Chapter 1

FOUNDATIONS OF LOGIC AND REASONING FOR KNOWLEDGE-BASED SYSTEMS

Brian R. Gaines and Mildred L. G. Shaw

ABSTRACT

This paper examines the foundations of logic and reasoning in future generation computing systems processing knowledge and interacting closely with people. It first provides a framework for the current research focus on knowledge-based systems in the context of the generational infrastructure of the computing industry. It then goes on to examine the role of computing in society and its implications for human-computer interaction. The nature of expert systems is then examined through examples demonstrating the need to be able to encode human processes of reasoning with imprecise rules and data. The role of standard logic as a foundation for arithmetic and its relevance as a foundation for knowledge processing are then examined. Psychological and nonstandard logical foundations for knowledge processing are discussed which allow unrestricted predication. Finally, the abstract forms of these foundations are presented as a calculus of distinctions.

TABLE OF CONTENTS

I. Introduction: Fifth Generation Computers and Expert Systems 4
- A. Discontinuities in Development 5
- B. Information is Not Knowledge 5

II. Paradigm Shifts in Information Technology 6
- A. Complexity Reduction 6
- B. Industrial Technology Imposes Simplicity 7
- C. People Cope with Residual Complexity 7
- D. Information and Knowledge 8

III. Expert Systems 9
- A. Meta-Knowledge 10
- B. Inducing Patterns in Information 10
- C. Imprecision in Expert Systems 11
- D. Information and Knowledge Processing 12

IV. Knowledge and Psychology as FGCS Technologies 13
- A. The Divergence of Psychology and Logic 13
- B. The Consequences of Antipsychologism in Logic 14

V. Psychological and Systemic Foundations 15
- A. Logics Based on Unrestricted Predication 15
- B. A Logic of Knowledge Systems 15
- C. Imprecision and Nontruth-Functionality 16
- D. An Epistemological Hierarchy 17
- E. A Knowledge Hierarchy 18

VI. Logics of Knowledge Systems 19
- A. Distinctions about the Origins of Distinctions 19
- B. Conversation Theory 20

C. Conceptual Dependencies 20

VII. Applications to Knowledge Engineering 20
A. Knowledge Engineering 22
B. Soft Systems Analysis 22
C. PLANET: a Knowledge Engineering System 24
1. Construct Elicitation 24
2. Single Construct System Analysis 24
3. Multiple Construct System Analysis 24
4. Database Administration 25
D. Applications to Expert System Development 25

VIII. Summary and Conclusions 25

Acknowledgments 26

References 26

I. INTRODUCTION: FIFTH GENERATION COMPUTERS AND EXPERT SYSTEMS

We are now in the fifth generation era of computing in which a transition is taking place from *information processing* to *knowledge processing*. The transition between each generation has involved the introduction of new technologies:[25,29]

1. The transition to the zeroth generation (1940 to 1948) was the introduction of *electronic device technology* with vacuum tubes replacing mechanical relays and allowing increased speed and reliability hardware operation.
2. The transition to the first generation (1948 to 1955) was the introduction of stored program control replacing manual patching and allowing instruction sets to be extended through subroutines to give a variety of *virtual machine architectures*.
3. The transition to the second generation (1956 to 1963) was the introduction of *problem-oriented languages* replacing manual translation of problem statements to machine code and allowing increased speed and reliability of software production.
4. The transition to the third generation (1964 to 1971) was the introduction of interactive, time-sharing systems replacing batch processing with direct *human-computer interaction* and allowing the development of personal support systems for problem solving and decision making.
5. The transition to the fourth generation (1972 to 1979) was the introduction of *knowledge-based systems* replacing information processing with knowledge processing and allowing problems to be represented in a form natural to people.
6. The transition to the fifth generation (1980 to 1987) was the introduction of *inductive-information systems* replacing the explicit expression of knowledge with its inference from case histories and allowing computers to acquire knowledge from experience as do people.

During the period of each generation, *research* activities are focused on the transitional technology of that generation; *product innovation* activities are focused on the transitional technologies of previous generations; and established *product lines* are drawn from the

technologies of two generations back. Thus, in the fifth generation era, machine learning and inductive inference have emerged again as major topics after neglect since the early 1960s; knowledge-based systems supported by, and supporting, effective human-computer interaction, high-level languages, machine architectures, and VLSI design are the primary objective of the Fifth Generation Computer System (FGCS) development program initiated by the Japanese;[52] and major product lines are typified by highly interactive integrated packages such as *Symphony* or *Framework* and highly interactive computers such as the *Macintosh*.

This analysis gives a rationale for the objectives of the FGCS program: (1) to provide knowledge processing rather than information processing and (2) to offer natural human-computer interfaces. These objectives have been accepted as realistic because of the successful development of *expert systems* (ES)[30,59] that encode human knowledge and cognitive skills for the computer and make them accessible to others. They are also an integral part of the developmental infrastructure of the computing industry.

A. Discontinuities in Development

The hardware, architecture, software, interaction, knowledge, induction, and sequence of transitions may be presented as the continuous evolution of an industry. However, the step to FGCS involves a transition that is different in kind and requires a more radical reappraisal than any before it. In the development of any major industry, there are periods of growth which consolidate existing technologies and extend them to a widening range of applications. During these periods, conceptual frameworks, design methodologies, and implementation techniques become well established, systematized, and widely known and accepted. However, there are also phases of discontinuity in which new technologies open up new areas of application that are qualitatively different from the old and require new modes of thinking. The conceptual frameworks, design methodologies, and implementation techniques after the discontinuity differ radically from those before it. During the transition, there is confusion as the older modes of thinking are applied to the new situation and the newer modes of thinking have to establish their credibility as preferred alternatives to those with a track record of past success.

Information technology has now come into a phase of discontinuity where the new modes of thinking are necessary. We have been through a phase of consolidation of techniques and expansion of applications in which computing hardware costs have dropped to very low levels and an immense range of systems and applications software has been made available. Information processing is now available in all those markets that may require it with few impediments of cost or availability of hardware and software. Even the limitation of the shortage of trained computer professionals is fast disappearing as a high proportion of students in all disciplines graduate with computing training and computers themselves become simpler to program. However, just as we are beginning to celebrate the arrival of the *information age*,[10] we are coming to realize that many of the predictions for that age are dependent on processing *knowledge* rather than information.

B. Information is Not Knowledge

It is not easy to give a simple and widely acceptable definition of the difference between knowledge and information, and this will be discussed in more detail later. However, an example will serve to illustrate some essential distinctions. *Encyclopedia Britannica* contains about 300M characters and a scientific book of 350 pages contains about 1M characters. With the availability of low-cost 140 Mbyte Winchester disks and text compression techniques, it is feasible to have the encyclopedia or some 300 books on-line to a personal computer for about $5,000 in hardware cost. The advent of low-cost optical disks will allow a scientist to have 2 Gbytes on line for this cost in a year or two: say, a general encyclopedia, a specialist one for his subject area, some 2,000 books and 30,000 papers. For most university

scientists, this would represent a complete information environment for the needs of their personal interaction with their discipline in research, teaching, and consultancy. If the numbers are too small, then project them a further 3 years into the future for another 10-fold increase in capacity.

We have passed, or are passing, the threshold of personal information availability requirements in terms of cost of storage. Speed of access is also no longer a problem since a fully inverted keyword index to the material can be provided at the cost of doubling the storage. The texts themselves will be made available at a price comparable to those of books once publishers can see a market of sufficient size, and prices should decline since the production, distribution, and inventory costs of optical disks are much lower than those of books. Compared with previous media, we can see that computer-based information technology has major advantages. The scientist has in a compact system a personal library that may be accessed easily through a complete index. It couples simply into his word processor making referencing and quoting other work very easy.

However, apart from ease of storage and access, what advantages does our computer-based system have over a library? All it does is mimic the previous medium of books and journals at a lower cost, in a smaller space, and with greater speed of access. These are significant advances, but they do not represent a qualitative change. The burden of using the information, recognizing patterns in it, linking it together to support an argument, and searching it for solutions to a problem all still reside with the scientist. It is the user who creates knowledge from the information.

