EXPERT SYSTEMS IN AGRICULTURE

Brian R. Gaines

Knowledge Science Institute, University of Calgary
Calgary, Alberta, Canada T2N 1N4

Abstract: The development of expert systems technology and applications is reviewed with particular attention to applications in agriculture. The role of different information technologies in agriculture and their relations to expert systems are reviewed. Key achievements of expert systems applications are outlined. The main features of the different approach to system development involved in expert systems are described. Technical aspects of acquiring, encoding, and applying knowledge are briefly explained.

1. INTRODUCTION

Knowledge has always been significant to our civilization. However, in the past fifty years this significance has grown to a level where it is coming to dominate other socio-economic factors. Machlup in 1962 drew attention to the major and growing role of information services in the US economy and estimated that in 1958 knowledge production accounted for 29% of the US GNP (Machlup 1980). Porat's (1977) detailed analysis of the 1967 US national income-and-accounts estimated that the information sector accounted for 46% of GNP. Drucker (1978) took this up as a major issue in his book, *The Age of Discontinuity* and estimated that by 1975 the knowledge sector would account for 50% of US GNP. In the 1970s the predicted transition commenced to a post-industrial economy less dependent on manufacturing and human labour, and more dependent on knowledge and information technology.

Agriculture is a knowledge-based industry. It draws upon science and technology. It is enmeshed in the economics of the market place. It is subject to fluctuating political direction. In addition to all these general factors, effective agricultural management requires knowledge of principles peculiar to agriculture and their application in the light of detailed knowledge of local conditions. Agriculture involves the control of a diversity of flows, of materials, information, finance and knowledge. It involves the application of human expertise to complex management problems. In the past computing and communication technology has been useful in supporting this expertise, particularly through accounting and inventory management, access to market information, and instrumentation and control of agricultural processes. It is now attractive to consider the application of recent developments in knowledge-based computing systems, so-called "expert systems," to support human expertise in agricultural management.

Expert systems involve a change in our approach to technological development. The standard approach to problems of decision and control has been to identify in detail a physical model of the system to be controlled and then to compute optimal control strategies for this model. This has been very effective in a wide range of situations but breaks down if the system to be controlled cannot be modelled in detail. The expert systems approach to decision and control is a new paradigm that notes that people can operate effectively many systems for which we do not have detailed models. Instead of

modeling the system, it then makes sense to model the people and compute decision and control strategies that mimic their expertise.

This approach has been remarkably successful in some industrial control situations where automation was previously felt to be impossible. It has also been successfully applied to medical, geological and other high-level expertise. In general it seems applicable to areas where monolithic scientific modeling is not possible and where human skills are available to encode. Many aspects of agriculture may be characterized in this way. In particular, the new paradigm avoids the need to use technology to simplify the world until we can model it.

Expert systems are part of the "fifth-generation" of knowledge-based computing technology that has been developed in the 1980s (Moto-oka 1982, Gaines 1984). This new technology developed under the pressure of socio-economic needs of post-industrial for a "new intellectual technology" (Bell 1973). Figure 1 shows Bell's model of the three major economic sectors of our present society, and Figure 2 shows the primary characteristics of the "post-industrial" sector.

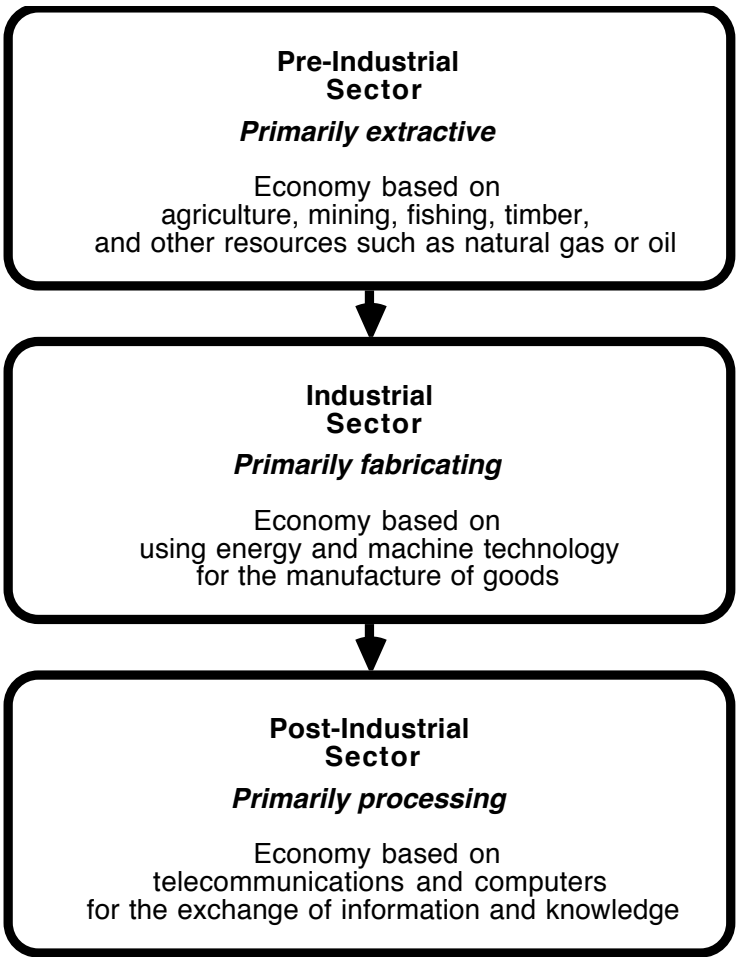

Fig.1 Major economic sectors of society

Economic Sector	Change from a goods producing to a service economy
Occupational Distribution	Pre-eminence of the professional and technical class
Axial Principle	Centrality of theoretical knowledge as the source of innovation and policy formulation for the society
Future Orientation	Control of technology and technological assessment
Decision-Making	Creation of a new "intellectual technology"

Fig.2 Features of the post-industrial sector

Figure 3 shows how the new technology impacts the three economic sectors (Gaines 1987b). The pre-industrial and industrial sectors have had an exchange relationship in which the pre-industrial provides materials for the industrial, and the industrial provides equipment for the pre-industrial. Now the industrial provides equipment for the post-industrial, and the post-industrial provides knowledge-based systems for the pre-industrial and the industrial, for example, expert systems for resource exploration and for energy-efficient farming, and computer-integrated manufacturing for low-cost automated indutries.

