THE HIERARCHY OF KNOWLEDGE

The systemic hierarchy based on a calculus of distinctions has an analog based on personal constructs (Gaines & Shaw 1981a) shown in Fig. 12. The lowest level is one of **constructs**, distinctions made in interacting with the world. The next level is one of **experiences**, 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.

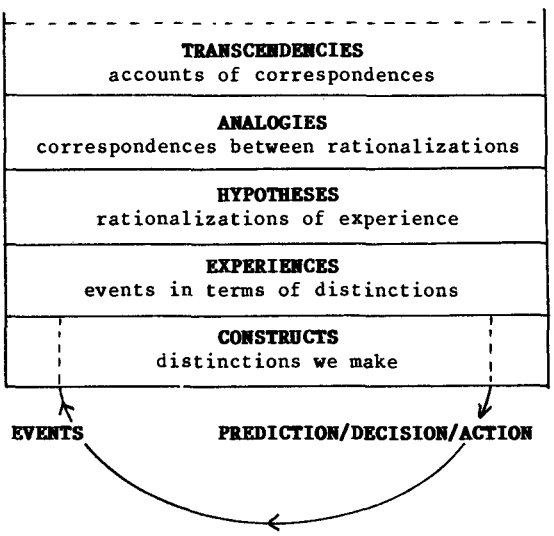

Fig. 12 Knowledge hierarchy of a personal scientist

Kelly (1955) places the major emphasis of his work on the notion of **constructive alternativisim**, that we have a choice in our construct systems at every level in the hierarchy and that real-world problems may often be solved by exercising this choice. Note that this should not be interpreted as an idealist position that ascribes all phenomena to our interpretation of them. Since the construct hierarchy also leads to decision and action, changes in it may equally affect the real world. Kelly and Brown are both neutral to any philosophical stance such as idealism versus realism; it is the distinctions which a philosopher makes that determines his stance and these can be analysed in terms of the epistemological and knowledge hierarchies. Kelly certainly saw his theory as reflexive and the only fundamental principle, apart from that of a construct itself, being that of constructive alternativism. This is important in its applications as a foundation for expert systems since one is attempting to capture the actual processes of the expert, not model them through some pre-conceived system that may not encompass them or distort them.

The hierarchy is strongly idiosynchratic and formalization does not imply uniformity. Constructive alternativism 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 constructs, in the relations between constructs in the hierarchy, and in those constructs most likely to change when 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 Section 4 arise because of the individualistic nature of construct systems. A person may be able to use his construction system whilst having no basis for communicating it to others. Two people may use exactly the same construct yet refer to it by different names. Two people may use the same names for a constructs and yet use it 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.

These systemic, psychological and computational foundations for a model of human cognition based on personal constructs lead to the hypothesis that the key to expertise is in the distinctions which the expert is making. The following sections examine techniques for eliciting these distinctions.

9. ELICITING CONSTRUCTS: 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

Fig. 13 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 (Customers, Actors, Transformation, Weltanshauung, Owner, Environment):

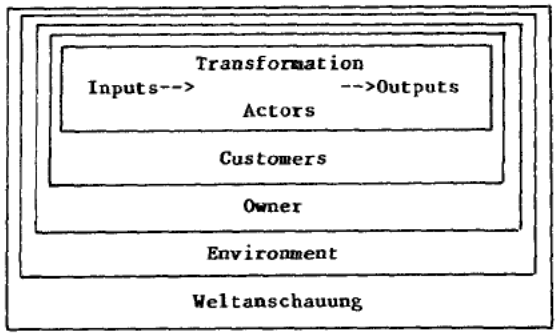

Fig. 14 Root Definitions of Relevant Systems

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.

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 oil extraction takes place, and the **weltanshauung** consists of energy policies, pollution concerns, and so on. These are the considerations which knowledge engineering has to take into account.

Checkland's methodology structures the task of knowledge engineering and provides techniques that are widely applicable. However, it leaves the burden of analysis with a person, and has no obvious operational form as a computer program. Basden (1983) has noted the utility of Checkland's analysis in developing expert systems, and we have shown how the techniques outlined in this paper may be applied to operationalizing it (Shaw & Gaines 1983b).