It would be reasonable to suppose that the additional features necessary to make the system itself a knowledge processor could be developed using the same conceptual frameworks, design methodologies, and implementation techniques that led to the present generation of information processors. However, there are some fundamental obstacles to this: conceptually, that we do not understand the processes of inference that a person uses in recognizing patterns, generating arguments, and solving problems; and technically, that these processes seem to involve computations of a complexity many orders of magnitude beyond those possible with current computers. In terms of our text storage example, the information needs to come out of a passive storage medium into an active processing medium that is accessing substantial parts of it simultaneously through an indexing framework that is itself active and conceptual rather than textual.

The conflict between old and new in a phase of discontinuity forces a resolution, or synthesis, in which the fundamentals common to both are reexamined. There are presuppositions to the old solutions that are no longer valid. They need to be made overt if the new solutions are to be formalized and understood.

II. PARADIGM SHIFTS IN INFORMATION TECHNOLOGY

The previous section introduces the notion that in developing knowledge systems the computer industry is making a change in technology that is discontinuous with that preceeding it. There has been a *paradigm shift*[40] that shows up as a change in the conceptual framework for all aspects of computing. However, such discontinuities do not occur for no reason, as causes of disruption; rather they are the effects of deeper underlying causes such as social need. There is much talk about the impact of information technology on society. However, we can also see information technology as generated by the needs of an increasingly overcrowded, complex, interdependent world with limited resources.

A. Complexity Reduction

The German sociologist Luhmann[42] encapsulates the dynamics of society by postulating *complexity-reduction* as the fundamental motivation for all our social institutions:

The world is overwhelmingly complex for every kind of real system... Its possibilities exceed those to which the system has the capacity to respond. A system locates itself in a selectively constituted 'environment' and will disintegrate in the case of disjunction between environment and 'world'. Human beings, however, and they alone, are conscious of the world's complexity and therefore of the possibility of selecting their environment — something which poses fundamental questions of self-preservation. Man has the capacity to comprehend the world, can see alternatives, possibilities, can realize his own ignorance, and can perceive himself as one who must make decisions.[42]

Luhmann's model seems to underlie De Bono's[9] optimism about the role of computers:

By great good fortune, and just in time, we have to hand a device that can rescue us from the mass of complexity. That device is the computer. The computer will be to the organization revolution what steam power was to the industrial revolution. The computer can extend our organizing power in the same way as steam extended muscle power . . . Of course we have to ensure that the result is more human rather than less human. Similarly we have to use the computer to reduce complexity rather than to increase complexity, by making it possible to cope with increased complexity.[9]

B. Industrial Technology Imposes Simplicity

The paradigm for industrial technology in the past has been to create Luhmann's simplified environment for man by dealing with the complexity of the world through removing it. The man-made environment is simple and regular and capable of being modeled easily. Where it interacts with nature, as in agriculture, the natural world is simplified by saturating it with variety-reducing mechanisms, such as weed killers and pesticides. The technology of the past has made simplistic scientific models true by constraining natural phenomena to fit the models. Industrial technology has made much of the world simple, homogenous, and comprehensible. It has minimized the knowledge necessary to operate in a complex world by simplifying the world. This has been an effective way to operate, but not an efficient one in the preservation of nonrenewable resources.

De Bono's optimism suggests that information technology is capable of extending our organizing capabilities to cope with a more complex world. This would be a reversal of the objectives of previous industrial technology and hence a major paradigm shift. Computers can be used to control the world and make it simple or they may be used to aid our interaction with a complex world. The second alternative is becoming important to us as the consequences of industrialization in destroying natural resources are realized.[54] It is becoming important to accept and preserve the variety of the natural world and use it to extract sufficient resources for society's needs while husbanding long-term resources.

C. People Cope with Residual Complexity

The major paradigm shift is thus from imposing simplicity on the world to coping with its complexity. A secondary paradigm shift results from this in that the previous system for coping with residual complexity has been people. When technological simplification breaks down, people are used to solve the remaining problems. Automatic machines need human operators to cope with broken tools, misplaced parts, and other contingencies outside the well-defined environment for which the machine has been designed. Some industrial processes, such as pig iron smelters and lime kilns, are not amenable to automatic control, but require human operators. The shift in our conceptual frameworks has been to realize that we may be able to automate such processes not by modeling the complex world, but instead by modeling the way in which people cope with it. Expert systems mimic specific problem-solving processes of people without analyzing the underlying basis for the success of these processes.

ES represent a major shift in the paradigms of science and technology because they allow control without explicit understanding. The classical control paradigm is to develop a model of the plant and then to use this model to calculate an optimal control strategy. This works where the plant can be engineered to have a simple model amenable to analytic techniques.

Even in this case there is no guarantee that an understandable plant is an efficient one in terms of yield. The paradigm increasingly falters as the plant involves natural processes not engineered for simplicity in modeling. It is symptomatic of our lack of respect for human capabilities that we failed to notice for so long that people could operate systems that they did not fully understand and that this capability could be transferred to machines by modeling the operator rather than the plant.

D. Information and Knowledge

We can now bring together the discussion of the last two sections to give a succinct account of the relation between information and knowledge in FGCS.

1. The multi-Gbyte information system described in the previous section has the capacity to encode a substantial part of the complexity of the world, sufficient say to contain that information used by a scientist during his lifetime. However, it only stores and indexes that information. It cannot process it as knowledge to recognize patterns, support arguments, and solve problems.
2. Expert systems encode some basic aspects of people's capabilities to solve problems and support arguments in very limited contexts, i.e., involving a small amount of information. They are primitive knowledge processors in low-information environments.
3. Fifth generation developments will enhance the capabilities of expert systems and exploit them in commercial situations where knowledge processing is highly significant and the information is low.
4. Once the immediate potential of simple expert systems has been realized, there will be a shift towards integrating the low-cost, high-capacity information storage and processing capabilities of fourth generation computing systems with the knowledge processing of FGCS. This will involve new hardware capable of actively modeling large information structures and new software capable of processing knowledge on a large scale.

With such *sixth generation* systems, the paradigm shifts will be complete. We shall be able to represent the complexity of the world as information and process it as knowledge. To complete this shift, we need to go back to fundamentals and understand the nature of knowledge. In particular, we need to understand the basis of the technological paradigm, why it has operated successfully, and why it has to change at the next stage. We also need to understand the basis of human knowledge processing and reverse the logic of the past that where people do not conform to our normative models of scientific reasoning that it is they who are in error.

It is interesting to link this discussion with Drucker's remarks comparing the development of the electricity industry with that of the information industry and making the analogy that "information is energy for mind work".[11] He remarks that if it had been fashionable to speak of generations of generators, then there "would have been a 'fifth-generation' or 'sixth-generation' generator before there was any widespread use of electric power". It was the invention of the light bulb that led to the use of electricity as a universal form of energy. Fifth and sixth generation developments in knowledge engineering will play the role analogous to that of the light bulb and bring into widespread use the "mind energy" of the new medium of computing.[21,28]

The remainder of this paper examines fundamental issues relating to the nature of knowledge and links them to information and knowledge processing. The next section reviews some key developments in knowledge-based systems and the following sections develop foundations for them.

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.

FIGURE 1. Rule from fuzzy logic controller.

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.

FIGURE 2. Rules from a lime kiln operator's manual.

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.

FIGURE 3. A MYCIN rule.

III. EXPERT SYSTEMS

One of the first expert system developments was that of Mamdani and Assilian[43] in 1974 that took human knowledge expressed verbally and encoded it directly as computer programs which acted on the environment. This work was undertaken as part of a study of machine learning in process control and the system being controlled was a small steam engine. The verbal rules were of the form shown in Figure 1.

Priming the learning machine by such rules was first introduced because learning took too long. It was expected that the rules would act like verbal instructions to a human operator and aid the machine in setting up an initial approximate control strategy. However, it was found 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. Additionally, the learning machine then proceeded only to learn less effective strategies. Hence, interest switched to the process whereby human expression of verbal rules that appear vague can lead to highly effective control strategies, i.e., to the development of expert systems.

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 2 is an excerpt from a training manual for lime kiln operators.[55] 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.

In parallel with the controller development, other expert systems were being developed for completely different domains. MYCIN[75] is a medical diagnosis expert system which aids a clinician to act as a consultant on infectious diseases. It uses rules of the form shown in Figure 3 to acquire information, make inferences from it, and recommend therapy.