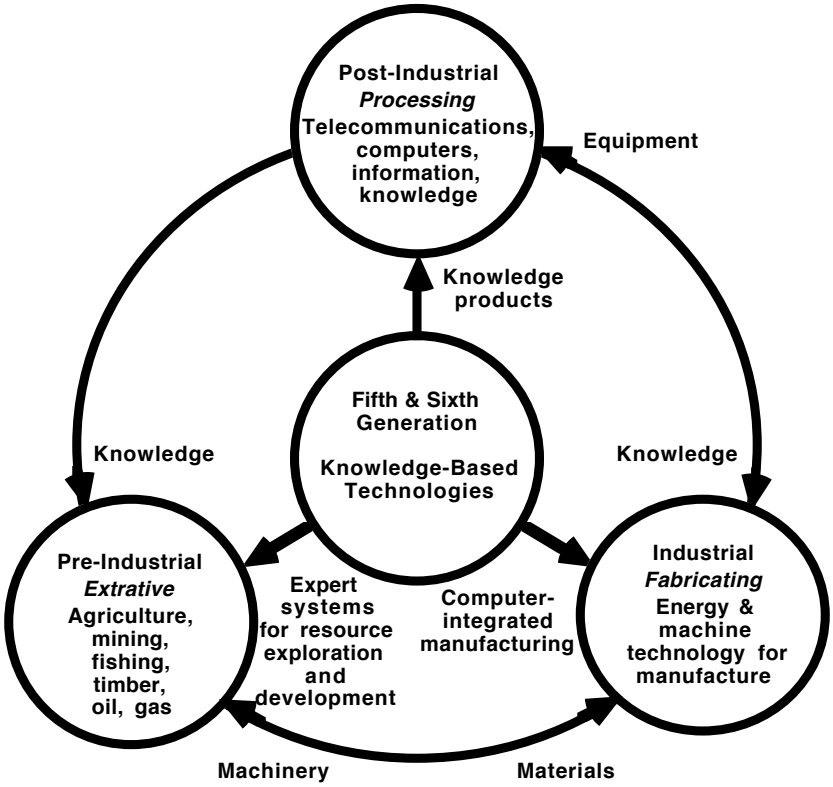

Fig.3 Relations between the major economic sectors and new technologies

The next section gives an overview of expert systems, their characteristics, structure and applications.

2. EXPERT SYSTEMS: AN OVERVIEW

The computer simulation of people in the roles of experts on some topic has become an important application of interactive computer systems. It has generated a new industry based on creating expert systems to make the practical working knowledge of a human expert in a specific subject area such as medicine or geology widely available to those without direct access to the original expert. Programs now exist that have made practical achievements in medical diagnosis, interpretation of mass spectroscopy results, analysis of geological survey data, and other problems where one would normally go to a human expert for advice (Michie 1979, Gevarter 1983, Reitman 1984, Waterman 1986). The initial success of expert system developments and the development of a number of reasonably domain-independent software support systems for the encoding and application of knowledge (Hayes-Roth, Waterman & Lenat 1983, Waterman 1986) has opened up the possibility of widespread usage.

2.1 EXPERT SYSTEMS ARCHITECTURES

Figure 4 shows the basic structure of an expert system:

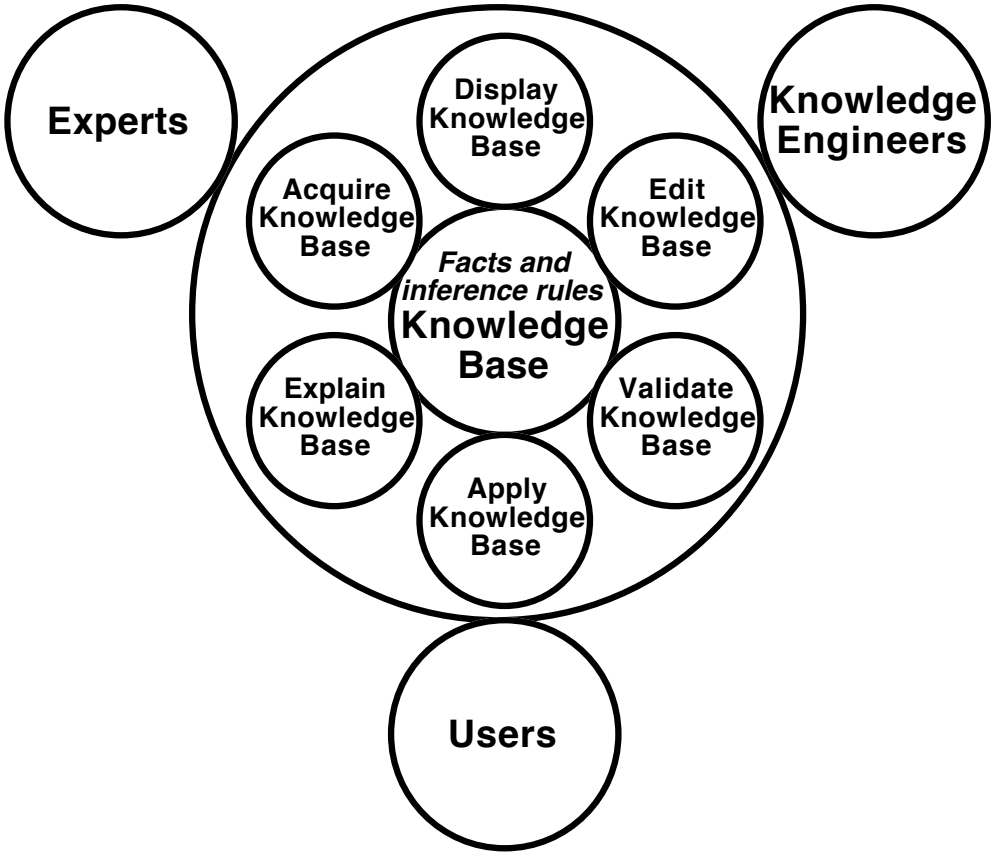

Fig.4 Structure of a knowledge-based system

• **Knowledge base**: the central ring shows the knowledge base of facts and inference rules. This extends the facilities of a conventional database by enabling the storage not only of facts but also of rules that enable further facts to be derived.

• **Operators**: the outer ring shows the three classes of operator of expert systems:

—**users** come to the system for advice on specific problems in a particular domain, supplying it with specific facts and their hypotheses about consequences or goals;

—**experts** come to the system to transfer their knowledge about the problem domain, supplying it with generally applicable facts and inference procedures;

—**knowledge engineers** act as intermediaries between the expert and the system, helping him to encode his knowledge and validate the operation of the expert system.

• **Expert system shell**: the middle ring shows the expert system shell that provides computational facilities for applying the knowledge base to user decision support:

—the **application system** at the bottom tests the user's hypotheses or seeks to satisfy his goals by deriving the consequences of facts about a particular situation reported by the user, using the general facts and inference rules supplied by the expert and edited by the knowledge engineer;

—the **explanation system** to the lower left answers queries about the way in which facts have been derived in terms of what information and inference rules have been used;

—the **acquisition system** to the upper left provides tools for interviewing the expert to obtain his vocabulary, and general facts and inference procedures in terms of it;

—the **display system** at the top provides tools for presenting the knowledge base in an understandable form—the relations between the facts and inference procedures;

—the **edit system** to the upper right provides tools for editing the knowledge base while maintaining its integrity in terms of the vocabulary, variables and operations used;

—the **validation system** to the lower right provides tools for checking the knowledge base against specific case histories with known consequences.

2.2 INITIAL STEPS IN EXPERT SYSTEMS TECHNOLOGY

One of the first expert system developments was the *fuzzy logic* control system developed in 1974 by Mamdani and Assilian. The system accepted human knowledge of control strategies expressed verbally and encoded it directly as computer programs which acted on the environment (Mamdani & Assilian 1981). This work was undertaken as part of a study of machine learning in process control and the system controlled was a small steam engine. The verbal 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

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, that is 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. The rules below are 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.

