Applied Expert Systems, E. Turban and P.R. Watkins (Editors)

EXPERT SYSTEM PROTOTYPING AS A RESEARCH TOOL

Daniel E. O'Leary

Graduate School of Business
University of Southern California
Los Angeles, California 90089-1421

Expert system prototyping is analyzed as a research tool. Because expert system prototypes allow for the storage and manipulation of symbolic information they can be used for research projects for which other tools are inappropriate. The experimental bases of expert system prototypes and experimental controls for expert system prototypes are analyzed. The use of expert system protoytpes as research databases also is discussed. Then the kinds of research issues that expert system prototypes can be used to examine and research issues in the use of prototypes as a research tool are analyzed.

1. INTRODUCTION

There is a well-established tradition of prototyping in engineering as a means of testing a product in order to determine if a product can be built and to ascertain if there are any operating anomalies in the design. Prototyping the system provides a "proof of concept."

However, prototyping is no longer the exclusive property of engineering. In recent years, prototyping has become accepted as a means of designing information systems (Earl [1]), decision support systems (Henderson and Ingraham [2]) and expert systems (Hayes-Roth et al. [3]). The extensive bibliography in Jenkins and Fellers [4] provides evidence of the growing use of prototyping as an application system design and development methodology.

Still *expert system* prototyping is relatively new. Its recent development is due, in part, to a number of changes in software and hardware. The development of artificial intelligence languages and expert systems shells and recent developments in hardware, e.g., personal computers and artificial intelligence work stations, have brought artificial intelligence-based computer power directly to the user. Its development also is due, in part, to the successful development of working expert systems (e.g., MYCIN).

However, the newness of expert system prototyping and the operational successes of building prototype expert systems may have camouflaged the value of prototype expert systems as a research tool. Accordingly, this paper is concerned with analyzing expert systems prototyping as a *research tool*, rather than as just an operational design and development tool. If expert system prototypes are to be regarded as a research tool, then there needs to be an understanding of

- the basic process and output of expert system prototyping (section 2),
- the unique advantages offered by expert system prototyping (section 3),
- the relationship of expert system prototyping to other forms of experimental analysis (section 4),
- some of the experimental controls useful in expert system prototyping (section 5),
- use of expert system prototypes as research databases (section 6),
- the kinds of issues that we can expect expert system prototypes to be useful in investigating (section 7),
- some of the limitations of expert systems prototypes (section 8), and
- some research questions into using expert systems prototypes as a research tool (section 9).

2. EXPERT SYSTEMS PROTOTYPES

An expert systems prototype is a small version of the expert system that is to be developed. Generally, such systems have few frills. The prototype provides a preliminary solution approach to the decision under consideration, without a substantial resource commitment.

In applications, the prototype is part of the process in the development of a larger expert system for a particular application. Expert systems prototypes provide an opportunity to test assumptions about the knowledge base, inference strategies of the expert and other characteristics of the system. In many cases the expert system prototype is thrown away, after these issues become better understood. Then, a new system is built using the knowledge garnered from the prototype development process.

Expert systems prototyping consists of designing and developing that computer program without an extensive preliminary examination of the problem area. This means iteratively discovering the expert's knowledge and developing the system's knowledge, i.e., knowledge acquisition. Prototyping is not mechanical, but generally is exploratory. Because of the iterative approach of knowledge acquisition, generally knowledge is discovered in layers.

Expert systems prototypes generally are designed for decision making situations requiring a certain amount of "intelligence." Typically, prototyping is done using an expert system shell to ease the development process, e.g., EXSYS, or an artificial intelligence language, e.g., LISP. Such artificial intelligence languages and expert system shells are designed to manipulate symbolic, rather than numeric variables.

3. PROTOTYPING AS A UNIQUE RESEARCH TOOL

Typically, statistical and mathematical representations are used to analyze research questions. Unfortunately, in some situations, mathematical or statistical representations of problems cannot be easily made. Generally, this is because the variables involved in the models do not take on numeric values or situations are not easily structured in a statistical or mathematical fashion. It may be that prototyping is the only type of tool that can be used to represent a problem because of the capabilities of representing symbolic information in expert systems.

3.1. Understanding and Learning from Expert Systems Prototypes

Piaget [5] noted that "All mathematical ideas begin by a qualitative construction before acquiring a metrical character." He also noted that by moving too rapidly from the qualitative structure to the quantitative or mathematical formulation often affects the ability to understand mathematics.