10. ELICITING CONSTRUCTS: CONVERSATION THEORY

What is required to operationalize soft systems analysis is a theory of the way in which different viewpoints at different levels come together in an overall system of knowledge. Pask's **conversation theory** provides the basis for this and has been applied by Coombs and Alty (1984) to the development of expert systems. Pask introduces the notion of a **P-Individual** or **psychologically characterized individual** in distinction to that of an **M-Individual** or **mechanically characterized individual.** He notes that an M-Individual is:

> "a human being (or his brain) as a biologically self-replicating system"

whereas:

> "a P-Individual...has many of the properties ascribed by anthropologists to a role, in society or industry, for example. A P-Individual is also a procedure and, as such, is run or executed in some M-Individual" (Pask 1975).

He notes that while **an M-Individual may support many P-Individuals**, i.e. a person may play many roles, **a P-Individual may also be distributed over many M-Individuals**, i.e. a social unit may require the interplay of several persons.

In terms of the discussion of Section 5 we may view P-Individuals as cross sections of knowledge hierarchies interfaced to the world through personal constructs. This gives a coherent systemic framework for the relations between the multiple viewpoints discussed in Section 4. For example, Fig. 15 shows John, a person with a number of roles. He has formed P-Individuals as a husband, a sales VP of his company, a backpacker and a fisherman. The **personal scientist** or process which is John as a husband, uses operators to make construct-forming distinctions. They are distinctions about perception, about models, about values, about action. It is with the formation of these constructs that a P-Individual comes into being. John uses different constructs when he is in the role of husband than when he is in the role of fisherman, but some of the constructs he uses in the role of fisherman are also used in the role of backpacker and we show this as overlapping P-Individuals.

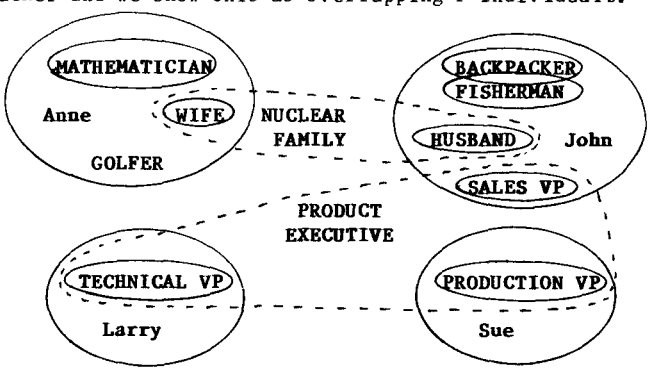

Fig. 15 P-Individuals formed by Groups

Figure 15 includes some of John's associates: Anne is his wife, a mathematician and a golfer; Larry is the technical VP in John's company, and Sue is production VP. Sometimes John is acting with the others as a joint P-Individual that has some constructs shared between them and others that arise from only one. For example, John and Anne form a nuclear family: a P-Individual which has legal rights, behavior, concepts and language which are not those of either party alone. The product executive of the company is a P-Individual which has behavior, authority, responsibility, and language which is unique to the entity and not that of any one of its participants.

As expected from soft systems analysis, Pask's theory of multiple-person P-Individuals shows that expertise may not reside in one person (Shaw & Gaines 1981a). In developing an expert system for product development in a company we have to take into account all the actors shown in the product executive of Fig. 15, the objectives set by the problem owner, the constraints of the environment in which they operate, and the overall ethos for the development.

We have shown the interaction of conversation theory and soft systems analysis with personal construct psychology and how this provides a basis for computer programs for knowledge engineering. The next section outlines briefly some computer techniques for doing this.

11. ELICITING CONSTRUCTS: PEGASUS AND SOCIOGRIDS

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 & Elicitation Techniques, Shaw 1982, Gaines & 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 individual, 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 (Gaines & Shaw 1981b). 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.

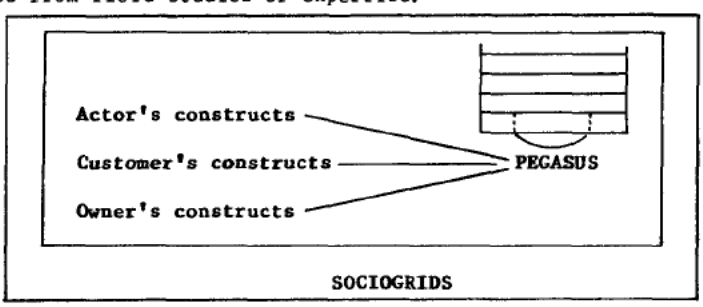

Fig. 16 PLANET automating soft system analysis

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 programs have been used for a wide variety of applications in management, education and clinical psychology (Shaw 1980, 1981, Shaw & Gaines 1981b, Pope & Keen 1980), and for knowledge engineering for expert systems (Shaw & Gaines 1983).