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.

In parallel with the controller development, other expert systems were being developed for completely different domains. MYCIN is a well-known medical diagnosis expert system which aids a clinician to act as a consultant on infectious diseases (Shortliffe 1976). It uses rules of the form:

RULE 50

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.

Such rules are obtained from specialists in microbial infections and their application to particular data is fairly simple data processing. The rules are validated through their application to many cases and revised when they fail to give the correct diagnosis. MYCIN is designed to interact with a clinician in order to make a diagnosis and suggest therapy for a particular patient with a suspected microbial infections. It first gathers data about the patient and then uses this to make inferences about the infections and their treatment. Figure 5 shows part of a dialog between MYCIN and a clinician responsible for testing and treating a patient.

--------PATIENT-1--------
(1) Patient's name: (first-last) **FRED SMITH**
(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? **GI (6)**

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]

Fig.5 A consultation with MYCIN

2.3 THE SPREAD OF EXPERT SYSTEMS TECHNOLOGY

Since the initial success with MYCIN a very wide range of expert systems have been developed. Figure 6 is a brief summary of some significant expert systems and their applications, and Figure 7 is a brief summary of some major areas of application of expert systems (Gevarter 1983). The numbers and domains have both increased greatly since these tables were prepared as research and applications of expert systems have developed world-wide. It is now impossible to make any accurate count since most universities have some activity and many industrial expert systems are regarded as highly proprietary.

DENDRAL	structure of chemical compounds from mass-spectral data and knowledge of mass spectrometry
META-DENDRAL	inferring rules of mass spectrometry
MYCIN	diagnosis of blood infections and recommendations for treatment
TEIRESIAS	knowledge acquisition and debugging
EMYCIN	empty (essential) MYCIN expert system shell
PUFF	EMYCIN with knowledge base for pulmonary function disorders
INTERNIST	advisors for general internal medicine (can deal with symptoms and diagnose 500 diseases)
CASNET	diagnosis of glaucoma
MOLGEN	scientist's assistant for planning experiments in molecular genetics
CRYSALIS	infers protein structure from electron density maps
PROSPECTOR	assisting evaluation of geological prospects
XCON (or R1)	system for configuring VAXs
AM	automated mathematician discovers interesting conjectures in mathematics

Fig.6 Some significant expert systems

2.4 THE IMPORTANCE OF META-KNOWLEDGE

MYCIN works well when the rules are 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 feels that something has gone wrong. To overcome these problems TEIRESIAS (Davis & Lenat 1982) was added as an auxiliary expert 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.

Analysis	electrical circuits, naval threat, mechanics, digital circuits, earthquake damage

Automatic programming	software generation
Concept formation	mathematics
Control	flexible manufacturing, lime kiln, smelting
Data analysis	geology, chemistry, protein crystallography, medicine, oil-well logs
Design	flexible manufacturing, computer system configuration, circuit, chemical synthesis
Diagnosis	medicine, computer faults, nuclear reactor accidents
Image analysis	weather forecasting, medical diagnosis
Instruction	electronics, mathematics, medical diagnosis, coaching, plant operation
Intelligent assistant	weapons assignment, medicine, radiology, computer sales, nuclear plant configuration
Knowledge acquisition	manufacturing, medical diagnosis, geology
Learning	plant diagnosis, chemistry, heuristics
Management	project management, automated factory
Monitoring	industrial plant, patient respiration
Planning	chemical synthesis, robotics, molecular genetics, medical diagnosis, scheduling

Fig.7 Areas of application of expert systems

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. The rules above are expressions of inferences from data about the world to conclusions about it, a basic level of knowledge processing. Davis discovered that similar rules could express inferences at a higher level about relations between these basic patterns of inference. A typical one of these *metarules* has the form:

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.

The effect of the metarules on the dialog is that the clinician and MYCIN-TEIRESIAS can discuss the rules and inferences made as well as the patient, organisms and infection. The dialog below shows a debugging sequence in which the inference made by MYCIN in Figure 5 is noted to be incorrect and TEIRESIAS aids the clinician in correcting it.

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.8 Knowledge-based knowledge correction with TEIRESIAS

In this dialog the clinician noticed an identity was missing in those deduced by MYCIN. He chose to allow TEIRESIAS to guide the debugging, and it used the metarule of starting with the identity which appeared incorrectly. The program then reconstructed the deductive process, asking the clinician to confirm its actions at each stage, or comment on anything incorrect. It confirmed the facts about the particular case in hand and checked that the rules were correct. Before very long the issue has become sharply focused for the clinician and the system.

The main guideline in debugging MYCIN 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.

2.5 INDUCTIVE ACQUISITION OF KNOWLEDGE

The interaction with TEIRESIAS demonstrates the dynamic nature of expert systems. 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 expert system 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 expert system 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.

The use of metarules allows expert systems 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* (Michalski & Carbonnel 1983).

Mamdani, Procyk and Baaklini (1976) 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:

> 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.

The use of a very high level of metarules for knowledge acquisition is illustrated by 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.

The rules below are 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 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).

> 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.

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 (1980) have reported an interesting comparison of two expert systems 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 structure from behavior. The expert system 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.

2.6. WHAT IS DIFFERENT ABOUT EXPERT SYSTEMS?

The previous section has shown how applications of expert systems to control and decision systems uses linguistic rules, approximate reasoning and metarules. It is clear that expert systems use some new computational techniques, but perhaps less apparent that the system design approach involved is itself fundamentally changed. However, major paradigm shifts are involved (Gaines & Shaw 1985). The classical approach in decision and control system design is the instrumentation, data collection, modeling and optimization sequence shown on the left of the diagram of Figure 9:

