Knowledge support sytems

Mildred L.G. Shaw
Department of Computer Science
University of Calgary
Calgary, Alberta
Canada T2N 1N4

ABSTRACT

This paper presents simulation as an aid to decision support systems and knowledge support systems. It includes the model of the communal scientist based on personal construct theory and the development of the model to that of the prosthetic brain. It presents the concept of knowledge support systems as interactive knowledge engineering tools, states the design criteria for such systems, and outlines the structure and key components of the KITTEN implementation. This includes tools for interactive repertory grid elicitation and entailment analysis that have been widely used for rapid prototyping of industrial expert systems. It also includes tools for text analysis, behavioral analysis and schema analysis, that offer complementary and alternative approaches to knowledge acquisition. The KITTEN implementation integrates these tools around a common database with utilities designed to give multiple perspectives on the knowledge base.

INTRODUCTION

Simulation is a particularly significant application of computing since, as well as the traditional uses it can provide decision support and aid in human knowledge processes by allowing theories and ideas to be tested in operational form. In recent years expert systems have also been used as an aid to human knowledge processes, since they can provide one person with interactive access to another person's skills encoded in a computer in some way. Most traditional simulation applications are concerned with physical variables that are far more concrete, and well understood, than the concepts of knowledge representation, inference and the reasoning processes of the human mind.

Also, the classical approach to designing a system solving some problem is to identify the problem in detail, model it mathematically, and then develop an algorithm for its solution either analytically or through simulation; whereas the expert systems approach is to assume that the problem is not identifiable in detail but that heuristic solutions to it exist and are implicitly available in the skills of experts who are known to be able to solve it. The expert system will then be built to model the behavior of the expert.

The success of this approach in solving previously intractable problems in industrial control, medicine, chemistry, and so on, has created widespread interest in the technology and acceptance of the new paradigm. It has many significant immediate applications and there are many similarities between the technologies of simulation and expert systems.

THE ARCHITECTURE OF EXPERT SYSTEMS

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 widely available to those without direct access to the original expert (1). 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. This section gives an overview of the structure of expert systems and knowledge engineering.

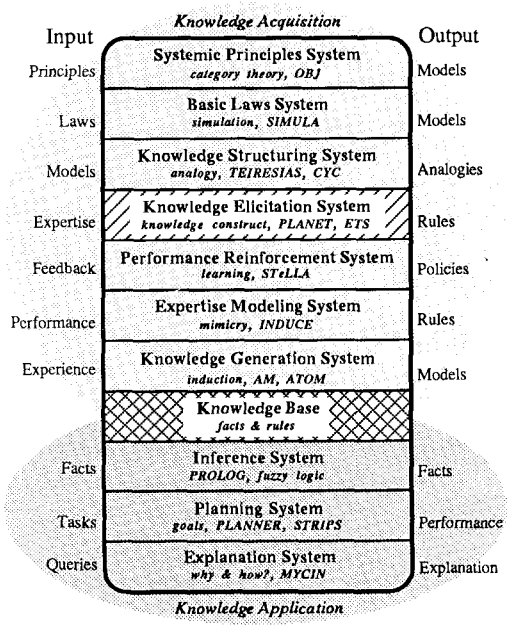

Fig.1 The knowledge engineering hierarchy

Figure 1 shows the basic architecture of an expert system: The central feature is a *knowledge base* of facts and inference rules. The lower oval shows the *expert system shell*, and the upper the *knowledge acquisition system.*

In the lower oval is the operational system for applying the knowledge base to decision support. The *inference system* derives the consequences of facts about a particular situation. Local facts or observations are used to derive further facts or conclusions in that situation. However, just generating facts is not sufficient, but a problem-solving capability is needed, so there is the *planning system* which determines how to use the inference system to satisfy the specified objectives. It is given the particular task and decides how to perform that task. The *explanation system* is an essential part of an expert system as it answers questions about the basis of the inferences made. This is essential to the user of the expert system who otherwise must take on trust the inferences and conclusions which are presented to him. An explanation system allows the user to trace the reasoning of the system by asking "how" and "why" questions about any result that the system produces. The knowledge base plus the expert system shell makes the total expert system.