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.

FIGURE 4. A TEIRESIAS metarule.

A. Meta-Knowledge

The rules of Figures 1 through 3 are expressions of inferences from data about the world to conclusions about it, a basic level of knowledge. It was rapidly discovered, however, that similar rules could express inferences at a higher level about the nature of these basic inferences. MYCIN worked well when the rules were complete and correct, but it was difficult to set up the rules initially. It was also difficult to trace the rules causing errors in the deductions of MYCIN if the clinician felt that something had gone wrong. TEIRESIAS[8] 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*.

Modifying and adding rules through interaction with TEIRESIAS demonstrates the dynamic nature of ES. They have often been presented in the literature as means of capturing expert knowledge as if that knowledge was a static, finished product, universal in nature and owned by one person. In practice, knowledge is fluid, situational, and generally distributed across groups of people. The initial priming of an ES with the knowledge of one or more experts is a way of getting started. Metarules cannot be effective unless the basic vocabulary and patterns of inference have been at least approximately established. However, an ES then becomes an entity that grows with the knowledge and experience of its users. It is a support tool for their knowledge acquisition and refinement as well as a way of making it more accessible to themselves and others.

Thus, the expression of knowledge as inference rules generalizes from basic factual knowledge to that about the nature of knowledge itself and enables changes to be made in ES through discussion and learning. The present generation of ES has already gone beyond being a static recording of the knowledge of experts to being a technique for developing and extending expertise and acquiring knowledge.

B. Inducing Patterns in Information

The use of metarules allows ES to be programmed interactively by their users. From one perspective, TEIRESIAS can be seen as an important development in automatic programming. From another it can be seen as a way in which a machine acquires knowledge through interaction with a person. Metarules may also be used to implement knowledge acquisition procedures in which machines acquire knowledge directly through their interaction with the environment, i.e., a process of *machine learning*.[49]

Inductive inference of knowledge through machine learning procedures is a significant alternative to direct access to the verbal knowledge of experts and comes closer to capturing the knowledge of main interest. We want to know the basis for what the experts *do* rather than what they *say* they do. Michalski and Chilausky[48] have reported an interesting comparison of two ES for plant disease diagnosis, one constructed directly from the expert's stated rules and the other constructed automatically by an algorithm for logical induction of

If time is small and error is negative big
  Then desired change is big.
If time is big and error is positive zero
  Then desired change is zero.

FIGURE 5. Rules from fuzzy logic learning controller.

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.

FIGURE 6. Two rules of high generality from AM.

structure from behavior. The ES with induced rules had a better performance in diagnosis and the expert agreed that its rules, although not obvious at first, were better than his own.

Inductive inference may also be used without an expert to acquire knowledge through experience. Mamdani et al.[44] found that learning could be introduced effectively in the steam engine controller through metarules that expressed the way in which the basic rules should be changed as a result of performance feedback. The learning level of their controller operator on rules of the form is shown in Figure 5.

The highest level of generality in inductive inference is shown in the rules of Davis and Lenat's AM[8] 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. Figure 6 shows two of the few hundred rules used by AM. Most of them are sufficiently general to apply to other situations. In an organizational context, for example, Rule 19 indicates that a committee which has decision-making powers that do not apply to its subcommittees below a certain size has an interesting property (that of requiring a quorum). Its subcommittees also have an interesting property (that of being quorate).

AM has been criticized for some of the shortcuts taken in its implementation, which casts doubts on the conclusions that may be drawn about the power and generality of the mechanisms used.[61] Lenat and Brown[41] have analyzed its success in mathematics and failure elsewhere and suggest that AM had easy access to much of what it appeared to discover through the natural expression of certain mathematical conjectures in LISP. However, the program serves to illustrate that even very high-level knowledge processing may be programmed using current techniques provided the space of information being processed is very small.

C. Imprecision in Expert Systems

It will be noted that the controller rules in Figures 1 and 5 are imprecise, using terms such as "medium" and "big". Similarly, the MYCIN rule of Figure 2 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 inferences from them using Zadeh's[81]

fuzzy sets theory and its application to decision processes.[82] Shortliffe[75] uses a similar method in MYCIN where the ".7" in the rule is a truth value in a multivalued logic, and Lenat and Brown[41] use 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.[23]

To see the significance of 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 kg at the rear entrance of Building 6A on February 15th 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, (1) 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: (2) 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.

D. Information and Knowledge Processing

The examples in this section illustrate how 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 are providing facilities to emulate the capabilities of a skilled colleague and advisor. However, they are capable of operating within only small domains of information and their operation is heuristic without adequate formal foundations.

IV. KNOWLEDGE AND PSYCHOLOGY AS FGCS TECHNOLOGIES

The dependence of FGCS developments on increased understanding of knowledge and human reasoning has been recognized by those concerned with the Japanese program, and the question, "Are you saying that the design of the fifth-generation may be modeled by learning more about the human thinking process?", has been answered: "Yes, we should have more research on human thinking processes, but we already have some basic structures. For more than 2000 years man has tried to find the basic operation of thinking and has established logic. The result is not necessarily sufficient; it's just the one that mankind found. At present we have only one solution — a system like predicate calculus. It is rather similar to the way man thinks. But we need more research. What, really, is going on in our brain? It's a rather difficult problem."[17]

This leads to a pragmatic approach to developing FGCS, basing them on the best available foundations for knowledge and psychology, even if these are inadequate. Technology is often in advance of science and the pressure to develop FGCS may be just what is needed to advance the sciences of knowledge and psychology. This approach leads to questions: (1) how adequate is predicate calculus as a science of knowledge and (2) how adequate is logic as a foundation for psychology?

A. The Divergence of Psychology and Logic

A philosophical dictionary definition of "knowledge" is that it is "apprehended truth" and that "certain knowledge is more than opinion, less than truth".[62] Knowledge stands between opinion which is in the domain of psychology (the study of mental phenomena) and truth which is the domain of logic (the study of truth preserving inferences); the underlying discipline is epistemology (the study of truth ascription). It would be reasonable to suppose that psychology and logic would be closely related aspects of a science of knowledge. Until recent times, this was so, but in the late 19th century, psychology and logic diverged and an attempt was made to set up a system of logic that stood alone, independent of any processes of human reasoning. The divergence stemmed from attempts to establish formal foundations for mathematics, particularly those of Frege[15] and Husserl.[34]

The stimulus for Frege's work was to reestablish irrevocably the position that mathematics was *a priori* (independent of experience) and analytic (based on inferences in which the conclusion is part of the premise), in contradiction to Mill's[50] position that arithmetic is not *a priori*, but empirical and induced from experience, and Kant's[36] position that it is *a priori*, but synthetic and made known through experience. Husserl's work shows the arguments most strongly because in response to Frege's[16] critical review of his 1891 book, he became concerned with the role of psychology in logic and formulated the basis of pure logic clearly demarcated from processes of the mind.[35] In his discussion and the distinctions he makes, Husserl generates a model of relationships between logic, psychology, and knowledge which underlies the development of modern epistemology.[51] Similarly, Frege generated a notion of pure logical inference that underlies the development of modern formal logic.[39]

These classical discussions may appear remote from modern computing and artificial intelligence (AI). However, Lenat and Brown's paper[41] puts forward the Kantian argument exactly. The problem is why AM succeeded in the first place and why it is so difficult to use the same paradigm to discover new heuristics. Their conclusion is that AM's success in making conjectures in basic mathematics was because set- and number-theoretic results

are simply stated in LISP. The conjectures were to a large extent available *a priori*, but were *synthetic* because AM needed its empirical investigations to trigger off their discovery. However, while the metarules of logic and arithmetic are embedded in LISP, those about discovery heuristics are not and could not be thus triggered.

This analysis of AM not only demonstrates the Kantian argument, but also suggests a line of attack that investigates the discovery heuristics as primitives in some operational calculus — a formal epistemology. Many of the distinctions developed in philosophy in relation to reasoning and the operation of the mind have technological implications when applied to AI.

