KNOWLEDGE SUPPORT SYSTEMS: INTEGRATED KNOWLEDGE ACQUISITION AND EXPERT SYSTEMS

Mildred L. G. Shaw & Brian R. Gaines
Knowledge Science Institute
Department of Computer Science
University of Calgary
Calgary, Alberta
Canada T2N 1N4

ABSTRACT

Growing recognition of the significance of knowledge-based computing systems has focused attention on processes of knowledge acquisition and transfer. Commercial applications of expert systems are being impeded by the knowledge engineering bottleneck and have led to the development of rapid prototyping tools. These have proved practically useful but the range of application of existing tools is limited, and it is not clear what limitations are inherent in the prototyping tools, what arises from the shells and their knowledge representation and inferencing procedures, and what arises from our lack of understanding of the underlying processes of knowledge acquisition and transfer.

This paper introduces the concept of knowledge support systems — integrated interactive knowledge acquisition systems and expert systems shells. A prototype knowledge support system that has been implemented is described with examples of some of the knowledge acquisition and application tools provided. It includes knowledge acquisition and transfer through interactive repertory grid elicitation; exchange techniques for understanding and agreement between experts; and different forms of analysis that have been widely used for comparing data between experts and for rapid prototyping of expert systems.

INTRODUCTION

Knowledge acquisition and transfer has long been thought of as the bottleneck in the process of building expert systems [8]. It seemed 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 which would be overcome eventually 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 technology is now in a mass-market situation where many organizations see the need for expert systems. This has led to a growth in demand that is far more rapid than the growth in supply of trained and experienced knowledge engineers. In addition, 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. Full exploitation of the potential of expert systems depends on the development of rapid prototyping systems directly usable by experts with the knowledge engineer acting only as manager, not intermediary.

KNOWLEDGE SUPPORT SYSTEMS

Knowledge support systems can be seen as integrated interactive knowledge acquisition systems and expert systems shells. They are seen as having a broad scope, encompassing both aids to knowledge acquisition and transfer, knowledge representation, and support of human knowledge processes [15]. It is desirable that the tools in a knowledge support system are domain independent and directly applicable by experts without intermediaries. They should be able to access a diversity of knowledge sources including text, interviews with experts, and observations of expert behavior and be able to encompass a diversity of perspectives including partial or contradictory input from different experts. Moreover, the knowledge support system should be able to encompass a diversity of forms of knowledge and relationships between knowledge and be able to present knowledge from a diversity of sources with clarity as to its derivation, consequences and structural relations. Users of the knowledge support system should be able to apply the knowledge in a variety of familiar domains and freely experiment with its implications. It should make provision for validation studies and as much of the operation of the knowledge support system as possible should be founded on well-developed and explicit theories of knowledge acquisition, elicitation and representation. As the overall knowledge support system develops it should converge to an integrated system.

Clearly, these are general requirements and may not apply in particular cases. For example, 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 systems. However, in general these requirements capture the key issues in knowledge support system design.

A KNOWLEDGE ACQUISITION TOOL

Figure 1 shows the architecture of one system that integrates a wide range of knowledge acquisition methodologies. The system is written in Pascal and runs on the Apple Macintosh family of computers to provide a highly interactive and graphic knowledge acquisition environment [3, 7]. At the heart of the system is an object-oriented knowledge base in which knowledge is formally represented as a multiple-inheritance structure of classes, objects, properties, values, and relations. Such a structure generalizes the entity-attribute datasets used in several early knowledge acquisition systems and has proved both general and powerful in a variety of applications.