The upper oval shows the knowledge transfer processes that may be used to establish the knowledge base. These are also the main areas of research activity. The *knowledge generation system* implements knowledge acquisition by the raw induction of models from experience without cultural support as exemplified by ATOM (2) and AM (3). That is, these systems are able to learn from the world and the events that they experience. The *expertise modeling system* implements knowledge acquisition by mimicking an expert's behavior as exemplified by INDUCE (4). For this process there must be available an expert who can do the task. The *performance*

reinforcement system implements knowledge acquisition from performance feedback as exemplified by STelLA (5), a learning machine which uses perceptual motor skills and is controlled by a reward and punishment regime. The *knowledge elicitation system* implements knowledge acquisition by interviewing experts as exemplified by PLANET (6) which was implemented on an Apple II and has been used in many other knowledge elicitation contexts as well as for expert systems, and ETS (7) which runs on an Xerox Lisp Machine. More recently KITTEN is being implemented on a network of Apollo workstations, and this is discussed in more detail later in the paper. The *knowledge structuring system* implements knowledge acquisition by analogy to transfer knowledge from one domain to another or from one situation to another, as exemplified by TEIRESIAS (3) and CYC (8). It is necessary for these systems to have some skeleton to work with, into which to put the knowledge. The *basic laws system* implements knowledge acquisition by building physical models as exemplified by simulation languages such as SIMULA (9). This is usually public domain knowledge and demonstrates the traditional model of simulation. The *systemic principles system* implements knowledge acquisition by derivation from abstract principles as exemplified by the category-theoretic language OBJ (10). This is the level of mathematical theories.

Most current knowledge acquisition is done by manual knowledge elicitation at the level highlighted. However, automatic tools for rapid prototyping are becoming increasingly important (11, 12, 13, 14) and are described by Gaines & Shaw (15).

THE PROSTHETIC BRAIN

Last year's paper (16) put forward a unifying framework for expert systems and simulation. It described the cognitive science foundations of a science of knowledge engineering in personal construct psychology. Personal construct psychology (17) provides foundations for cognitive science based on a conventionalist epistemological approach that has deep systemic foundations and is applicable to any anticipatory system (18, 19). This includes social systems and technological systems as well as the individual person.

Personal construct psychology can provide foundations for knowledge science based on a social model. Kelly's (17) focus was the individual as he was concerned primarily with clinical psychology. However, through his emphasis on the systemic principles underlying anticipation, the psychology he develops is applicable to organizations of individuals acting together as anticipatory systems. I have previously termed the personal construct psychology model of an individual a *personal scientist* (21) to emphasize that the basic human drive is modeling, hypothesis formation and testing, in order to antipate the future. However, this is also the drive of human society as a survival-driven system, and one may take the underlying system to be a *communal scientist* , interacting with the world, experiencing events, generating actions, and modeling the interaction so as to better anticipate future events.

Knowledge engineering may be defined as the acquisition, elicitation, structuring and encoding of knowledge for application in inferential, goal-directed, explanatory, decision and action support systems. Expert systems and simulations can be combined in decision support systems within the communal scientist model together with other support systems to help a person acquire knowledge. In Figure 2 the person is seen as supported by other people, media and various forms of information technology. For example, media may be representational media, such as radio or television; representational storage media, such as film or videotape; symbolic media, such as a book or research paper; or simulation media, such as computer graphics. One-way communication is shown by a single-headed arrow, and two-way communication by a double-headed arrow. Thus, whereas representational media which can transcend space, and symbolic media which can transcend time are one way communication systems: that is, the user cannot influence them; interactions with other people, simulation systems, expert systems and control systems involve two-way communication. These can be seen as complementary developments