B. The Consequences of Antipsychologism in Logic

The antipsychologistic basis of modern logic is natural in the light of the objective of formalizing arithmetic as a demonstration that it is *a priori* and synthetic. It rejects as irrelevant the questions: (1) how do we come to know arithmetic and (2) why is arithmetic instantiated in real-world phenomena?, although a variety of answers to these questions are part of Frege's and Husserl's discussions. It also does not take into account the distinction between what is known and what might be inferred from this; Wittgenstein notes: "If I know an object I also know all its possible occurrences in states of affairs,"[80] drawing attention to the gap between formal logic, where all consequences are "known", and human reasoning, where they are not. It raises the questions: (3) what is the significance of not knowing all the logical consequences of a state of affairs and (4) how do we come to know particular consequences?

All of these questions are highly significant for knowledge-based systems yet outside the framework of standard formal logic. Questions 1 and 4 may be regarded as primarily within the domain of psychology and Question 2 as within the domain of philosophy. The significance of Question 3 was not realized until the notion of *computational complexity* was developed and formalized. The superexponential combinatorial explosion of any consequence derivation system creates a boundary that is real and extremely "hard"; the notion of a hard proof is not just psychological, a result of human frailty, but is systemic, a property of the problem for any derivation system.

The boundary created by computational complexity means that existence proofs relating to inference mean little in practical terms. It is not sufficient to know that solutions are eventually accessible; we need to know they can be reached using reasonable resources. The knowledge representation and inference schema used must make the types of inference required easy. This is an extralogical requirement, but a critical one in practice. Hardware parallelism and speed increase cannot be used to overcome the problem because superexponential growth overwhelms even large multiplicative factors. A danger in developing FGCS is that any logic we build into the hardware may prove inappropriate to emulate human inference.

These considerations suggest that we should regard with some suspicion formal logics developed as foundations for mathematics and deliberately disregarding psychological considerations. It could be that mathematics epitomizes human intellectual achievement and standard predicate calculus (SPC) captures all that is significant. Even this is suspect since we know of many problems with SPC: in the foundations of set theory;[46] relevance of implication;[1] ontology of existence;[64] and extension beyond first order.[39]

However, while one may criticize SPC for its known faults, it is not easy to suggest an alternative system for knowledge representation and inference. There are many studies of human reasoning and alternative logics, but there has not been the intensive study of a coherent system for knowledge in general as there has been for that subset which is mathematics. Such a study has its place in the development of FGCS, and in the remainder of this paper we survey some possible bases for such a system with strong logical and psychological foundations.

V. PSYCHOLOGICAL AND SYSTEMIC FOUNDATIONS

Kelly's *Personal Construct Psychology* (PCP) develops a complete theory of the normal and abnormal, which has strong systemic foundations[37] and has been operationalized through computer programs.[66] 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", and emphasizes the use of constructs 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".[37]

When Kelly developed his theory, there were neither the mathematics to formalize it nor the computer technology to operationalize it. Despite this, PCP has developed to become a major psychological theory and to provide a variety of techniques used in clinical, educational, and managerial psychology.[45,67] PCP has strong links with Brown's logical system based 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.[5]

A. Logics Based on Unrestricted Predication

Kelly's psychology[37] and Brown's logic[5] both take as primitive an axiom of unrestricted predication. In set-theoretic terms, they take *comprehension*[14] as *a priori* and derive all other concepts from this. This is not the standard and familiar approach, although logicians such as Fitch[13] have argued for it, and this makes it misleading to translate the Kelly/Brown 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 prior knowledge. 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[22] that what is distinguished is a *system* and that this notion defines a system. 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.

B. A Logic of Knowledge Systems

Kelly's personal construct psychology[37] and Brown's calculus of distinctions[5] both point to a knowledge representation system based on unrestricted predication as a primitive. The significance of such an axiom 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[15] 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 prop-

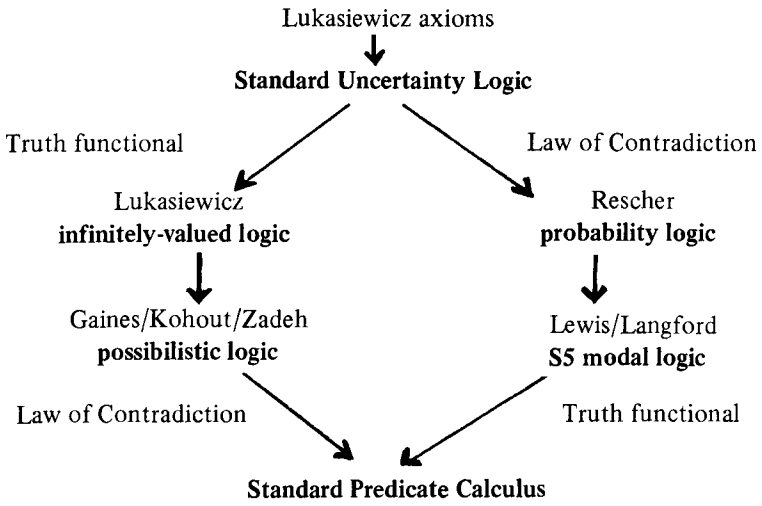

FIGURE 7. Logics of uncertainty generated from calculus of distinctions.

osition. Maydole[47] proved that most contenders for an alternative logic led to inconsistencies in set theory, and White[79] showed that Lukasiewicz's 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[81] took as a basis for fuzzy sets theory, Mamdani and Assilian[43] used in their expert system for linguistic control, and Shortliffe[75] discussed as a formal basis for MYCIN. The links from Brown's calculus of distinctions to Lukasiewicz logic have been established: Varela[77] used the calculus to analyze self-referential systems using a construction based on Russell's paradox which gave a three-valued logic; Gaines[19] used a recursive form of this construction to generate the infinite valued Lukasiewicz logic; and Gaines[24] shows that the Lukasiewicz axioms encompass modal, probability, and fuzzy logics.

The significance of this construction is that the axioms are taken first as defining a nontruth-functional *standard uncertainty logic*; fuzzy logic is then derived by adding the constraint that the logic be truth functional, and probability logic is derived by the alternative constraint that the law of contradiction apply. Figure 7 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.

C. Imprecision and Nontruth-Functionality

The need for a system adequate to represent knowledge to encompass both imprecision and nontruth-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.

Nontruth-functionality arises because it is not possible to 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 0% if we have prearranged, to 50% if the 2 events are not causally connected, to 100% if I do not want to talk to you.

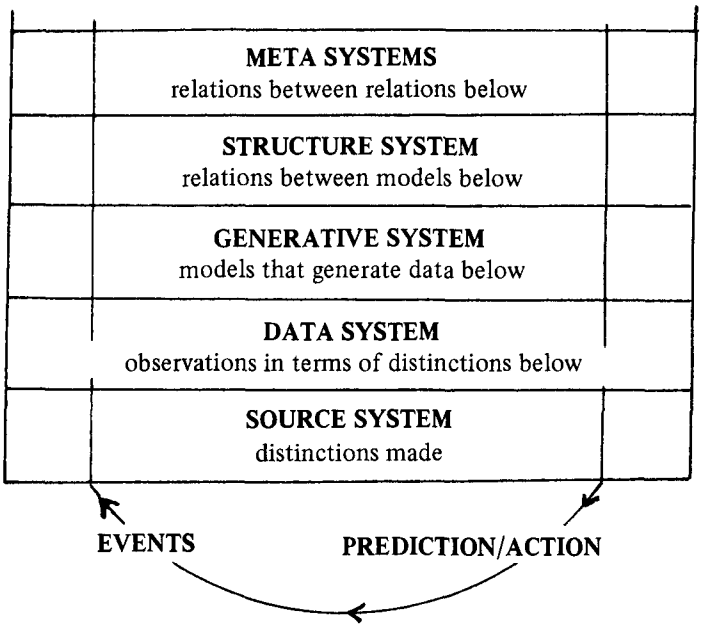

FIGURE 8. Epistemological hierarchy of a personal scientist.

It is possible to compute the constraints upon the result, but not a single-valued result. Quinlan[57] has shown that the constraint systems of such nontruth-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.