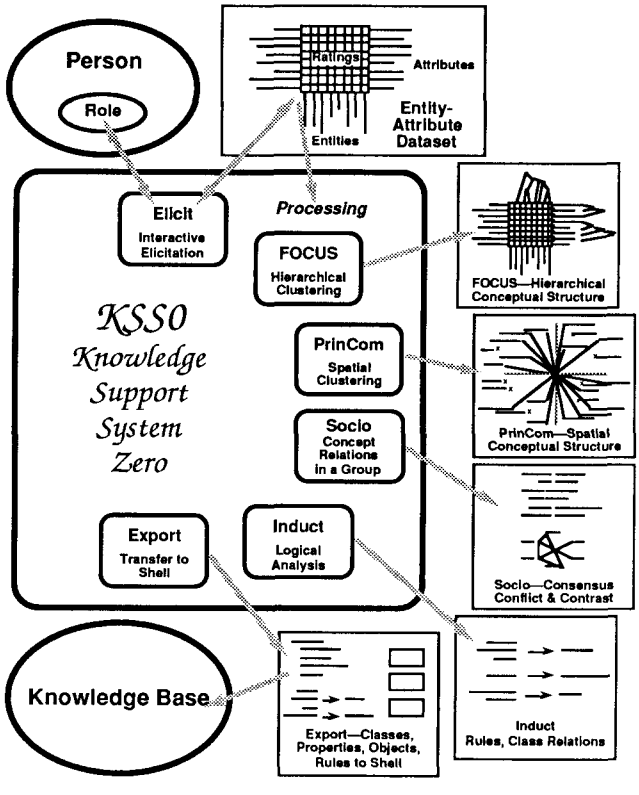

Fig.1 KSS0, an integrated knowledge acquisition system

The elicitation tools are based on Shaw's [13, 14] computer-based interviewing techniques extended through the use of graphical rather than numerical data entry. The visualization and direct manipulation of knowledge structures through the graphic and click-and-drag facilities of modern workstations is an important development at the user interface giving experts, knowledge engineers and clients improved access to the knowledge structures.

- *Elicit* accepts specifications of cases within a sub-domain and provides an interactive graphical elicitation environment within which the experts can distinguish cases to derive their constructs. The resultant class is continuously analyzed to provide feedback prompting the expert to enter further cases and constructs.

The visualization tools consist of an interactive interface to represent the abstractions derived from those cases in terms of hierarchical clusters using Shaw's [13] *FOCUS* algorithm, and relational diagrams such as a non-hierarchical conceptual maps derived through principal components analysis [17, 18] The objectives are to validate the raw domain knowledge and suggest further structure at a higher level through interactive topological induction [12]:

- *FOCUS* hierarchically clusters cases and constructs within a sub-domain prompting the experts to add higher-level constructs structuring the domain.
- *PrinCom* spatially clusters cases and constructs within a sub-domain prompting the experts to add higher-level constructs structuring the domain.

The group comparison tools consist of an interactive interface to represent the relations between the terminologies and conceptual systems of different experts, or experts and clients. The objectives are to determine the consensus, conflict, correspondence and contrast between different conceptual systems [16]:

- *Socio* compares the structures for the same sub-domain generated by different experts, or the same expert at different times or from varying perspectives.

The inductive part consists in the derivation of constraints within the conceptual structures through logical entailment analyses [1, 4, 5, 6, 9]. The objective is to suggest further structure at a higher level that translates into class inclusions or rules in the expert system shell:

- *Induct* induces logical entailments enabling the constructs of an case or the evaluations of a decision-making situation in a domain, to be derived from other constructs.

The generative part consists in the transformation of the knowledge analysis made by the previous tools into formalisms understandable by knowledge-based system shells such as *NEXPERT* [10, 11] and *Babylon* [2]:

- *Export* formats the specifications of sub-domains as classes, of cases as objects, of constructs as properties, and of entailments as methods, and transfers them to the performance tool.

The next section illustrates the use of the interactive elicitation and visualization tools in eliciting an adequate dataset or repertory grid, and shows the use of the induction and export tools in KSS0 to create a computational knowledge base from grid giving the required performance in an expert system shell.

A SAMPLE ELICITATION

The dataset used is this paper is taken from a paper by Cendrowska [1] on the inductive analysis of a set of ophthalmic data. The problem is to infer the type of lens to be prescribed by determining relevant qualities of the client. The example is simple and the results well-defined which makes it a good test case for the basic techniques. However, the requirements and results are sufficiently complex for the test case to be non-trivial and to demonstrate many major features of the knowledge acquisition tools and their application.

The *Elicit* facilities in the knowledge acquisition tool are designed to elicit the dataset at a stage when the expert may not be able to state the relevant constructs and a complete and consistent set of critical cases. This is done by helping him or her to describe the relevant constructs and critical cases to build up a repertory grid. He or she chooses "New Grid", specifies a dataset name of "Lens", moves to the Elicit status screen and specifies the domain, context and names for cases as shown in Figure 2. The expert leaves the default rating scale range of 1 through 9, allowing intermediate or "fuzzy" constructs to be entered. He or she then clicks on "Done", moves to the *Elicit* case selection screen, clicks "Done" again and moves to the *Elicit* case entry screen, entering some initial cases as shown in Figure 3.