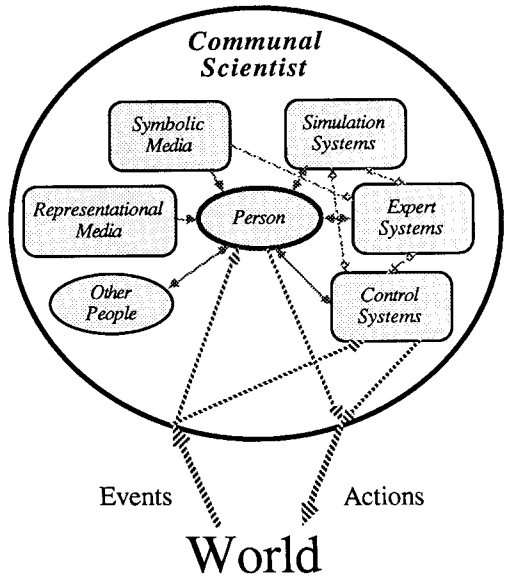

Fig.2 Knowledge support system for a personal scientist—the prosthetic brain

of computing to provide new media to support the knowledge, decision and control processes of people. From the individual's point of view each of these may be regarded as extensions of the processes within his brain, and the total structure may be termed a *prosthetic brain.* From a social point of view the entire sub-system shown is part of the knowledge process of society and may be understood in terms of *knowledge science.*

INTERACTIVE KNOWLEDGE ENGINEERING

It is clear to see from the previous diagram that expert systems and simulations can be combined in decision support systems and the varied interplay possible between people, simulated experts, simulations and the real world can be used to design new forms of system with enhanced capabilities. Most simulation is that of the physical world, but an expert system can be considered as the simulation of the mind of an expert.

Such considerations 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. In 1980 it was proposed that personal construct psychology (17, 21) could provide foundations for expert systems, particularly in systems that combined interactivity with database access and expert advice to provide decision support, and examples of algorithms and programs that extracted entailment rules from repertory grid data were given (22). This work led to industrial studies of the methodology applied to the development of expert systems: Boeing Computer Services (23, 7, 24) and Lockheed Software Technology Center (25) have reported success in applications; and validation has been reported in a statistics domain (26).

This paper gives a state-of-the-art report on the use of techniques based on personal construct psychology to automate knowledge engineering for expert systems. It is based on four areas of advance since previous papers:

- Improved techniques for the derivation of rules from repertory grid data which give: a natural knowledge representation for uncertain data combining fuzzy and probabilistic logics; and an information-theoretic measure of the significance of a derived rule (14);

- Widespread applications experience in prototyping expert systems using the methodology (7, 27);
- Improved interactive techniques for on-line knowledge engineering from groups of domain experts interacting through a computer network (28, 29);
- The development of a knowledge engineering workbench, KITTEN, that provides next generation knowledge support system facilities including textual analysis, induction of models from behavior, multi-level and multi-expert repertory grid elicitation, and hierarchical construct laddering, to automate knowledge engineering for a wide range of problem domains.

KNOWLEDGE SUPPORT SYSTEM DESIGN

Knowledge support systems are seen as having a broad scope, encompassing both aids to knowledge engineering and support of human knowledge processes—in the long term the division between knowledge engineering tools and expert system shells will break down, and integrated systems will be necessary. The general requirements for a knowledge support system are that:

1. The tools should be domain independent;
2. The tools should be directly applicable by experts without intermediaries;
3. The tools should be able to access a diversity of knowledge sources including text, interviews with experts, and observations of expert behavior;
4. The system should be able to encompass a diversity of perspectives including partial or contradictory input from different experts;
5. The system should be able to encompass a diversity of forms of knowledge and relationships between knowledge;
6. The system should be able to present knowledge from a diversity of sources with clarity as to its derivation, consequences and structural relations;
7. Users should be able to apply the knowledge in a variety of familiar domains and freely experiment with its implications;
8. The knowledge support system should make provision for validation studies;
9. As much of the operation as possible should be founded on well-developed and explicit theories of knowledge acquisition, elicitation and representation;
10. As the overall knowledge support system develops it should converge to an integrated system.

All of these requirements are subject to caveats—some domain dependence may be appropriate for efficiency in specific knowledge support systems—some human intervention may be helpful or necessary when an expert is using a knowledge support system—and so on. However, the broad design goals stated capture the key issues in knowledge support system design currently.