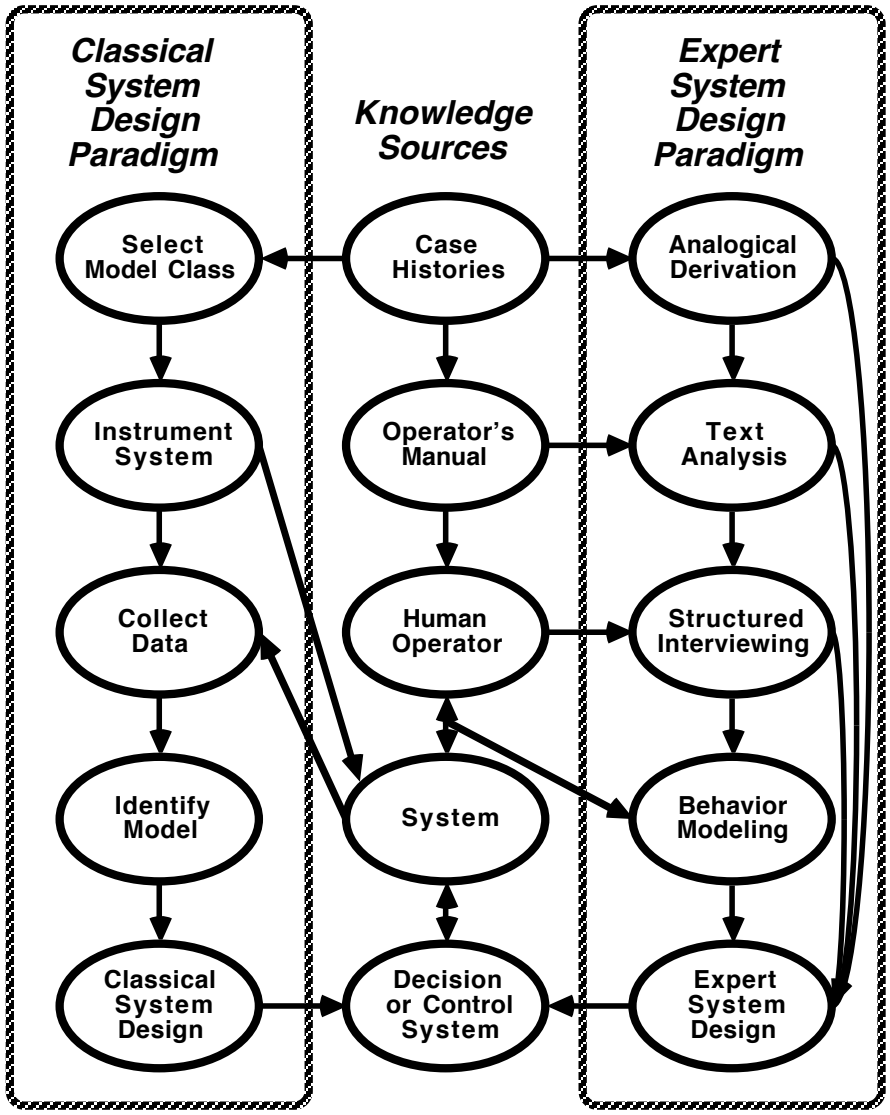

Fig.9 Classical and expert system design paradigms

- Knowledge of past case histories is used to select a class of system models
- The information required to discriminate within this class determines how the system should be instrumented for data acquisition
- Data is collected from the system through the instrumentation

• A model is identified from the model class which best fits the data
• This model is used to design a decision or control system for optimal performance

This approach to system design underlies the methodologies of the physical sciences and technologies based on them. It has the merit that it has been extremely successful in engineering much of the technological infrastructure of our current society including our manufacturing industries.

However, this approach is successful only to the extent that the systems under consideration are amenable to instrumentation and modeling. Its greatest successes have been where this amenability can be achieved normatively, that is in cases where the system to be controlled is itself a human artifact. For example, linear system theory has become a major tool in systems engineering not because most natural systems are linear—they are not. The implication is in the opposite direction: that linear systems are mathematically tractable and that we design artificial systems to be linear so that we may model them readily.

The application of "a linear model with quadratic performance criterion" to natural systems is often attempted but, in general, it does not work. We have done so not because the tool was appropriate, but because it was the only one we had. However, the use of a hammer to insert screws, although partially effective, tends to distort, destroy, and generally defeat the purpose of using a screw. Similarly, the use of an inappropriate system theory to model a system may give useful, but limited, results when we have no other, but it distorts reality, destroys information and generally defeats the purpose of modeling that system.

Much of our current technology succeeds to the extent that it is normative. In agriculture we reduce the complexity of a natural ecology to a comprehensible simplicity by the use of pesticides, herbicides and fertilizers. We reduce the system to be amenable to our modeling techniques. That simpler is not necessarily better and that re-engineering nature to impose uniformity destroys variety which is itself valuable have become major issues in recent years.

The expert systems paradigm may be seen to be particularly applicable when it is not possible to model the system but there is an alternative source of data available because human operators are able to performance the decision or control task. The right hand column shows the use of knowledge sources in the expert system design paradigm:

• Structured interviewing may be used to acquire knowledge directly from the operators
• Behavior modeling may be used to identify the operators' strategies even if they are unaware of them or give incorrect ones in interviews
• Text analysis may be used with instructional material such as the operators' manuals
• Reasoning by analogy may be used based on the case histories without the data collection of the classical system design paradigm

Note that the classical and expert system design paradigms, and the various acquisition techniques for expert systems, need not be regarded as competitive. They may all be used in the system design, and it may be that a collection of weak methods is necessary because there is no overall strong method that can be applied.

2.7 SUMMARY

Expert systems have developed in the past decade from experiments in artificial intelligence to a major commercial industry. They allow human expertise in knowledge representation and inference to be encoded for computers and used for problem-solving. They provide not only a static recording of expertise but also the capability to develop and extend that expertise and tailor it to local situations. Successful applications include not only high-level human cognitive skills such as medical diagnosis, but also important industrial control skills where people carry out tasks that could not be automated using previously available techniques. The most recent developments of expert system include knowledge acquisition not only through interaction with experts but also through inductive inference from experience. Expert systems technology involves fundamental paradigm shifts in system development and is an essential foundation to a new science of manufacturing.

The following section reviews approaches to the development of expert systems.

3. KNOWLEDGE ENGINEERING: EXPERT SYSTEMS DEVELOPMENT

What Feigenbaum (1980) terms *knowledge engineering*, the reduction of a large body of knowledge to a precise set of facts and rules, has already become a major bottleneck impeding the application of expert systems in new domains. We need to understand more about the nature of expertise in itself (Hawkins 1983) and to able to apply this knowledge to the elicitation of expertise in specific domains.

The problems of knowledge engineering have been stated clearly:

> "Knowledge acquisition is a bottleneck in the construction of expert systems. The knowledge engineer's job is to act as a go-between to help an expert build a system. Since the knowledge engineer has far less knowledge of the domain than the expert, however, communication problems impede the process of transferring expertise into a program. The vocabulary initially used by the expert to talk about the domain with a novice is often inadequate for problem-solving; thus the knowledge engineer and expert must work together to extend and refine it. One of the most difficult aspects of the knowledge engineer's task is helping the expert to structure the domain knowledge, to identify and formalize the domain concepts." (Hayes-Roth, Waterman & Lenat 1983)

This bottleneck is the major impediment to the application of expert systems, and to the realization of their full potential. The automation of knowledge acquisition and transfer has become a major research and technological goal. This section presents some techniques for the automation of the knowledge engineering task with examples of their application and indications of their further development and integration.

3.1 AUTOMATING KNOWLEDGE ENGINEERING

The problem of expertise elicitation from a skilled person is well-known in the literature of psychology (Nisbett & Wilson 1977, Broadbent, Fitzgerald and Broadbent 1986). Hawkins (1983) has analyzed the nature of expertise and emphasizes its severe limitations and dependence on critical assumptions which are often implicit. Dixon (1981) has surveyed studies showing that much human activity is not accessible to

awareness. Collins (1985) has studied knowledge transfer processes among scientists and suggests that some knowledge may not be accessible through the expert, not only because he cannot express it, but also because he may not be aware of its significance to his activity. Bainbridge (1979, 1986) has reviewed the difficulties of verbal debriefing and notes that there is no necessary correlation between verbal reports and mental behavior, and that many psychologists feel strongly that verbal data are useless.