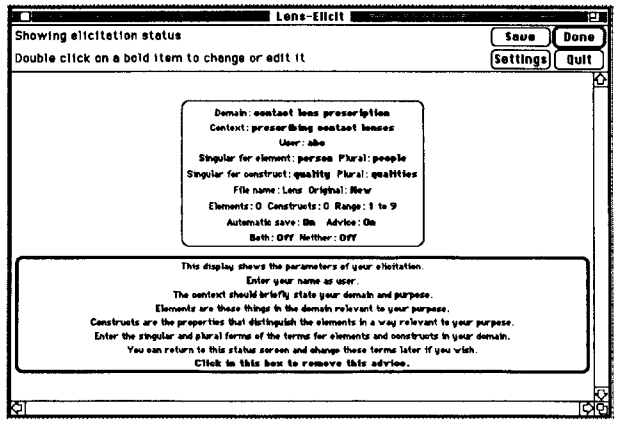

Fig.2 Status screen for contact lens knowledge elicitation

What initial cases does it make sense for the expert to choose? He or she should think in terms of those which represent critical features of the problem. For example, in the contact lens problem, the fact that clients with reduced tear production should not have lenses fitted should be represented, as should the fact that astigmatic clients in general have hard lens and non-astigmatic clients soft lens.

The expert may next choose to add some initial constructs that are clearly relevant to the problem. Figure 4 shows the constructs "not soft—soft", "not hard—hard", "normal—reduced" (tear production), and "not astigmatic—astigmatic", added as initial constructs.

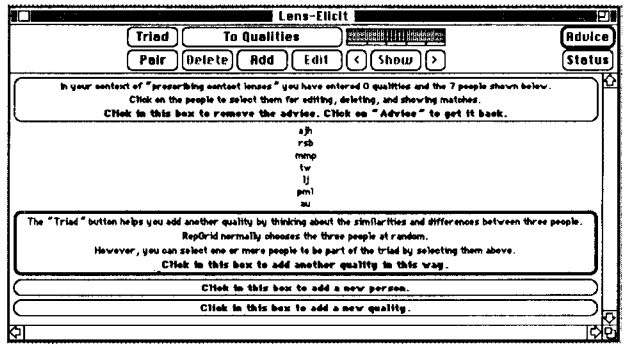

Fig.3 Initial clients entry

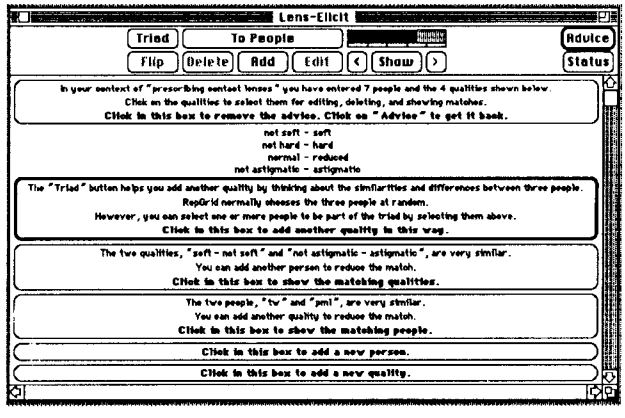

Fig.4 Initial qualities entry

Note that the expert has not assumed that the decisions to fit soft or hard contact lenses are mutually exclusive. Entering them as separate decisions allows for the possibility that either may be suitable for some types of patient. This example will show that this makes no difference to the behavior or performance of the knowledge acquisition tool or resulting expert system.

The constructs entered are not sufficient to solve the problem. How does the acquisition tool help the expert to realize this? The most important capability is that of having interactive analysis at any time during an elicitation. The expert can see his or her repertory grid as a knowledge base and come to understand what is missing. Figure 5 shows the tool's hierarchical cluster analysis of the initial dataset with 6 cases and 3 constructs. It can be seen that "astigmatic" is clustered with "hard" and "not astigmatic" with "soft". It is apparent from the case clusters that groups of clients are not being distinguished. Note from the data in Figure 5 that the expert has chosen to use only the two extreme values of the rating scale, for example, in the top row the value 1 means "hard" is selected and the value 9 means "not hard" is selected. He or she could also have expressed doubt about some cases by using intermediate values.