Since expert system prototypes allow the use of symbolic variables, the user does not have to quantify the knowledge and the decision making processes as part of the modeling process or move from the qualitative to the quantitative until the user is ready (Howe [6]). However, prototyping does allow the user to structure the problem in a manner that ultimately may lead to quantification.

An example of this learning process recently occurred in the development of an expert system prototype for an auditing problem. The firm trying to solve the problem developed and threw away two expert systems prototypes without finding a satisifactory solution. Then, based on the understanding gleaned from those prototypes, the developers realized that the problem could be solved as a linear program.

3.2. Resource Requirements to Employ Expert Systems Prototyping

Expert systems prototypes also provide the opportunity to analyze problems without a substantial resource commitment of the researcher. Watkins and O'Leary [7], in a field study, examined the ability of domain experts to develop expert system prototypes using an expert system shell. They found that domain experts, with little computer experience, were able to rapidly develop expert systems prototypes with little formal training. Thus, this would suggest that researchers, also experts in particular domains, would be able to rapidly develop expert systems prototypes.

3.3. Theory Biases of Expert Systems Prototypes

However, as with most research tools, researcher bias can be embedded into the expert system prototype. The bias can occur in at least three different ways. First, the researcher can bias the system based on the choice of "first principles" or theory on which the system is based. As noted by Davis [8, p. 403] in a discussion of electronic trouble shooting, first principles refer to "... an understanding of the structure and function of the devises they are examining." The use of an alternative theory could have a substantial impact on the knowledge used in the system.

Second, the knowledge representation scheme may affect the expert system prototype. Wilensky [9] reviews the importance of knowledge representation.

Third, the inference engine impacts the manner in which the knowledge base is analyzed by the system, in response to inquiries made to the system prototype. Many inference engines simply search through the knowledge base in a "top to bottom" or "side to side" manner (e.g., Hayes-Roth et al. [3]). Few would argue that humans process information in that manner.

In any case, the biases of the researcher can be controlled, in part, by implementing the appropriate research methodologies and controls. These are discussed in more detail in sections 4 and 5.

3.4. Theory Laden Tools

Since expert system prototypes can accommodate symbolic variables, they are a very flexible tool. As a result, prototyping can offer an alternative tool to investigate a number of research situations.

Often, it can be used as an alternative tool to investigate questions where the original tool(s) used in the research is "theory laden." For example, linear regression is used in conjunction with the Brunswik Lens model (Dudycha and Naylor [10]). Measurements of human performance are variables in a linear regression model. In this case linear regression is a theory laden tool in the study of human information processing. Expert systems prototyping offers an alternative tool that is not wed to that approach to understanding human information decision making.

4. EXPERIMENTAL BASIS FOR KNOWLEDGE ACQUISITION FOR EXPERT SYSTEM PROTOTYPES

The development of the expert system prototype requires the acquisition of the knowledge base for the system. The approach used to gather knowledge will impact the knowledge base of the expert system prototype that is developed (e.g., Burton et al. [11]). As a result, if expert system prototypes are to be used as a research tool then it is critical to ensure that the knowledge is gathered using a systematic and accepted research methodology.

Further, identifying expert systems prototyping as making use of particular experimental techniques is important. This allows the use of controls and established literature associated with each of these methodologies in knowledge acquisition for expert system prototypes.

The primary knowledge acquisition methods potentially include using any of a number of experimental techniques including (e.g., Simon [12]): Deductive Reasoning, Case Study, Participant Observer, Content Analysis, Simulation. Other experimentally accepted techniques, such as time series analysis or surveys are rarely used in expert systems analysis because these approaches do not generate enough detail or qualitative information for knowledge acquisition.

4.1. Deductive Reasoning

As noted by Simon [12, pp. 203-204], "The principle of deductive reasoning for obtaining knowledge is that, if A is true and if B is true then, under specified conditions one can safely say that C is true". Thus, deductive reasoning takes the form of "If ... Then ..." rules, which probably is the best understood form of knowledge representation. This suggests that expert systems prototypes are an ideal vehicle for encoding and analysis of deductive reasoning. Unfortunately, decision makers rarely think and reason using "If ... Then ..." rules (Biggs et al. [13]).

4.2. Case Study

At another level, knowledge acquisition for expert system prototypes is analogous to a case study. Oftentimes in the development of an expert system prototype a particular decision making situation is used to elicit the knowledge require to solve a problem. Clearly, "case study" could be replaced with "expert system prototype" in the following definition of a case study.