Clinical psychologists see the problem as one of cognitive defences that impede internal communication, and have developed techniques of verbal interaction to identify underlying cognitive processes (Freud 1914, Kelly 1955, Rogers 1967). These can be used to by-pass cognitive defences, including those resulting from automization of skilled behavior. Hayes-Roth, Waterman & Lenat (1983) have given guidelines about knowledge transfer from experts to expert systems, and Welbank (1983) has surveyed the psychological problems of doing this.

In an industrial context the application of techniques derived from clinical psychology to problems of hard system engineering is not an obvious or probable step to take, and the early approach to knowledge engineering was to "manage" it by inventing the profession of knowledge engineer as defined in the quotation above. It has been possible that knowledge engineering might develop as a profession on a par with systems analysis and programming, and that an initial shortage of skilled knowledge engineers would cause problems to be overcome eventually be as the profession developed. However, this scenario now appears less and less likely. There is certainly a shortage of knowledge engineers and problems in developing applications, but doubts have been cast on the notion that human labor is the appropriate solution to the knowledge engineering problem:

- The decline in costs of both hardware and software support for expert systems has brought the technology into a mass-market situation far more rapidly than originally envisioned;
- This has lead to a growth in demand for expert systems that is proceeding far more rapidly than the growth in supply of trained and experienced knowledge engineers;
- The declining costs of expert system technology are also making the expense of human labour in tailoring the technology for particular applications appear to be the dominating constraint and an excessive cost;
- A move towards a labor-intensive activity such as knowledge engineering is contrary to all trends in industry;
- In particular it is contrary to the trend towards automatic programming techniques in the computing industry;
- The role of the knowledge engineer as an intermediary between the expert and the technology is being questioned not only on cost grounds but also in relation to its effectiveness—knowledge may be lost through the intermediary and the expert's lack of knowledge of the technology may be less of a detriment than the knowledge engineer's lack of domain knowledge.

These considerations of the previous section have heightened interest in the possibility of providing knowledge support systems to automate knowledge engineering as a process of direct interaction between domain experts and the computer.

3.2 KNOWLEDGE ACQUISITION FOR EXPERT SYSTEMS

A number of tools have been developed for the rapid prototyping of expert systems automating some or all of the methodologies on the right of Figure 9. The objective of systems such as PLANET (Shaw & Gaines 1983), ETS (Boose 1984), MOLE (Eshelman, Ehret, McDermott & Tan 1986), SALT (Marcus 1986), KITTEN (Shaw & Gaines 1986), KNACK (Klinker Bentolila, Genetet, Grimes & McDermott 1986), KRITON (Diederich, Ruhmann & May 1986), OPAL (Musen, Fagan, Combs & Shortliffe 1986), AQUINAS (Boose & Bradshaw 1987) and KSS0 (Gaines 1987a) is to expedite the process of knowledge acquisition and transfer to knowledge-based systems.

Figure 10 shows the architecture of one system that integrates a wide range of knowledge acquisition methodologies.

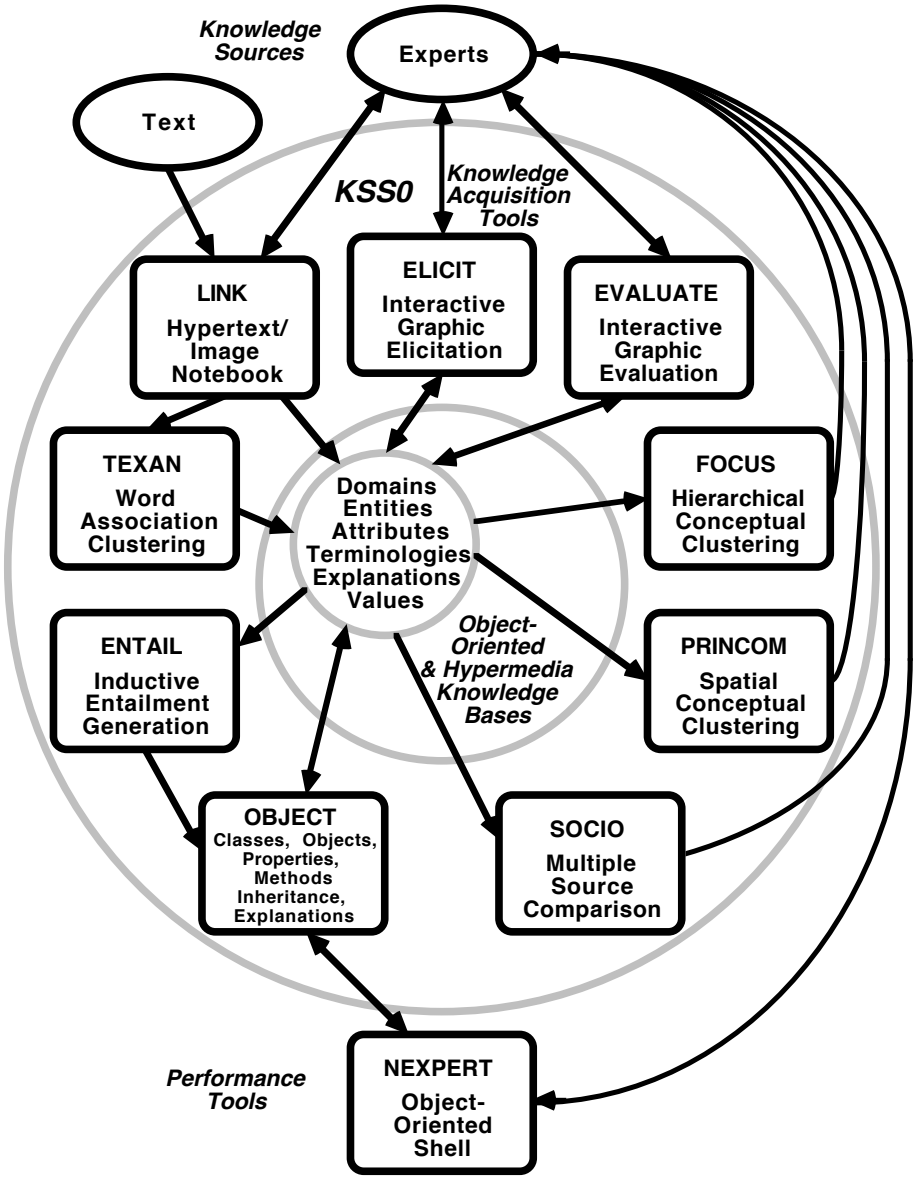

Fig.10 KSS0—an integrated knowledge acquisition system

KSS0 is written in Pascal and runs on the Apple Macintosh family of computers to provide a highly interactive and graphic knowledge acquisition environment. At the heart of KSS0 is an object-oriented knowledge base in which knowledge is formally represented as a multiple-inheritance digraph of classes, objects and properties. Such a structure generalizes the entity-attribute grids used in several early knowledge acquisition systems (Shaw & Gaines 1983, Boose 1984) and has proved both general and powerful in a variety of applications (Boose & Bradshaw 1986, Gaines 1987a). The main tools in KSS0 are:

- **LINK** provides a hypertext/image notebook allowing knowledge to be entered in free form. In particular, this allows more formal structures to be annotated with clarifications and explanations of vocabulary and lines of reasoning.
- **TEXAN** accepts text input directly and from the notebook, analyzes it by word frequency and association and produces a conceptual clustering of the most frequent highly-linked words. This is presented to the experts for click-and-drag graphic editing, and used to suggest sub-domains and associated entities and attributes.
- **ELICIT** accepts specifications of entities within a sub-domain and provides an interactive graphical elicitation environment within which the experts can distinguish entities to derive their attributes. The resultant class is continuously analyzed to provide feedback prompting the expert to enter further entities and attributes.
- **EVALUATE** accepts specifications of decision-making situations in terms of specifications of relevant classes, generally of problems and potential solutions. It provides an interactive graphical elicitation environment within evaluative attributes can be added, either by the expert, or by the knowledge engineer from protocols of their behavior.
- **FOCUS** hierarchically clusters entities and attributes within a sub-domain prompting the experts to add higher-level entities structuring the domain.
- **PRINCOM** spatially clusters entities and attributes within a sub-domain prompting the experts to add higher-level entities structuring the domain.
- **ENTAIL** induces logical entailments enabling the attributes of an entity, or the evaluations of a decision-making situation in a domain, to be derived from other attributes.
- **SOCIO** compares the structures for the same sub-domain generated by different experts, or the same expert at different times or from varying perspectives.
- **OBJECT** formats the specifications of sub-domains as classes, of entities as objects, of attributes as properties, of entailments as methods, and of links as explanations, and transfers them to a performance tool, currently Object Nexpert (Rappaport 1986). A return path is provided for the knowledge base to be reloaded to KSS0 after editing within the shell support environment.

Figure 11 shows KSS0 in action eliciting knowledge about mapping techniques from a geographer as part of the development and validation of a geologic expert system (Shaw & Woodward 1987):

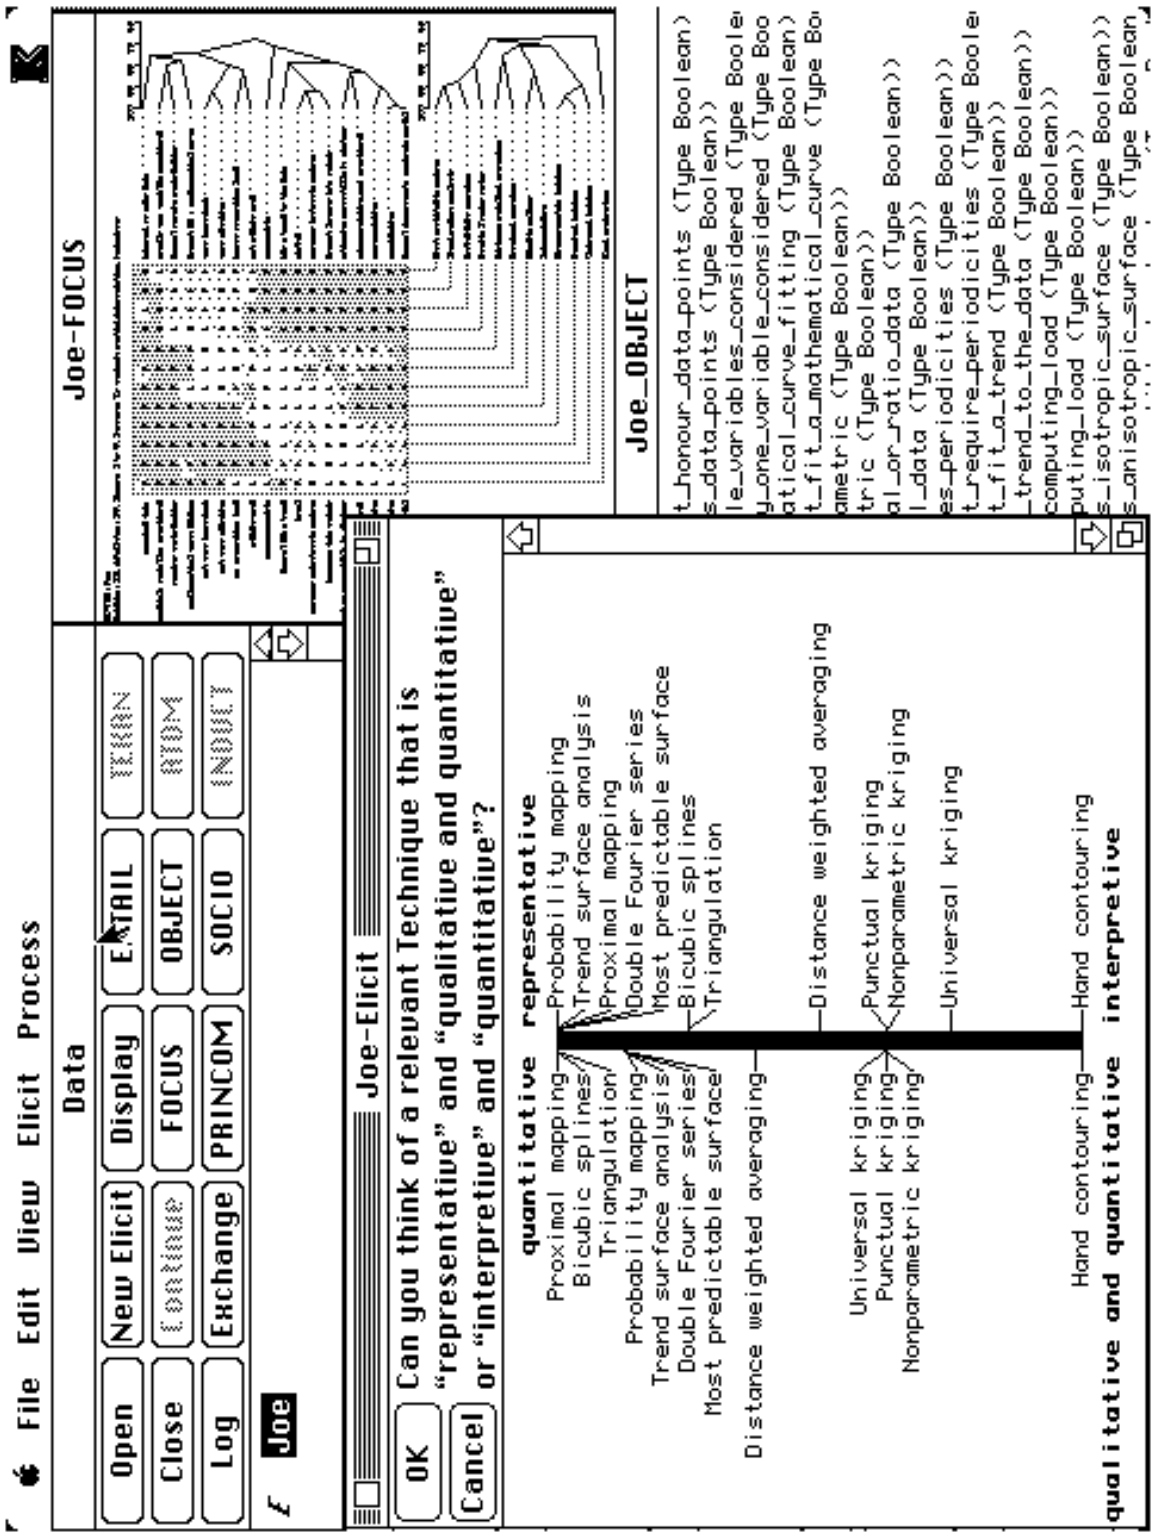

Fig.11 KSS0 in action

- The window at the top left is for session control giving access to the datasets and programs