The PLANET system for repertory grid elicitation and analysis (21, 30, 31, 32) is a primitive knowledge support system satisfying requirements 1 and 2 for domain independence and direct use. Its foundations in personal construct psychology, which itself has strong systemic and cognitive science foundations (18, 16), are attractive in terms of requirement 9. Boose (7) in evaluating ETS has noted the limitations of basic repertory grid techniques in terms of requirement 5—that the methodology is better suited for analysis than for synthesis problems, for example, debugging, diagnosis, interpretation and classification rather than design and planning, and that it is difficult to apply to deep causal knowledge or strategic knowledge—and is attempting to overcome these by the use of grid hierarchies in NeoETS (33). The TEIRESIAS extension to MYCIN is an early form of knowledge support system providing debugging support for an expert system using basic analogical reasoning (3). The development of knowledge support systems has become a major area of activity recently, for example, MORE (34), SALT (35), SEAR (36), and MOLE (37).

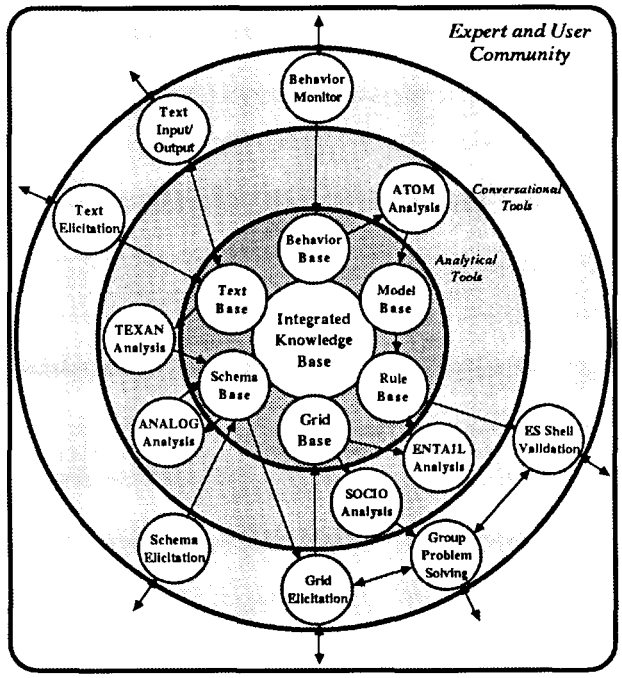

Fig.3 KITTEN
Knowledge Inititiation & Transfer Tools for Experts & Novices

THE KITTEN SYSTEM

Figure 3 shows the structure of KITTEN (Knowledge Initiation & Transfer Tools for Experts and Novices). It draws on many concepts and techniques for knowledge engineering to begin to encompass requirements 3 through 8, while attempting to satisfy 9 by relating them all through personal construct psychology, and 10 by building a workbench of tools around a common database. It consists of a: knowledge base; various analytical tools for building and transforming the knowledge base; and a number of conversational tools for interacting with the knowledge base. The KITTEN implementation is written in Pascal and currently runs on a network of Apollo workstations. Its structure is best understood by following sequences of activity that lead to the generation of a rule base and its loading into an expert system shell.

A typical sequence is text input followed by text analysis through TEXAN which clusters associated words leading to a schema from which the expert can select related elements and initial constructs with which to commence grid elicitation. TEXAN is a text analysis program designed to pump-prime the grid elicitation process when a manual or text book is available that the expert regards as having reasonable coverage of the domain. It uses techniques that were originally designed to map subject matter concepts against student concepts in computer-managed instruction systems (40). The text is fully indexed by all non-noise words grouped by their stems, and a coupling matrix of word associations is calculated using a simple distance-in-text measure. The high-frequency associations in the text are clustered and presented to the expert as a prototypical schema for the subject area which he can edit for spurious words and associations, and then use to suggest knowledge islands and associated elements and constructs.

Figure 4 shows the grid resulting from the elicitation of constructs (17) from an expert on staff appraisal, and Figure 5 the rank ordered entailments derived from this data. Such grids are analyzed by ENTAIL which induces the underlying knowledge structure as production rules that can be loaded directly into an expert system shell (14). The entailment analysis of a repertory grid treats each pole of a construct as a fuzzy predicate to which the elements have degrees of membership given by their ratings, and