D. An Epistemological Hierarchy

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[31] category-theoretic formulation of fuzzy sets, Ralescu's[58] application of these to a general theory of modeling, and the development by Goguen et al.[32] of *OBJ*, a category-theoretic programming language. The role of distinctions at the base level of all models is evident in Klir's[38] hierarchy of modeling shown in Figure 8. The loop from events through distinctions up through the modeling hierarchy and then down again to predictions and actions characterizes what Shaw[66] 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[20] has shown that it leads to general knowledge acquisition algorithms. 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

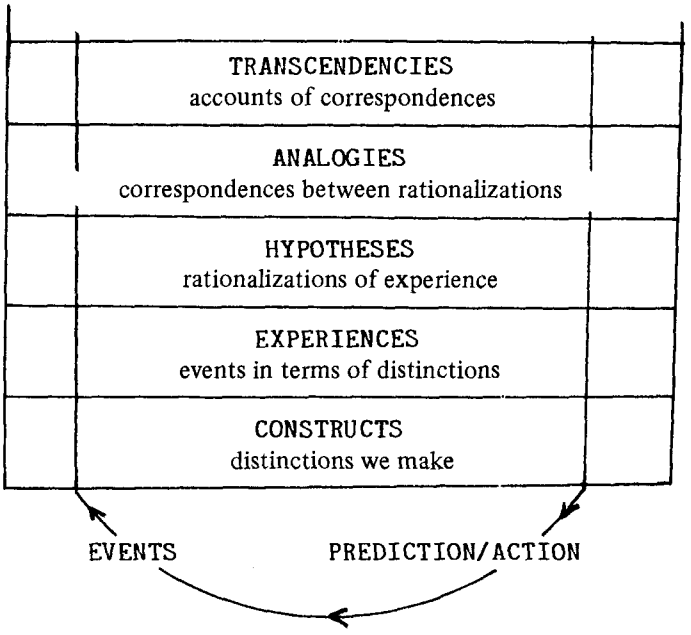

FIGURE 9. Knowledge hierarchy of a personal scientist.

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 whether a level accounts for the information flowing through it and we may term this distinction one of *surprise*[20] after the economist Shackle.[65] 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.[18] 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, to inadequate approximation to predict events; at the next, to inadequate simplicity to explain events; and at the next, to inadequate comprehensiveness to account for events. AMs "interest"[8] may be explicated as surprise generated when rules such as property inheritance are violated.

E. A Knowledge Hierarchy

The systemic hierarchy based on a calculus of distinctions has an analog based on personal constructs[27] shown in Figure 9. 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.

Kelly[37] places the major emphasis of his work on the notion of *constructive alternativism*, 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 vs. realism; it is the distinctions which a philosopher makes that determine his stance and these can be analyzed 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 preconceived 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.[70] Problems arise because of the individualistic nature of construct systems. A person may be able to use his construction system while 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 construct 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.

VI. LOGICS OF KNOWLEDGE SYSTEMS

A logical system is a means for making inferences that derive true statements (conclusions) from true statements (premises). Thus, to define a logic, we need to establish what are its statements, its truth criteria, and its inference procedures. An axiomatically defined logic has a syntax stated to generate well-formed statements, axioms stated to be primitive truths, and inference rules stated to derive true conclusions from true premises.

Logics may be established within the hierarchy of distinctions by taking a statement to be a distinction made, truth to be a distinction about statements, and inference rules to be a distinction about what may be inferred from a statement. The hierarchy is naturally recursive in that it results from making distinctions about the system whereby we acquire and process knowledge. Different forms of logic and different notions of truth and inference arise as we make different forms of basic distinction.

A. Distinctions about the Origins of Distinctions

Some very basic distinctions about the origins of distinctions at the highest level of the hierarchy lead to the main classifications of philosophical logic. If we make no distinctions about the origins of distinctions, then truth is what is asserted to be true and inferences are legitimated by making them. This leads either to a anarchistic theory in which anything goes[12] or to a pragmatic theory of reasoning and a utilitarian theory of truth based on an assumption of rationality that we distinguish something as true because it is useful to do so and reason in such a way that it works for our purposes. Three distinctions about the origins of distinctions underly the deeper theories of reasoning and truth:[60]

1. If we distinguish those distinctions that are *necessary* in that they seem to impose themselves upon us, then we generate what Popper[56] terms World 1, that of physical reality with a *correspondence* theory of truth and *causal* theory of inference.
2. If we distinguish those distinctions that are made by *agents* in that they seem to be freely chosen, then we generate what Popper terms World 2, that of subjective experience with a *performative* theory of truth and *conventionalist* theory of inference.
3. If we distinguish those distinctions that are *independent* in that they seem to exist in their own right, then we generate what Popper terms World 3, that of statements in themselves with a *coherence* theory of truth and *structuralist* theory of inference.

B. Conversation Theory

Figure 10 shows how these distinctions about the origins of distinctions lead to notions of worlds, truth, and inference. Considering entities as coherent fragments of the worlds and individuals as autonomous entities leads to the notions of Pask's[53] *conversation theory* which has been taken as a basis for ES development.[7] Pask's concepts of conversations and individuals are sufficiently widened for the theory to encompass all aspects of knowledge acquisition and processing. Conversations with nature model experience of the physical world and conversations with oneself model the restructuring of information at one level into knowledge at a higher level. The definitions of entities and individuals can be formalized by noting that: an entity is a distinguished part of a world; a coherent entity is one that is relatively stable in that the relations of its parts one with another are more stable than with those outside it; an autonomous entity is one which is relatively uncontrollable in that the relations of its parts to those outside it are less stable than the relations of parts outside it with one another. Stability can be formalized in relation to pairs of distinctions in opposition except for unstable entities, and opposition can be formalized as pairs of distinctions distinguished in that they never both mark the same entity.

C. Conceptual Dependencies

These additional distinctions generate notions of time and the various relationships between objects, people, and ideas that Schank has taken as a basis for conceptual dependency theory.[63] Each distinction defines a subworld of space, time, social relations, and so on that inherits systems of truth and inference from its parent world and adds further constraints to them. In this way, various modal logics arise of spatiality, temporality, morality, and so on.

The basic axioms of inference will usually start with requirements that it is transitive, so that if A implies B and B implies C, then A implies C. Other axioms such as A implies A, A implies that B implies A, and so on will be added as the semantics of the situation requires. The standard uncertainty logic described previously that results from a nontruth-functional reading of the axioms for Lukasiewicz's infinitely valued logic seems a suitable starting point for most situations since its semantics are extremely weak, yet its inference power is adequate for many of the procedures used in expert systems.[24]

VII. APPLICATIONS TO KNOWLEDGE ENGINEERING

The previous sections have concentrated on foundations and theory in presenting the problems of knowledge-based systems and in developing a logic for such systems through the calculus of distinctions. However, returning to Drucker's analogy, we are expected fifth and sixth generation computing to result in the equivalent of a lamp bulb to make widespread use of the "mind energy" available through computing. What are the signs that such a practical component for knowledge engineering is becoming available?

The developments discussed in Section III provide some evidence of the progression