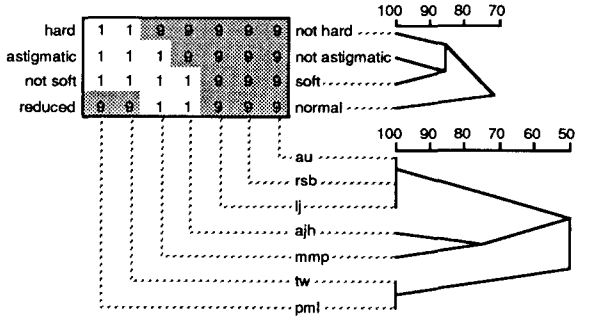

Fig.5 *FOCUS* analysis of initial data

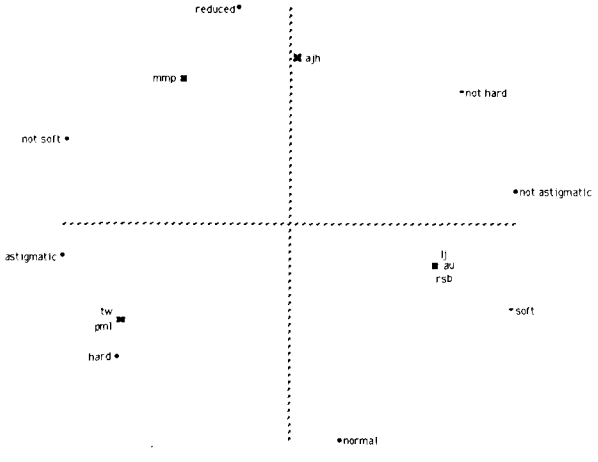

Fig.6 *PrinCom* analysis of initial data

Figure 6 shows the tool's alternative spatial analysis of the initial dataset. It is apparent that neither type of lens should be fitted if tear production is "reduced", and that "astigmatic" goes with "hard" and "not astigmatic" with "soft". Again the lack of distinction between groups of client is apparent. These two forms of analysis are closely related as will be apparent from checking that cases and constructs that are close together in one are also close together in the other. The hierarchical analysis has the advantage that the relation between the raw data and the clusters is apparent. The spatial analysis has the advantage that the relation between the cases and constructs is explicit.

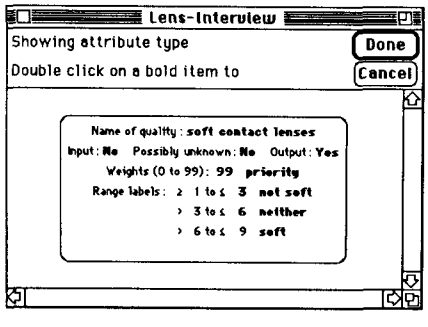

Fig.7 Setting the type of each quality

To obtain a logical analysis of the entailments between constructs the expert or knowledge engineer has first to specify that the type of the first two constructs is "Output", not "Input", through the *Type* screen shown in Figure 7.

Running the *Induct* logical entailment analyis program then produces the rules shown in Figure 8. These conclude that soft lenses should be fitted when tear production is normal and the astigmatism is not astigmatic, and that hard lenses should be fitted when tear production is normal and the astigmatism is astigmatic.

astigmatism=astigmatic -> soft contact lenses=not soft 100% 18.7%
tear production=reduced -> soft contact lenses=not soft 100% 32.7%

tear production=normal & astigmatism=not astigmatic -> soft contact lenses=soft 100% 7.87%

astigmatism=not astigmatic -> hard contact lenses=not hard 100% 26.0%
tear production=reduced -> hard contact lenses=not hard 100% 51.0%

tear production=normal & astigmatism=astigmatic -> hard contact lenses=hard 100% 8.16%

Fig.8 *Induct* analysis of initial data

How does the expert proceed to add more constructs and cases? Several possibilities are open. He or she may continue the elicitation using triads or breaking case and construct matches, or may think about the graphic representation of the knowledge base, or the rules produced so far. It does not really matter what prompts the expert to remember that age and prescription are also relevant variables and that there are some special cases that should be entered.