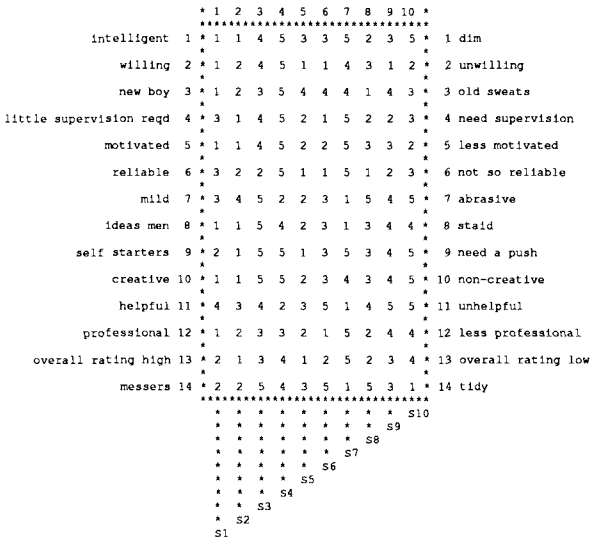

Fig.4 Repertory grid elicited on staff appraisal

induces the logical implications between these predicates. The original ENTAIL program produced all entailments consistent with the grid and allowed the expert to prune any that seemed spurious before using them as inference rules in an expert system. ENTAIL II rank orders entailments in terms of the uncertainty reduction they induce in the distribution of the data, and hence tends to reject spurious entailments (14).

Entail	Truth	Inf.	(Cutoff 0.18) Implication
L 1 → L 9	0.80	0.29	intelligent→self starters
L 9 → L13	1.00	0.29	self starters→overall rating high
R 9 → R 1	0.80	0.28	need a push→dim
L10 → L 8	1.00	0.28	creative→ideas men
L 1 → L10	0.80	0.26	intelligent→creative
R 8 → R10	1.00	0.26	staid→non-creative
L10 → L 9	0.80	0.26	creative→self starters
R13 → R 6	0.80	0.26	overall rating low→not so reliable
L 9 → L10	0.80	0.24	self starters→creative
R10 → R 1	0.80	0.24	non-creative→dim
L10 → L 1	0.80	0.23	creative→intelligent
R13 → R 9	1.00	0.23	overall rating low→need a push
R 4 → R13	0.80	0.22	need supervision→overall rating low
R 5 → R 4	0.80	0.22	less motivated→need supervision
R 5 → R13	0.80	0.22	less motivated→overall rating low
R 9 → R10	0.80	0.22	need a push→non-creative
L 1 → L 3	0.80	0.21	intelligent→new boy
L 6 → L13	0.80	0.21	reliable→overall rating high
R10 → R 9	0.80	0.20	non-creative→need a push
R 1 → R 6	0.60	0.19	dim→not so reliable
R 1 → R10	0.80	0.19	dim→non-creative
R 9 → R 4	0.60	0.19	need a push→need supervision
R 9 → R12	0.60	0.19	need a push→less professional
R 9 → R13	0.60	0.19	need a push→overall rating low
R12 → R13	0.80	0.19	less professional→overall rating low
R13 → R 4	0.80	0.19	overall rating low→need supervision
R13 → R12	0.80	0.19	overall rating low→less professional
L 4 → L 5	0.80	0.18	little supervision reqd→motivated
L 4 → L 9	0.60	0.18	little supervision reqd→self starters
R 6 → R 4	0.80	0.18	not so reliable→need supervision
R 6 → R13	0.80	0.18	not so reliable→overall rating low
L12 → L 9	0.60	0.18	professional→self starters
L13 → L 4	0.80	0.18	overall rating high→little supervision reqd
L13 → L 5	0.80	0.18	overall rating high→motivated
L13 → L 9	0.60	0.18	overall rating high→self starters

Fig.5 ENTAIL analysis of repertory grid on staff appraisal

A repertory grid, such as that of Figure 4, may be viewed as a component of a database in entity-antribute form with elements as entities, constructs as attributes and ratings of elements on poles of constructs as values. The entailments are shown with two values in the range from 0 to 1: first, the truth value of the hypothesis; and second, the information content (uncertainty reduction generated) of asserting the hypothesis. For example, L1 -> L9 has a truth value of 0.80, and an information content of 0.29. The information content measures the significance of the hypothesis and is used to ensure that trivial entailments consistent with the data are pruned. The ENTAIL analysis of Figure 5 shows that L1, L4, L6, L9, L10 and L12 imply L13, that intelligent, creative, reliable and professional self-starters requiring little supervision receive a high overall rating, whereas R4, R5, R6, R9 and R12 imply R13, that being less motivated, not so reliable, less professional, needing supervision and needing a push leads to a low overall rating.