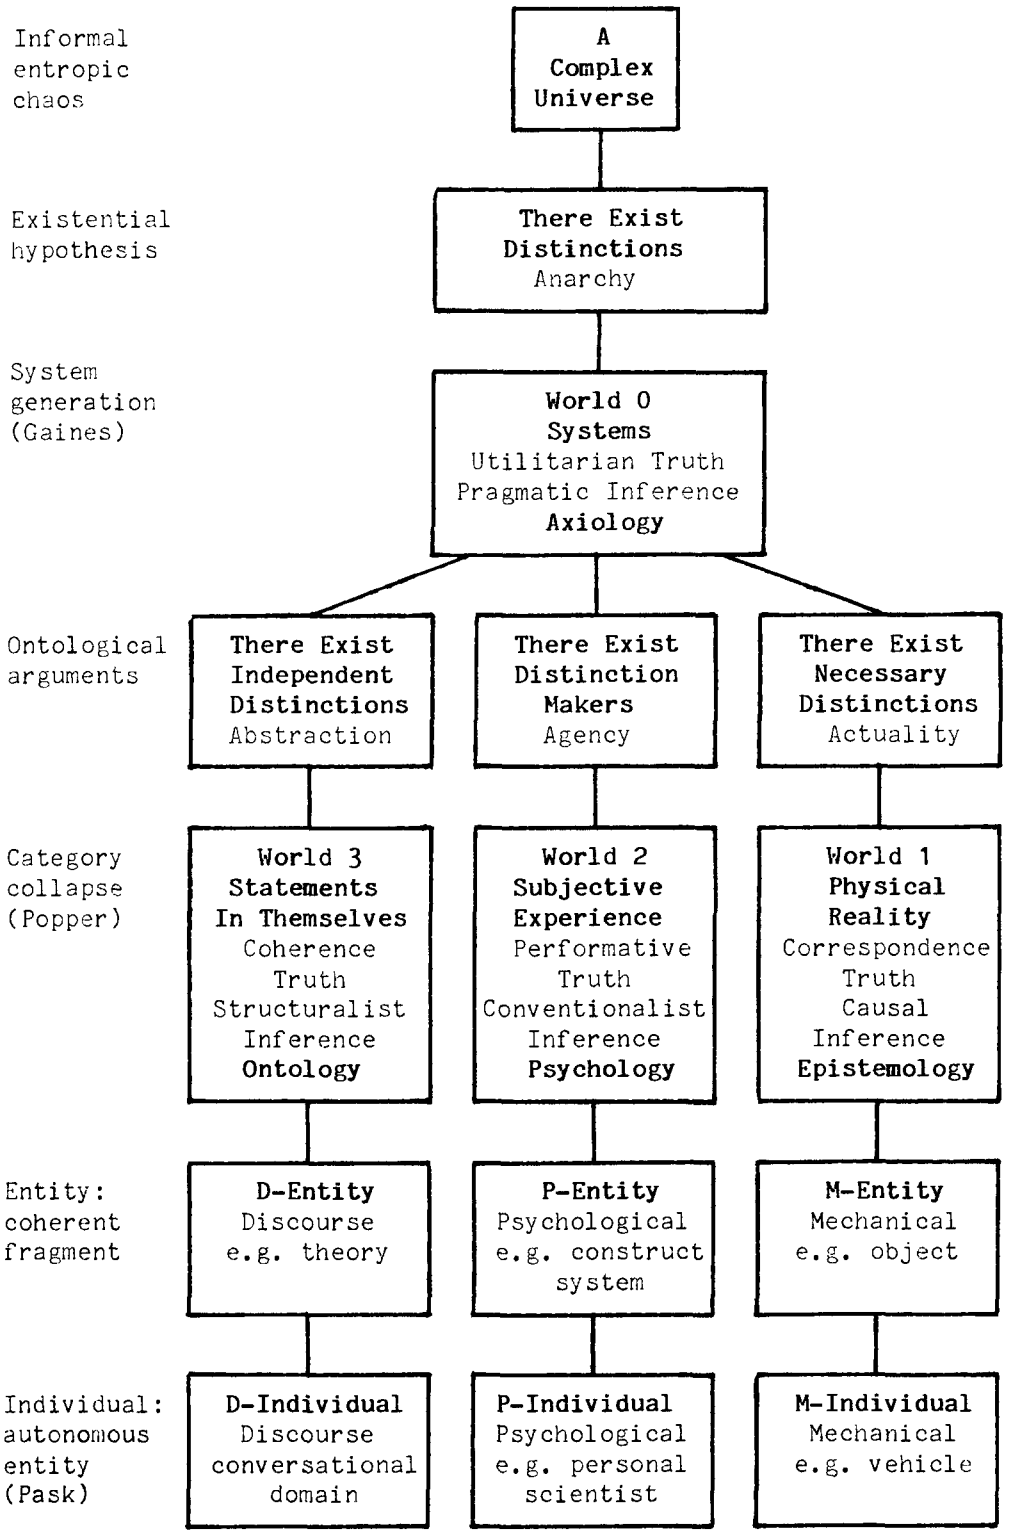

FIGURE 10. Distinctions about the origins of distinctions generating worlds, notions of truth and inference, conversation theory, and notions of entities and individuals.

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 11. Stages in soft systems methodology.

towards an applicable technology. The fuzzy controller shows that expert rules expressed verbally can be encoded simply and directly for reproduction by a computer. MYCIN shows that this can be done in the domain of cognitive skills. TEIRESIAS shows that metalevel knowledge can be used to debug and extend such expert systems through discussion with experts and users. The inductive inference results show that meta-level knowledge may be used to acquire domain knowledge either by examining expert behavior or directly from experience.

A. Knowledge Engineering

As we have noted, none of the techniques for knowledge processing, such as fuzzy linguistic translation, meta-rules, and inductive inference, can be used unless the basic distinctions being used, the expert's constructs, have already been elicited. This problem of elicitation has been a major limitation 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.[33]

The problem of knowledge elicitation from a skilled person is well known in the literature of psychology. Bainbridge[2] 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, e.g., through interviewing techniques that attempt to by-pass cognitive defenses, including those resulting from automization of skilled behavior. Welbank[78] has reviewed many of these techniques in the context of knowledge engineering. All of them may be seen as involving some pathway to the constructs of the expert as noted in Section V.D.

B. 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 dialogue between expert and knowledge engineer. Checkland[6] 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 as shown in Figure 11.

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 analyzing

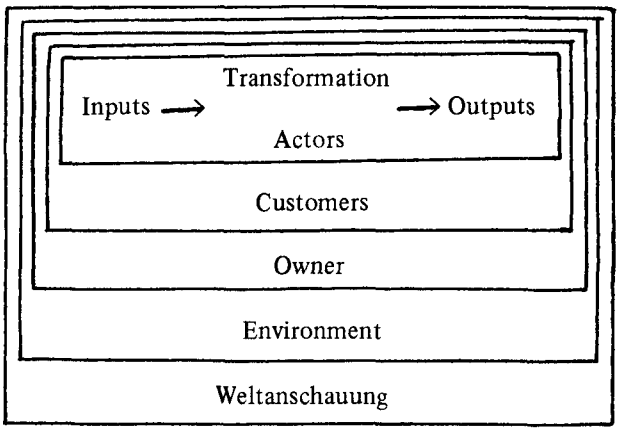

FIGURE 12. Root definitions of relevant systems.

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 the nested set shown in Figure 12 to which Checkland gives the mnemonic CATWOE.

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[3] 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.[73]

In terms of Figure 10, Checkland's CATWOE arises from the interative application of distinctions of reality and agency in turn: the *Weltanschauung* is an agent creating an *Environment* which forms the reality for the *Owner* who is an agent creating a *System* which forms the reality for the *Actors* who are agents creating a *Transformation* which forms the reality for the *Customers*. Note how the *System* is at the center with three levels on either side of it. This is a systemic basis for Checkland's distinctions. There is also reality feedback in that the transformation is part of the reality of the actors and hence part of the system; the system is part of the reality of the owner and hence part of the environment; and the environment is part of the reality of the weltanschuung.

Applying the CATWOE 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 *Weltanschauung* 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.[71] We have shown elsewhere how it accounts for the way in which communities are formed by individuals and the way in which social norms in communities also act to form individuals.[72] The individual can be regarded as a cross section of the epistemological hierarchy of the community in which he is operating.[69] These considerations are those of Pask's *conversation theory*,[53] particularly the notion of expertise distributed across a multiperson *P-Individual*, and it is notable that this theory is being used as a foundation for the development of expert systems.[7]

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.

C. PLANET: a Knowledge Engineering System

The program system that we have developed for doing soft systems analysis from a variety of viewpoints and then comparing and contrasting the systems elicited is called *PLANET* (Personal Learning, Analysis, Negotiation and Elicitation Techniques).[26,68] 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 construct system elicitation and analysis and divide naturally into four groups.

1. 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.

2. Single Construct System Analysis

FOCUS provides a hierarchical clustering of an expert's construct system that preserves the data elicited from him so that the sources of the analysis are evident and can be discussed. *PRINGRID* is a nonhierarchical cluster analysis based on principal components that can be used to gauge the major dimensions along which distinctions are being made. *ENTAIL* is a logical analysis of the construct system taking the expert's distinctions to be fuzzy predicates. 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.

3. Multiple Construct System Analysis

MINUS compares two construct systems based on the same elements and constructs to highlight the differences between them. *CORE* provides an interactive comparison of such systems to determine where there are differences and extract invariancies, e.g., over time or over a group of people. *SOCIOGRIDS* analyzes the total construct structure generated by a number of people construing the same elements from their individual viewpoints. It produces *socionets* showing the relationships between individual construct systems within the group and a *mode* system reflecting the constructs shared across the group.