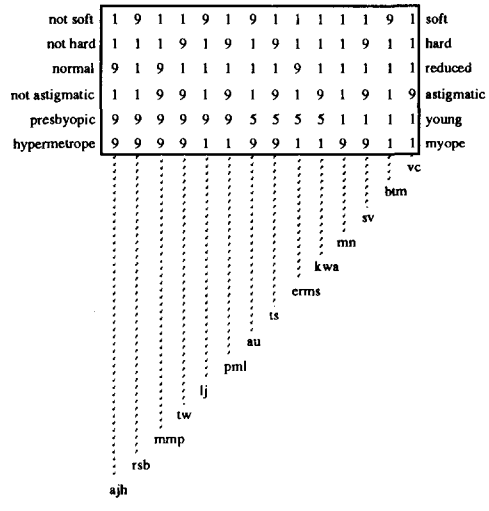

Fig.9 *Display* of final dataset

Note that it also does not matter if the expert enters a duplicate case. It will affect the result if the expert enters an incorrect decision. If this conflicts with an existing case it will be noticed quickly. Otherwise, it will be apparent through studying the graphic representations or rules.

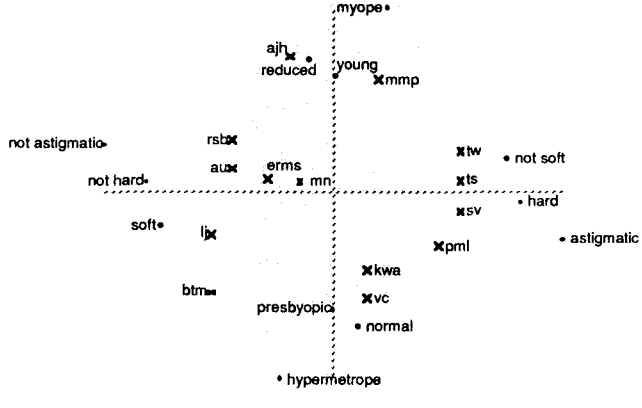

Fig.10 *PrinCom* analysis of final dataset

Figure 9 shows a display of the dataset when the expert has entered another 7 cases and the age and prescription constructs, and Figure 10 shows the spatial analysis of the grid.

The area behind the clients for whom neither form of lens are recommended has been shaded so that the structure of the knowledge base stands out. Clients rsb, au, lj and btm are standard cases for soft lenses, tw, ts, sv and pml are standard cases for hard. Clients ajh, mmp and erms are standard cases for no lenses because their tear production is reduced. Client mn is the anomalous case where soft lenses might otherwise be fitted. Clients kwa and xc are the anomalous cases where hard lenses might otherwise be fitted.

Figure 11 shows the rules produced by an *Induct* analysis of the dataset.

```
tear production=reduced -> soft contact lenses=not soft 100% 36.4%
age=presbyopic & prescription=myope -> soft contact lenses=not soft 100%
    51.0%

tear production=normal & astigmatism=not astigmatic & age=young -> soft
    contact lenses=soft 100% 8.16%
tear production=normal & astigmatism=not astigmatic & age=pre-presbyopic
    -> soft contact lenses=soft 100% 28.6%
astigmatism=not astigmatic & age=presbyopic & prescription=hypermetrope -
    > soft contact lenses=soft 100% 28.6%

astigmatism=not astigmatic -> hard contact lenses=not hard 100% 9.49%
tear production=reduced -> hard contact lenses=not hard 100% 36.4%
age=presbyopic & prescription=hypermetrope -> hard contact lenses=not
    hard 100% 51.0%
age=pre-presbyopic & prescription=hypermetrope -> hard contact lenses=not
    hard 100% 51.0%

tear production=normal & astigmatism=astigmatic & prescription=myope ->
    hard contact lenses=hard 100% 2.33%
astigmatism=astigmatic & age=young & prescription=hypermetrope -> hard
    contact lenses=hard 100% 28.6%
```

Fig.11 *Induct* analysis of final dataset

Suppose the expert does not realize that the knowledge base is adequate at this stage and adds further cases? Provided the correct lens prescription is given according to Cendrowska's original model, adding additional cases will have no effect on the rules produced by *Induct*. It may vary the cluster outputs since the additional cases will weight the analyses differently, but the broad structures apparent in these analyses will remain the same. If the expert adds further constructs in principle the *Induct* analysis will not be affected. However, it is possible for an arbitrary construct, such as hair color, to correlate highly with a significant construct or the decision constructs. On a small dataset of cases this may result in the arbitrary construct appearing to be significant to *Induct*. Increasing the number of cases should get rid of spurious correlations. However, it is probably of much greater importance that the expert notices a spurious construct appearing in the rules and decides to remove it. The acquisition tool is primarily designed to support the transfer of expertise by interviewing experts. It also has powerful techniques for statistically sound empirical induction. However, the number of cases required to be entered increases very rapidly with the proportion of incorrect decisions and irrelevant constructs, and it is better for the expert to clean up the data than to rely on statistics applied to poor quality data.