An alternative route is to monitor the expert's behavior through a verbal protocol giving information used and decisions resulting and analyze this through ATOM. Michalski and Chilausky (4) have demonstrated that inductive modeling of an expert's behavior may produce effective rules when those elicited by interview techniques are clearly inadequate. ATOM is an algorithm for inducing the structure of a system from its behavior using a search over a model space ordered by complexity and goodness of fit. As in ENTAIL, models are evaluated in terms of the uncertainty reduction induced by the model in the distribution of the modeled behavior (42, 2, 43).

These two routes can be combined. KITTEN attempts to make each stage as explicit as possible, and, in particular, to make the rule base accessible as natural textual statements rather than technical production rules. The expert system shell being used in KITTEN currently is Nexpert (38) which gives a variety of textual and graphical presentations of the rule base enabling the expert to see the impact of different fragments of knowledge. Figure 6 shows Nexpert in operation loaded with the entailments of Figure 5.

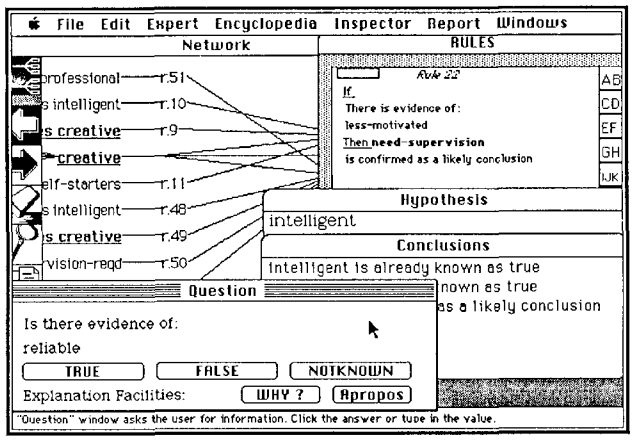

Fig.6 Inference rules, derived by ENTAIL from repertory grid data on staff appraisal, in use in Nexpert expert system shell

Interaction with Nexpert enables the expert to see the derived rules in action. He can determine their consequences with test data, analyze new hypothetical cases, and see the inter-relations between rules presented graphically. The logging and explanation facilities of Nexpert enable him to track down spurious inferences that may arise with the rules derived by ENTAIL, or proper inferences that are missing. He can then edit the rules and test the revized system using Nexpert's facilities.

The groupings resulting from TEXAN, when combined with the construct classifications of repertory grids as in Figure 4, may be viewed as schema structuring a knowledge domain. ANALOG is a method that maps schema to schema based on their structure without regard to content. It is based on a theory of analogy that explicates analogies as pullbacks of faithful functors between categories (41)

and generates maximal sub-graph isomorphisms between two classificatory data structures. It may be regarded as a generalization of the copy-edit process being used in the encoding of commonsense knowledge in CYC (8). ANALOG produces meaningful results on artificial examples and grids in related domains. It will also find meaningless analogies between unrelated domains which cannot be rejected by information-theoretic statistical procedures such as those used in ENTAIL and ATOM. It seems likely that effective application of ANALOG depends on the expert pump-priming the matching with known or hypothesized relations and the program extending these rather than attempting to generate them completely. At present, this is still in the experimental stage of development.

The group problem-solving component of KITTEN is particularly important because it goes beyond the stereotype of an "expert" and "users", and allows the system to be used to support an interactive community in their acquisition and transfer of knowledge and mutual understanding. The SOCIO analysis allows members of a community to explore their agreement and understanding with other members, and to make overt the knowledge network involved (21, 39). Much expertise only resides within the social context of cooperating individuals and requires elicitation across the group. The SOCIO analysis program supports group elicitation techniques in which the construct systems of a number of users are compared. Grids are elicited separately but then exchanged in two ways: a user can place elements on a colleague's constructs from his own point of view, and the analysis system then allows him to explore their *agreement*; or he can attempt to place them from his colleague's point of view and hence explore his *understanding*.