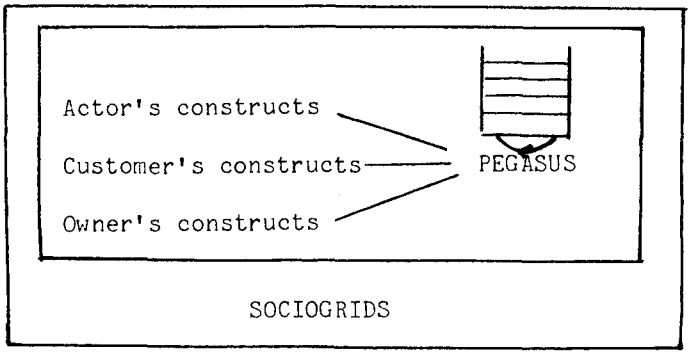

FIGURE 13. PLANET automating soft system analysis.

4. Database Administration

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

D. Applications to Expert System Development

As shown in Figure 13, PLANET provides a set of tools for the knowledge engineer enabling him to operate at Stage 3 of Checkland's soft systems analysis semiautomatically. 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 validity of knowledge elicitation with PLANET has been studied using a situation for which a set of expert constructs exists that is known to be effective. The BIAIT methodology for analyzing business record-keeping needs uses seven constructs to determine the types of transactions carried out by a business and infers the necessary records from the positioning of the business on these constructs.[76] Studies using PLANET with accountants, businessmen, and accounting students have shown that the BIAIT constructs may be elicited from those who do not know this methodology.[70,74]

Boose[4] has used an interactive system based on construct elicitation and analysis procedures derived from PLANET as a knowledge elicitation front-end to a variety of ES shells. He reports that his *Expertise Transfer System* (ETS) has been used to prototype over 100 ES at Boeing Computer Services and that the time taken has been hours compared with months for other nonautomated methodologies. ETS does not produce complete, large-scale ES, but it produces usable initial systems that can be developed further using the metarule tools of the ES shells.

VIII. SUMMARY AND CONCLUSIONS

This paper has argued that we are in a phase of discontinuity in the development of computing systems in which it is necessary to go back to fundamentals. We have developed information-processing systems with the storage capacity to match the complexity of a substantial part of the world and this technology is well founded. We are beginning to develop knowledge-processing systems that can mimic people in coping with the complexity of the world, but they can operate only in small information domains and the underlying knowledge technology is not well founded.

The dependence of FGCS on knowledge processing, coupled with the assumption that logic programming will be a key technology, suggests that we should examine the rationale behind standard logic. Modern formal logic was developed as a foundation for arithmetic,

and psychological considerations were ruled out as irrelevant. Many aspects of the debate about this decision are relevant to developments in artificial intelligence.

Nonstandard logics and models of human knowledge processing are recent developments that lack the extensive, coherent literature of standard logic. However, it is important to consider, particularly in hardware design, that there may be systems of logic more appropriate to knowledge processing. An inappropriate system will not just be slower, but may run up against computational complexity barriers that make useful inference impossible.

Logical and psychological system theories based on an axiom of unrestricted predication are plausible and attractive alternatives to standard logic. In conclusion, it is suggested that the architecture of logical inference processors should take such systems into account and that further investigation of the logical, psychological, and systemic foundations of knowledge should be a major component of the development of expert systems and future generation computing systems.

ACKNOWLEDGMENTS

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

REFERENCES

1. **Anderson, A. R. and Belnap, N. D.,** *Entailment: The Logic of Relevance and Necessity,* Princeton University Press, Princeton, N. J., 1975.
2. **Bainbridge, L.,** Verbal reports as evidence of the process operator's knowledge, *Int. J. Man-Machine Studies,* 11, 411, 1979.
3. **Basden, A.,** On the application of expert systems, *Int. J. Man-Mach. Stud.,* 19, 461, 1983.
4. **Boose, J. H.,** Personal construct theory and the transfer of human expertise, in *Proc. AAAI-84,* American Association for Artificial Intelligence, Menlo Park, Calif., 1984, 27.
5. **Brown, G. S.,** *Laws of Form,* George Allen & Unwin Ltd., London, 1969.
6. **Checkland, P.,** *Systems Thinking, Systems Practice,* John Wiley & Sons, New York, 1981.
7. **Coombs, M. and Alty, J.,** Expert systems: an alternative paradigm, *Int. J. Man-Mach. Stud.,* 20, 21, 1984.
8. **Davis, R. and Lenat, D. B.,** *Knowledge-Based Systems in Artificial Intelligence,* McGraw-Hill, New York, 1982.
9. **De Bono, E.,** *Future Positive,* Maurice Temple Smith, London, 1979.
10. **Dizard, W. P.,** *The Coming Information Age: An Overview of Technology, Economics and Politics,* Longman, New York, 1982.
11. **Drucker, P. F.,** *The Age of Discontinuity,* Harper & Row, New York, 1978.
12. **Feyerabend, P.,** *Against Method,* NLB, London, 1975.
13. **Fitch, F. B.,** *Symbolic Logic, An Introduction,* Ronald Press, New York, 1952.
14. **Fraenkel, A. A., Bar-Hillel, Y., and Levy, A.,** *Foundations of Set Theory,* North-Holland, Amsterdam, 1973.
15. **Frege, G.,** *Die Grundlagen der Arithmetik,* Wilhelm Koebner, Breslau, 1884.
16. **Frege, G. and Husserl, R. von E.,** Philosophie der Arithmetik, *Z. Philosophie Philosophische Kritik,* 103, 313, 1894.
17. **Fuchi, K., Sato, S., and Miller, E.,** Japanese approaches to high-technology R&D, *Computer,* 17, 14, 1984.
18. **Gaines, B. R.,** Behaviour/structure transformations under uncertainty, *Int. J. Man-Mach. Stud.,* 8, 337, 1976.
19. **Gaines, B. R.,** Foundations of fuzzy reasoning, *Int. J. Man-Mach. Stud.,* 8, 623, 1976.
20. **Gaines, B. R.,** System identification, approximation and complexity, *Int. J. Gen. Syst.,* 3, 145, 1977.
21. **Gaines, B. R.,** Computers and society: the basis for a systemic model, in *Improving the Human Condition: Quality and Stability in Social Systems,* Ericson, R. F., Ed., Society for General Systems Research, Louisville, Ky., 1979, 523.