- The window at the bottom left is for interactive entity and attribute elicitation. It is showing that two attributes are closely matched and prompting the expert to enter a new entity. All of the entities in the window are interactive—the expert can drag entities to new ratings on attributes, edit names, and so on.
- The window at the top right shows a hierarchical clustering of the entity-attribute data already elicited. This is used to elicit the object hierarchy from the expert.
- The window at the bottom right shows some of the properties, classes, objects and rules generated by the system for loading to Nexpert.

The role of the knowledge engineer in using systems such as KSS0 in one of managing the knowledge elicitation process. As shown in Figure 12, the knowledge engineers' roles include:

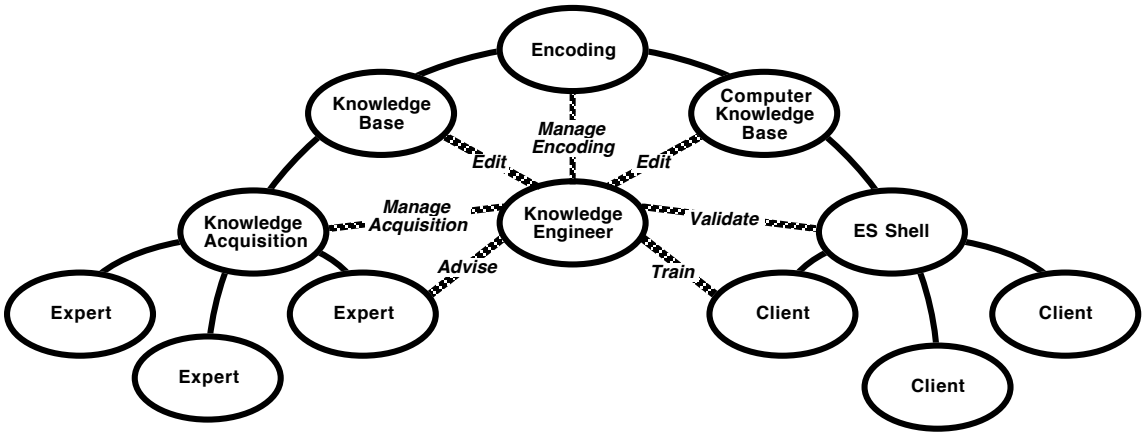

Fig.12 Extended model for knowledge acquisition

- Advising the expert on the process of interactive knowledge elicitation.
- Managing the interactive knowledge acquisition tools, setting them up appropriately.
- Editing the unencoded knowledge base in collaboration with the expert.
- Managing the knowledge encoding tools, setting them up appropriately.
- Editing the encoded knowledge base in collaboration with the expert.
- Validating the application of the knowledge base in collaboration with the expert.
- Training the clients in the effective use of the knowledge base in collaboration with the expert by developing operational and training procedures.

This use of interactive elicitation can be combined with manual elicitation, and with the use of the interactive tools by the knowledge engineer rather than, or in addition to, the expert. Thus, the role also includes:

- Directly eliciting knowledge from the expert.
- Using the interaction elicitation tools to enter knowledge into the knowledge base.

3.8 SUMMARY

This section has argued that while knowledge engineering is certainly the area to emphasize in the commercial exploitation of expert systems technology, it is misleading to assume that it necessarily involves the detailed involvement of knowledge engineers.

Knowledge engineering can, and should, be automated through knowledge support systems that are used directly be the subject domain experts. The knowledge engineer should operate only at the meta-level in helping the expert to select and apply the tools within these knowledge support systems. This has a number of significant advantages:

- efficient transfer of expertise without the overhead of intermediaries
- accurate transfer of expertise without distortion by intermediaries
- increased awareness by experts of the nature of their knowledge and expertise.

There are now a number of practical knowledge engineering tools demonstrating various possibilities for automating knowledge engineering. Some of these tools have already been used widely for rapid prototyping of industrial expert systems, and others are experimental illustrating the frontiers of current research. The integration of such tools into versatile knowledge support systems is a major research thrust at a number of organizations concerned with expert systems development. Any project involved in expert systems applications should monitor the development and availability of such tools, and regard them as a basic technology of at least equal importance to those of expert system shells.

5. CONCLUSIONS

It has been possible in this paper to give only a brief overview of expert systems, the underlying technology, its innovative features, applications, and approaches to development. Let me conclude with some points relevant to agriculture:

- The primary fifth generation objective is ease of use of computers—to make computers available for use by everyone, not just by specialists.
- Information technology is generated by the needs of society. 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 over-crowded, complex, inter-dependent world with limited resources.
- The industrial model constrains natural phenomena to fit models. The industrial model for agriculture has served us well. However, like the industrial model for many other aspects of our society, it is becoming out of date. The technology of the past made simplistic scientific models true by constraining natural phenomena to fit the models.
- We have simplified the world with technology: Industrial technology made the world simple, homogeneous and comprehensible. It 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 non- renewable resources.
- Information technology can model a complex world. It is capable of accepting the variety of the natural world and using it to extract sufficient resources for society's needs whilst husbanding long-term resources.
- Present day computer systems are primitive. We are at a very early stage in our understanding of information technology, in our capabilities to use it effectively, and in the power of the technology available to us. Present day computer systems are

primitive, difficult to apply and awkward to use. The fifth generation objective is to overcome these problems but we still have far to go.

- The key technology now is knowledge. It is knowledge technology, not computer technology, which limits our progress. The computer industry will make low-cost knowledge processing facilities widely available. It cannot tell us what knowledge we need to encode for them, and how we should process it, in order to solve fundamental problems in agriculture. Only those with a deep knowledge of agriculture can do this.
- Expert systems encode knowledge. Currently "expert systems" provide examples of computer-based knowledge processing in operation. We can think of them simplistically as means for encoding the expertise of a person and making it available to others. This is important in its own right. However, the technology provides much more than this.
- Expert systems serve to store knowledge. Expert systems are "knowledge banks" that can be used by individuals or groups to acquire, preserve and utilize knowledge. They create "knowledge liquidity" in the way that banks create financial liquidity.
- Expert systems provide a medium for communicating knowledge. Expert systems are "knowledge media" that can be used to communicate knowledge not as the static prescriptions of books but rather as the interactive partnership of an expert assistant.
- Expert systems avoid information overload. Information services provide the data necessary to solve problems but overload us with the task of searching out and structuring the information necessary to solve our particular problem. Expert systems technology provides an interface into information services that does the searching and structuring for us.
- We can use information technology to cope with a complex world. How will this affect agriculture? There are now many opportunities to use information technology to make the world simple and comprehensible without actually changing the world.
- We can reduce the use of chemical agents. For example, in the natural world we should be able to reduce the use of chemical agents to create simple and uniform environments for growing food and instead use feedback from local conditions coupled to expert knowledge from many domains to farm optimally in terms of the local environment. The trade-off between short-term yields and long-term conservation, and hence long-term yields, will be part of the expertise applied, and we might expect the "trade-off" to disappear—to be able to farm in such a way as to develop resources rather than destroy them, while gaining efficiency now.
- We can cope with complex legislation. For example, in the man-made world of complex legislation, government intervention and market forces, we should be able to use encoded expertise to carry out strategies that are productive and cost-effective while concentrating on farming rather than law, accounting, lobbying and other symbolic realities.
- The uncertainties of man and of nature will remain. Uncertainty, variety, change and development are the dynamics of our cultural evolution. We have spent many years reducing the world to the uniformity of the industrial paradigm. In future we shall