CONCLUSION

The example given in this paper demonstrates how a dataset representing an expert's conceptual structure for a problem domain can be elicited in a knowledge acquisition tool and converted automatically to a knowledge base that performs as expert system solving the problem in an expert system shell. In large scale system development knowledge structures for different sub-domains have to be developed to support a complete system. The acquisition tool can be used to develop and validate knowledge structures for each of these sub-domains. The modular and well-structured knowledge bases produced can then be combined through appropriate control structures.

Understanding and communicating expertise are not easy tasks. Knowledge acquisition tools are designed to support and guide the expert and knowledge engineer in that understanding and communication. Any real system development involves exploration, false starts, inadequate datasets, and so on. The expert has to learn the skills of expertise transfer through trial and error. A good tool provides an interesting, friendly and supportive environment for that learning experience. Once that learning has converged upon a solution, the tool can also provide formally correct and well-structured mechanisms for transferring the solution to a knowledge base driving an expert system shell.

ACKNOWLEDGEMENTS

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

REFERENCES

[1] Cendrowska, J. (1987) An algorithm for inducing modular rules. **International Journal of Man-Machine Studies 27** (4), 349-370 (October).

[2] Christaller, T., di Primio, F. & Voss, A. (1989). **Die KI-Werbank Babylon: Eine Offene und Portable Entwicklungsumgebung für Expertensysteme**. Bonn: Addison-Wesley.

[3] Gaines, B.R. (1987). Rapid prototyping for expert systems. Oliff, M.D., Ed. **Intelligent Manufacturing: Proceedings from First International Conference on Expert Systems and the Leading Edge in Production Planning and Control**. pp.45-73. Menlo Park, California, Benjamin Cummins.

[4] Gaines, B.R. (1989) An Ounce of Knowledge is Worth a Ton of Data: Quantitative Studies of the Trade-Off between Expertise and Data based on Statistically Well-Founded Empirical Induction. **Proceedings of 6th International Workshop on Machine Learning**, pp.156-159. San Mateo, California: Morgan Kaufmann (June).

[5] 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).

[6] Gaines, B.R. & Shaw, M.L.G. (1986). Induction of inference rules for expert systems. **Fuzzy Sets and Systems**, 8(3), 315-328 (April).

[7] Gaines, B.R. & Shaw, M.L.G. (1987). Knowledge support systems. **Proceedings of ACM MCC-University Research Symposium**. Austin, Texas: MCC. pp.47-66.

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

[9] Quinlan, J.R. (1987) Simplifying decision trees. **International Journal of Man-Machine Studies 27** (3), 221-234 (September).

[10] Rappaport, A. (1987a) Multiple-problem subspaces in the knowledge design process. **International Journal of Man-Machine Studies 26**(4), 435-452 (April).

[11] Rappaport, A. (1987b) Cognitive primitives. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Second AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.15-0-15-13. Banff (October).

[12] Rappaport, A. & Gaines, B.R. (1988). Integration of acquisition and performance systems. **Proceedings of the Third AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.25-1-25-20. Banff (November).

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

[14] Shaw, M.L.G., Ed. (1981) **Recent Advances in Personal Construct Technology**. London: Academic Press.

[15] Shaw, M.L.G. & Gaines, B.R. (1987) KITTEN: Knowledge Initiation & Transfer Tools for Experts & Novices. **International Journal of Man-Machine Studies 27**(3), 251-280 (September).

[16] Shaw, M.L.G. & Gaines, B.R. (1988). A methodology for recognizing consensus, correspondence, conflict and contrast in a knowledge acquisition system. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Third AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.30-1-30-19. Banff (November).

[17] Slater, P., Ed. (1976). **Explorations of Intrapersonal Space: Volume 1**. London: John Wiley.

[18] Slater, P., Ed. (1977). **Dimensions of Intrapersonal Space: Volume 2**. London: John Wiley.