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.

We have reported (Shaw 1984) experiments 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. The By including the BIAIT constructs themselves as part of the data for analysis we were able to show that SOCIOGRIDS analysis regenerates the known hierarchy of expertise in business accounting of those taking part. By using FOCUS to cluster each person's constructs with those of BIAIT we were able to show that the BIAIT constructs all exist in the data, but not with one individual. These experiments show the validity of the approach and the possibility of automating soft systems analysis as a tool for knowledge engineering.

12. SUMMARY AND CONCLUSIONS

The use of computer-based expert systems as management decision support tools will be a major trend in the application of fifth generation computers during the next decade. However, the knowledge technology underlying such systems is still largely pragmatic and lacks coherent foundations. In this paper we have illustrated the operation and application of expert systems, surveyed the foundations of expertise from both psychological and system theoretic points of view and proposed an overall model of expertise. We have demonstrated the need for any system representing knowledge to be able to cope with imprecision and non-truth-functionality. We have related this to the underlying knowledge structures and show how these are based on a calculus of distinctions that itself underlies the theory of fuzzy sets and possibilistic systems.

We have shown how the calculus of distinctions corresponds to the psychological theory of personal constructs and how these generate a hierarchy of knowledge which is a basis for modeling and knowledge acquisition. We have described appropriate soft systems analysis techniques for investigating the knowledge hierarchy of an expert and shown how these link to conversation theory and multiple-person communities of expertise.

Finally we have described an interactive computer aid based on personal construct psychology and the calculus of distinctions which expedites the initial stages of knowledge engineering. 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.

In conclusion, knowledge is power, and the significance of fifth generation computer systems in providing automatic aids to knowledge processing cannot be over-estimated. The management decision support tools of the next generation will provide capabilities that go beyond our current imaginations and those organisations which first make effective use of them will be in a very powerful position. Other papers in this book provide in-depth examples of the mathematical techniques and applications involved. We hope this paper gives some practical and theoretical justifications for interest in these techniques.

13. ACKNOWLEDGEMENTS

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

14. REFERENCES

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

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

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

Mamdani, E.H. & Assilian, S. (1981). An experiment in linguistic synthesis with a fuzzy logic controller. Mamdani, E.H. & Gaines, B.R., Eds. **Fuzzy Reasoning and its Applications.** pp. 311-323. London: Academic Press.

Maydole, R.E. (1975). Paradoxes and many-valued set theory. **Journal of Philosophical Logic, 4,** 269-291.

Mancuso, J.R. & 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.

Peray, K.E. & Waddell, J.J. (1972). **The Rotary Cement Kiln.** New York: Chemical Publishing Co..

Pope, M.L. & 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. **Improving the Human Condition: Quality and Stability in Social Systems.** Kentucky: Society for General Systems Research.

Shackle, G.L.S. (1955). **Uncertainty in Economics.** Cambridge University Press.

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. (1984). Interactive knowledge elicitation. **Proceedings of CIPS SESSION 84.** Calgary: Canadian Information Processing Society (May).

Shaw, M.L.G. & Gaines, B.R. (1981a). The personal scientist in the community of science. Reckmeyer, W.J., Ed. **General Systems Research and Design: Precursors and Futures.** pp. 59-68. Society for General Systems Research.

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

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

Shaw, M.L.G. & Gaines, B.R. (1983b). A computer aid to knowledge engineering. **Proceedings of Expert Systems 83.** pp. 263-271. London: British Computer Society.

Shepherd, E. & Watson, J.P., Eds. (1982). **Personal Meanings.** London: John Wiley.

Shortliffe, E.H. (1976). **Computer-Based Medical Consultations: MYCIN.** New York: Elsevier.

Sowa, J.F. (1984). **Conceptual Structures: Information Processing in Mind and Machine.** Reading, Massachusetts: Addison-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.** BTRL, Ipswich: Martlesham Consultancy Services.

White, R.B. (1979). The consistency of the axiom of comprehension in the infinite-valued predicate logic of Lukasiewicz. **Journal of Philosophical Logic, 8,** 509-534.

Zadeh, L.A. (1965). Fuzzy sets. **Information and Control, 8,** 338-353.

Zadeh, L.A. (1973). Outline of a new approach to the analysis of complex systems and decision processes. **IEEE Transactions on Systems, Man and Cybernetics, 2,** 28-44.