CONCLUSIONS

This paper has discussed the relationship of simulation to expert systems in two ways. The first was that together with people and different forms of media they form a support system for the individual to make decisions in the world. The second was that an expert system can be seen as the simulation of the mind of an expert. It also presented the concept of knowledge support systems as interactive knowledge engineering tools, stated the design criteria for such systems, and outlined the structure and key components of the KITTEN implementation. The development of KITTEN has contributed the first steps towards an integrated knowledge support system by building the tools around a common database, providing access to the same data in each of its intermediate forms, and providing conversion utilities between different data forms.

ACKNOWLEDGEMENTS

Financial assistance for this work has been made available by the National Sciences and Engineering Research Council of Canada and the work has been done jointly with Brian Gaines. We are grateful to John Boose and Jeff Bradshaw of Boeing AI Center, and Ernie Chang of Alberta Research Council for stimulating discussions relating to knowledge support systems.

REFERENCES

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

(2). Gaines, B.R. (1977). System identification, approximation and complexity. *International Journal of General Systems*, 3, 145-174.

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

(4). 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.

(5). Gaines, B.R. & Andreae, J.H. (1966). A learning machine in the context of the general control problem. *Proceedings of the 3rd Congress of the International Federation for Automatic Control*. London: Butterworths.

(6). 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.

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

(8). Lenat, D., Prakash, M. & Shepherd, M. (1986). CYC: Using common sense knowledge to overcome brittleness and knowledge acquisition bottlenecks. *AI Magazine* 6(4), 65-85.

(9). Nygaard, K. & Dahl, O-J. (1980). The development of the SIMULA languages. Wexelblat, R.L. (Ed.) (1981). *History of Programming Languages*. pp. 439-480. New York: Academic Press.

(10). Goguen, J.A. & Meseguer, J. (1983). *Programming with parametrized abstract objects in OBJ. Theory and Practice of Programming Technology*. Amsterdam: North-Holland.

(11). Gaines, B.R. & Shaw, M.L.G. (1984). Logical foundations of expert systems. *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, 238-247 (October). Halifax, Nova Scotia.

(12). Gaines, B.R. & Shaw, M.L.G. (1985). Foundations of expert systems. Kacprzyk, J. & Yager, R.R., Eds. *Management Decision Support Systems Using Fuzzy Sets and Possibility Theory*. pp. 3-24. Rheinland: ISR Interdisciplinary Systems Reserch, Verlag TUV.

(13). Gaines, B.R. & Shaw, M.L.G. (1985). Systemic foundations for reasoning in expert systems. Gupta, M.M., Kandel, A., Bandler, W. & Kiszka, J.B., Eds. *Approximate Reasoning in Expert Systems*. 271-281. Amsterdam: North-Holland.

(14). Gaines, B.R. & Shaw, M.L.G. (1986). Induction of inference rules for expert systems. *Fuzzy Sets and Systems* 18(3), 315-328 (April).

(15). Gaines, B.R. & Shaw, M.L.G. (1986). Knowledge engineering techniques. *Proceedings of AUTOFACT'86*. Detroit (November).

(16). Shaw, M.L.G. & Gaines, B.R. (1986). A framework for knowledge-based systems unifying expert systems and simulation. Luker, P.A. & Adelsberger, H.H., Eds. *Intelligent Simulation Environments*, pp. 38-43 (January). La Jolla, California: Society for Computer Simulation.

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

(18). Gaines, B.R. & Shaw, M.L.G. (1981). A programme for the development of a systems methodology of knowledge and action. Reckmeyer, W.J., Ed. *General Systems Research and Design: Precursors and Futures*. pp. 255-264. Society for General Systems Research (January).

(19). Mancuso, J.R. & Adams-Webber, J.R., Eds. (1982). *The Construing Person*. New York: Praeger.

(20). Shaw, M.L.G. (1985). Communities of knowledge. Epting, F. & Landfield, A.W., Eds. *Anticipating Personal Construct Psychology*. pp. 25-35. Lincoln: University of Nebraska Press.

(21). Shaw, M.L.G. (1980). *On Becoming a Personal Scientist*. London: Academic Press.