restore and develop its indefinite variety and learn to use it to create what we need as a by-product of this variety.

- Agriculture will be a key test of the new technology. We have created information technology as the tool necessary to make this shift in paradigms. Agriculture will be one of the major arenas in which we first see this shift take place.

6. ACKNOWLEDGEMENTS

Financial assistance for this work has been made available by the Natural Sciences and Engineering Research Council of Canada. I am grateful to Mildred Shaw of University of Calgary, John Boose and Jeff Bradshaw of Boeing AI Center for access to their research relating to knowledge-based systems and their development.

7. BIBLIOGRAPHY

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

Bainbridge, L. (1986). Asking questions and accessing knowledge. **Future Computing Systems 1**(2), 143-149.

Bell, D. (1973). **The Coming of Post-Industrial Society**. New York: Basic Books.

Boose, J.H. (1984). Personal construct theory and the transfer of human expertise. **Proceedings AAAI-84**, 27-33. California: American Association for Artificial Intelligence.

Boose, J.H. & Bradshaw, J.M. (1986). Expertise transfer and complex problems using AQUINAS as a knowledge acquisition workbench for expert systems. **Proceedings of AAAI Workshop on Knowledge Acquisition for Knowledge-Based Systems.** 5-0-14-16. Banff (November).

Broadbent, D.E., FitzGerald, P. & Broadbent, M.H.P. (1986). Implicit and explicit knowledge in the control of complex systems. **British Journal of Psychology 77**, 33-50.

Collins, H.M. (1985). **Changing Order: Replication and Induction in Scientific Practice**. London: SAGE.

Davis, R. (1983). TEIRESIAS: experiments in communicating with a knowledge-based system. Sime, M.E. & Coombs, M.J., Eds. **Designing for Human-Computer Communication**. pp. 87-137. London: Academic Press.

Davis, R. & Lenat, D.B. (1982). **Knowledge-Based Systems in Artificial Intelligence**. New York: McGraw-Hill.

Diederich, J., Ruhmann, I. & May, M. (1986) KRITON: A knowledge acquisition tool for expert systems. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the First AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.12-0-12-12. Banff (November).

Dixon, N. (1981). **Preconscious Processing**. Chichester: Wiley.

Drucker, P.F. (1978). **The Age of Discontinuity**. New York: Harper & Row.

Eshelman, L., Ehret, D., McDermott, J. & Tan, M. (1986) MOLE: A tenacious knowledge acquisition tool. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the First AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.13-0-13-12. Banff (November).

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

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

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

Gaines, B.R. (1986) An overview of knowledge acquisition and transfer. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the First AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.14-0-14-18. Banff (November).

Gaines, B.R. (1987a). Rapid prototyping for expert systems. Oliff, M., Ed. **Proceedings of International Conference on Expert Systems and the Leading Edge in Production Planning and Control**. pp.213-241. University of South Carolina.

Gaines, B.R. (1987b). Engineering knowledge: artificial intelligence, robotics and future society. **Canadian Society for Mechanical Engineering Proceedings of Canadian Engineering Centennial Convention**. pp.19.1-19.23.

Gaines, B.R. & Shaw, M.L.G. (1985). From fuzzy sets to expert systems. **Information Science 36** (1-2), 5-16 (July).

Gevarter, W.B. (1983). Expert systems: limited but powerful. **IEEE Spectrum**,**18**, 39-45.

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

Hayes-Roth, F., Waterman, D.A. & Lenat, D.B., Eds. (1983). **Building Expert Systems**. Reading, Massachusetts: Addison-Wesley.

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

Klinker, G., Bentolila, J., Genetet, S., Grimes, M. & McDermott, J. (1986) KNACK—report-driven knowledge acquisition. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the First AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.23-0-23-13. Banff (November).

Machlup, F. (1980). **Knowledge and Knowledge Production**. Princeton University Press.

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.

Mamdani, E.H., Procyk, T. & Baaklini, N. (1976). Application of fuzzy logic to controller design based on linguistic protocol. Mamdani, E.H. & Gaines, B.R., Eds. **Discrete Systems and Fuzzy Reasoning**. pp. 125-149. London: Queen Mary College.

Marcus, S. (1986) Taking backtracking with a grain of SALT. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the First AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.29-0-29-12. Banff (November).

Michalski, R.S. & Carbonell, J.G., Eds. (1983). **Machine Learning: An Artificial Intelligence Approach**. Palo Alto, California: Tioga.

Michalski, R.S. & Chilausky, R.L. (1980). Knowledge acquisition by encoding expert rules versus computer induction from examples—A case study involving soyabean pathology. **International Journal of Man-Machine Studies**, 12, 63-87.

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.

Musen, M.A., Fagan, L.M., Combs, D.M. & Shortliffe, E.H. (1986) Using a domain model to drive an interactive knowledge editing tool. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the First AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.33-0-33-11. Banff (November).

Nisbett, R.E. & Wilson, T.D.. (1977). Telling more than we can know: verbal reports on mental processes. **Psychological Review 84**, 231-259.

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

Porat, M.U. (1977). **The Information Economy**. Washington: US Department of Commerce.

Rappaport (1986). Multiple problem spaces in the knowledge design process. **Proceedings of AAAI Workshop on Knowledge Acquisition for Knowledge-Based Systems.** 37-0-37-15. Banff (November).

Reitman, W., Ed. (1984). **Artificial Intelligence Applications for Business**. Norwood, New Jersey: Ablex.

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

Shaw, M.L.G. & Gaines, B.R. (1983). A computer aid to knowledge engineering. **Proceedings of British Computer Society Conference on Expert Systems**, 263-271 (December). Cambridge.

Shaw, M.L.G. & Gaines, B.R. (1986) Techniques for knowledge acquisition and transfer. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the First AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.39-0-39-13. Banff (November).

Shaw, M.L.G. & Woodward, J.B. (1987) Validation of a knowledge support system. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Second AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.18-0-18-15. Banff (October).

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

Waterman, D.A. (1986). **A Guide to Expert Systems**. Reading, Masachusetts: Addison-Wesley.

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