22. **Gaines, B. R.,** General systems research: quo vadis?, in *General Systems 1979,* Gaines, B. R., Ed., Society for General Systems Research, Louisville, Ky., 1980, 1.
23. **Gaines, B. R.,** Logical foundations for database systems, in *Fuzzy Reasoning and Its Applications,* Mamdani, E. H. and Gaines, B. R., Eds., Academic Press, London, 1981, 289.
24. **Gaines, B. R.,** Precise past — fuzzy future, *Int. J. Man-Mach. Stud.,* 19, 117, 1983.
25. **Gaines, B. R.,** Foundations of fifth generation computing, in *Proc. of CIPS SESSION 84,* Canadian Information Processing Society, Calgary, 1984, 500.
26. **Gaines, B. R. and Shaw, M. L. G.,** New directions in the analysis and interactive elicitation of personal construct systems, in *Recent Advances in Personal Construct Technology,* Shaw, M. L. G., Ed., Academic Press, London, 1981, 147.
27. **Gaines, B. R. and Shaw, M. L. G.,** A programme for the development of a systems methodology of knowledge and action, in *General Systems Research and Design: Precursors and Futures,* Reckmeyer, W. J., Ed., Society for General Systems Research, Louisville, Ky., 1981, 255.
28. **Gaines, B. R. and Shaw, M. L. G.,** *The Art of Computer Conversation: A New Medium for Communication,* Prentice-Hall, Englewood Cliffs, N.J., 1984.
29. **Gaines, B. R. and Shaw, M. L. G.,** A learning model for forecasting the future of information technology, *Future Computing Systems,* 1, 31, 1986.
30. **Gevarter, W. B.,** Expert systems: limited but powerful, *IEEE Spectrum,* 18, 39, 1983.
31. **Goguen, J. A.,** Concept representation in natural and artificial languages: axioms, extensions and applications for fuzzy sets, in *Fuzzy Reasoning and Its Applications,* Mamdani, E. H. and Gaines, B. R., Eds., Academic Press, London, 1981, 67.
32. **Goguen, J. A., Meseguer, J., and Plaisted, D.,** Programming with parametrized abstract objects in OBJ, in *Theory and Practice of Software Technology,* Ferrari, D., Bolgnani, M., and Goguen, J., Eds., North-Holland, Amsterdam, 1982, 163.
33. **Hayes-Roth, F., Waterman, D. A., and Lenat, D. B., Eds.,** *Building Expert Systems,* Addison-Wesley, Reading, Mass., 1983.
34. **Husserl, E.,** *Philosophie der Arithmetik. I. Psychologische und Logische Untersuchungen,* Pfeffer, Halle, 1891.
35. **Husserl, E.,** *Logische Untersuchungen. I. Prologomena zur Reinen Logik,* Max Niemeyer, Halle, 1900.
36. **Kant, I.,** *Critique of Pure Reason,* George Bell, London, 1781.
37. **Kelly, G. A.,** *The Psychology of Personal Constructs,* Norton, New York, 1955.
38. **Klir, G. J.,** Identification of generative structures in empirical data, *Int. J. Gen. Syst.,* 3, 89, 1976.
39. **Kneale, W. and Kneale, M.,** *The Development of Logic,* Clarendon Press, Oxford, 1962.
40. **Kuhn, T. S.,** *The Structure of Scientific Revolutions,* University of Chicago Press, Chicago, 1962.
41. **Lenat, D. B. and Brown, J. S.,** Why AM and EURISKO appear to Work, *Artif. Intell.,* 23, 269, 1984.
42. **Luhmann, N.,** *Trust and Power,* John Wiley & Sons, New York, 1979.
43. **Mamdani, E. H. and Assilian, S.,** An experiment in linguistic synthesis with a fuzzy logic controller, in *Fuzzy Reasoning and Its Applications,* Mamdani, E. H. and Gaines, B. R., Eds., Academic Press, London, 1981, 311.
44. **Mamdani, E. H., Procyk, T., and Baaklini, N.,** Application of fuzzy logic to controller design based on linguistic protocol, in *Discrete Systems and Fuzzy Reasoning,* Mamdani, E. H. and Gaines, B. R., Eds., Queen Mary College, London, 1976, 125.
45. **Mancuso, J. R. and Adams-Webber, J. R., Eds.,** *The Construing Person,* Praeger, New York, 1982.
46. **Martin, R. L., Ed.,** *The Paradox of the Liar,* Yale University Press, New Haven, Conn., 1970.
47. **Maydole, R. E.,** Paradoxes and many-valued set theory, *J. Philos. Logic,* 4, 269, 1975.
48. **Michalski, R. S. and Chilausky, R. L.,** Knowledge acquisition by encoding expert rules versus computer induction from examples — a case study involving soyabean pathology, *Int. J. Man-Mach. Stud.,* 12, 63, 1980.
49. **Michalski, R. S. and Carbonell, J. G., Eds.,** *Machine Learning: An Artificial Intelligence Approach,* Tioga, Palo Alto, Calif., 1983.
50. **Mill, J. S.,** *A System of Logic,* Longmans, London, 1843.
51. **Mohanty, J. N.,** *Husserl & Frege,* Indiana University Press, Bloomington, 1982.
52. **Moto-oka, T., Ed.,** *Fifth Generation Computer Systems,* North-Holland, Amsterdam, 1982.
53. **Pask, G.,** *Conversation, Cognition and Learning,* Elsevier, Amsterdam, 1975.
54. **Peccei, A.,** *One Hundred Pages for the Future: Reflections of the President of the Club of Rome,* Pergamon Press, Oxford, 1982.
55. **Peray, K. E. and Waddell, J. J.,** *The Rotary Cement Kiln,* Chemical Publishing, New York, 1972.
56. **Popper, K. R.,** Epistemology without a knowing subject, in *Logic, Methodology and Philosophy of Science III,* van Rootselaar, B. and Staal, J. F., Eds., North-Holland, Amsterdam, 333.
57. **Quinlan, J. R.,** Inferno: a cautious approach to uncertain inference, *Comput. J.,* 26, 255, 1983.
58. **Ralescu, D.,** A system theoretic view of social identification, in *Improving the Human Condition: Quality and Stability in Social Systems,* Society for General Systems Research, Louisville, Ky., 1979, 531.

59. **Reitman, W., Ed.,** *Artificial Intelligence Applications for Business,* Ablex, Norwood, N.J., 1984.
60. **Rescher, N.,** *The Coherence Theory of Truth,* Clarendon Press, Oxford, 1973.
61. **Ritchie, G. D. and Hanna, F. K.,** AM: a case study in AI methodology, *Artif. Intell.*, 23, 249, 1984.
62. **Runes, D. D.,** *Dictionary of Philosophy,* Peter Owen, London, 1960.
63. **Schank, R. C.,** *Conceptual Information Processing,* North-Holland, Amsterdam, 1975.
64. **Schock, R.,** *Logics Without Existence Presuppositions,* Almquist & Wiksell, Stockholm, 1968.
65. **Shackle, G. L. S.,** *Uncertainty in Economics,* Cambridge University Press, Cambridge, 1955.
66. **Shaw, M. L. G.,** *On Becoming a Personal Scientist,* Academic Press, London, 1980.
67. **Shaw, M. L. G., Ed.,** *Recent Advances in Personal Construct Technology,* Academic Press, London, 1981.
68. **Shaw, M. L. G.,** PLANET: some experience in creating an integrated system for repertory grid applications on a microcomputer, *Int. J. Man-Mach. Stud.*, 17, 345, 1982.
69. **Shaw, M. L. G.,** Epistemological hierarchies in knowledge and expert systems, in *Metamodels and Metasystems,* van Gigch, J. P., Ed., Abacus Press, Philadelphia, in press.
70. **Shaw, M. L. G.,** Interactive knowledge elicitation, in *Proc. CIPS SESSION 84,* Canadian Information Processing Society, Calgary, 1984, 202.
71. **Shaw, M. L. G. and Gaines, B. R.,** The personal scientist in the community of science, in *General Systems Research and Design: Precursors and Futures,* Reckmeyer, W. J., Ed., Society for General Systems Research, Louisville, Ky., 1981, 59.
72. **Shaw, M. L. G. and Gaines, B. R.,** Exploring personal semantic space, in *Empirical Semantics II,* Rieger, B., Ed., Studienverlag Brockmeyer, Bochum, West Germany, 1981, 712.
73. **Shaw, M. L. G. and Gaines, B. R.,** Eliciting the real problem, in *Adequate Modelling of Systems,* Wedde, H., Ed., Pergamon Press, Oxford, 1983, 100.
74. **Shaw, M. L. G. and Gaines, B. R.,** A computer aid to knowledge engineering, in *Proc. British Computer Soc. Conf. on Expert Systems,* Cambridge, 1983, 263.
75. **Shortliffe, E. H.,** *Computer-Based Medical Consultations: MYCIN,* Elsevier, New York, 1976.
76. **Sowa, J. F.,** *Conceptual Structures: Information Processing in Mind and Machine,* Addison-Wesley, Reading, Mass., 1984.
77. **Varela, F. J.,** A calculus for self-reference, *Int. J. Gen. Syst.*, 2, 5, 1975.
78. **Welbank, M.,** *A Review of Knowledge Acquisition Techniques for Expert Systems,* Martlesham Consultancy Services, BTRL, Ipswich, 1983.
79. **White, R. B.,** The consistency of the axiom of comprehension in the infinite-valued predicate logic of Lukasiewicz, *J. Philos. Logic,* 8, 509, 1979.
80. **Wittgenstein, L.,** *Tractatus Logico-Philosophicus,* Routledge & Kegan Paul, London, 1921.
81. **Zadeh, L. A.,** Fuzzy sets, *Inf. Control,* 8, 338, 1965.
82. **Zadeh, L. A.,** Outline of a new approach to the analysis of complex systems and decision processes, *IEEE Trans. Syst. Man Cybern.*, 2, 28, 1973.