(22). Gaines, B.R. & Shaw, M.L.G. (1980). New directions in the analysis and interactive elicitation of personal construct systems. *International Journal of Man-Machine Studies* **13**(1) 81-116 (July).

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

(24). Boose, J.H. (1986). Rapid acquisition and combination of knowledge from multiple experts in the same domain. *Future Computing Systems*, In press.

(25). Wahl, D. (1986). An application of declarative modeling to aircraft fault isolation and diagnosis. Luker, P.A. & Adelsberger, H.H., Eds. *Intelligent Simulation Environments*, pp. 25-28 (January). La Jolla, California: Society for Computer Simulation.

(26). Gammack, J.G. & Young, R.M. (1985). Psychological techniques for eliciting expert knowledge. Bramer, M., Ed. *Research and Development in Expert Systems*, pp. 105-116. Cambridge University Press.

(27). Gaines, B.R. & Shaw, M.L.G. (1986b). Knowledge Engineering for an FMS Advisory System, *Proceedings of AMS 86*, to appear (June).
(28). Shaw, M.L.G. (1986). PCS: a knowledge-based interactive system for group problem solving. *Proceedings of 1986 International Conference on Systems, Man and Cybernetics*, to appear (October).
(29). Shaw, M.L.G. & Chang, E. (1986). A participant construct system. *Proceedings of North American Personal Construct Network Second Biennial Conference*. pp. 131-140. University of Calgary: Department of Computer Science (June).
(30). Shaw, M.L.G. (1982). PLANET: some experience in creating an integrated system for repertory grid applications on a microcomputer. *International Journal of Man-Machine Studies*, 17(3), 345-360.
(31). Shaw, M.L.G. & Gaines, B.R. (1986). Interactive elicitation of knowledge from experts. *Future Computing Systems*, 1(2), to appear.
(32). Shaw, M.L.G. & Gaines, B.R. (1986). An interactive knowledge elicitation technique using personal construct technology. Kidd, A., Ed. *Knowledge Elicitation for Expert Systems: A Practical Handbook*. To appear. Plenum Press.
(33). Bradshaw, J.M. & Boose, J.H. (1986). NeoETS. *Proceedings of North American Personal Construct Network Second Biennial Conference*. pp. 27-41. University of Calgary: Department of Computer Science (June).
(34). Kahn, G., Nowlan, S. & McDermott, J. (1985). MORE: an intelligent knowledge acquisition tool. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence* . pp. 581-584. California: Morgan Kaufmann.
(35). Marcus, S., McDermott, J. & Wang, T. (1985). Knowledge acquisition for constructive systems. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence* . pp. 637-639. California: Morgan Kaufmann.
(36). van de Brug, A, Bachant, J. & McDermott, J. (1985). Doing R1 with style. *Proceedings of the Second Conference on Artificial Intelligence Applications*. IEEE 85CH2215-2, pp. 244-249. Washington: IEEE Computer Society Press.
(37). Eshelman, L. & McDermott, J. (1986). MOLE: a knowledge acquisition tool that uses its head. *Technical Report*. Carnegie-Mellon University: Department of Computer Science.
(38). Roy, J. (1986). Expert systems in Nexpert. *MacTutor*, 2(2), 48-51 (February).
(39). Shaw, M.L.G. (1981). Conversational heuristics for eliciting shared understanding. Shaw, M.L.G., Ed. *Recent Advances in Personal Construct Technology*. pp. 31-44. London: Academic Press.
(40). Smith, R.A. (1976). Computer-based structural analysis in the development and administration of educational materials. *International Journal of Man-Machine Studies*, 8(4), 439-463 (July).
(41). Gaines, B.R. & Shaw, M.L.G. (1982). Analysing analogy. Trappl, R., Ricciardi, L. & Pask, G., Eds. *Progress in Cybernetics and Systems Research*. Vol. IX, pp. 379-386. Washington: Hemisphere.
(42). Gaines, B.R. (1976). Behaviour/structure transformations under uncertainty. *International Journal of Man-Machine Studies* , 8(3), 337-365 (May).
(43). Gaines, B.R. (1979). Sequential fuzzy system identification. *International Journal of Fuzzy Sets and Systems*, 2(1), 15-24 (January).