> The case study is ... the method of choice when you want to obtain a wealth of detail about your subject. You are likely to want such detail when you do not know exactly what you are looking for. The case study is therefore appropriate when you are trying to find clues and ideas for further research; in this respect, it serves a purpose similar to the clue-providing function of expert opinion. (Simon [12, p. 206])

Case study methodology can lack rigor without appropriate attention to its implementation. Case study methodology has been discussed by Yin [14] and others.

4.3 Participant Observer

In the development of some expert systems, developers have found the need to become "near experts" before they could develop an appropriate system (Lethan and Jacobsen [15]). Alternatively, experts may be trained to develop an expert system prototype (e.g., Watkins and O'Leary [7]). This is suggestive of the participant observer methodology. As noted by Simon [12, p. 207], "If you wish to understand the full complexity of a case situation in social science, you may have no alternative but to get yourself involved ...." Thus, the participant observer approach may be necessary both in order to understand the problem and to solicit the knowledge.

4.4. Content Analysis

Content analysis is aimed at objectively and systematically analyzing symbolic information ([Holsti [16]). As noted by Simon [12, p.212],

"The content analyst sets up various classification schemes, which he then applies to speeches or writings. These classifications either count particular kinds of words or ideas, or they measure the amount of words or time that is devoted to particular ideas."

A straight content analysis approach may not be appropriate. Since content analysis is limited to counting occurences, it may be limited in its ability to elicit key heuristic decision rules. As a result, other forms of content analysis have been developed.

One of the primary content analysis tools used in developing expert system prototypes is protocol analysis (e.g., Newell and Simon [17]). Protocols may derive from either audio or video tape. Protocol analysis has been used in the development of a number of expert system prototypes (e.g., Meservy [18]).

An alternative method of content analysis is the linguistic approach of Frederiksen [19]. That approach employs a grammatical structure that represents the logical and semantic structure of knowledge acquired from discourse.

4.5. Simulation

Simulation is a term that includes any laboratory experiment or game study (Simon [12, p. 215]). Games or experiments are used to simulate the way that the participants would behave in various aspects of real life competition (e.g., Rowe et al. [20]). The researcher can vary the simulation to ascertain what actions are taken under what circumstances. This knowledge can then be embedded in the context of the knowledge base of an expert system prototype. This technique has been used in a few situations to elicit knowledge for an expert system prototype (e.g., Bouwman [21]).

5. EXPERIMENTAL CONTROLS IN KNOWLEDGE ACQUISITION

If an expert system prototype is used as a research tool then in order to ensure that the system is developed in a rigorous manner, the developer must implement the system in an environment with the appropriate experimental controls. There are a number of controls that are associated with each of the above forms of knowledge acquisition. In addition, there are some other experimental controls that derive from the uniqueness of the need to gather knowledge, as opposed to numeric data.

First, whether unobtrusive or obtrusive means of knowledge acquisition is preferred likely is a function of the research project or the group of experts. Oftentimes, information should be gathered using unobtrusive means. The development of one expert system prototype used phone taps to gather the knowledge without altering the expert and question askers' behavior. Alternatively, sometimes obtrusive means may provide the best method of obtaining the necessary knowledge (e.g., Sphilberg et al. [22]).

Second, Kolodner and Reisbeck [23] note that memory organization is constantly changing to reflect new events. As a result, the expert may alter his approach to problem solving based on the analysis of problem by the researcher. To test this the researcher can perform both pretests and posttests. For example, the researcher can use a problem solved before developing the knowledge base and an analoguous problem solved after developing the the knowledge base. Then the researcher can compare the solution methodologies and knowledge to ascertain if the expert altered the initial approach. If the two approaches are not the same then the knowledge gathered prior to this latest knowledge acquisition may not reflect the expert's current approach to problem solving.

Third, the physical location of the expert during the knowledge acquisition may impact the system. If the system was not developed on site (e.g., at the laboratory) then some contingent knowledge that the expert associated with the site may not be placed in the knowledge base. Alternatively, if the knowledge is gathered on site then there is little control over the environment (e.g., interruptions and the amount of time that the expert can spend).

6. EXPERT SYSTEM PROTOTYPES AS RESEARCH DATABASES

Expert system prototypes can be used as databases for researchers in at least two different manners. Each system can be treated as an independent database or they can be aggregated with other prototypes. Then these databases can be used to test theories or generate theories.

6.1. Experimental Use of Single Expert System Prototype Knowledge Bases

Cook and Campbell [24] use the term "passive" observation to refer to studies where experiments are performed with existing data. Prototype expert systems knowledge bases form "databases" that can be analyzed. Researchers can investigate the knowledge base of an expert system prototype in the same sense that a user would consult the system for a solution. In this case, each use of the expert system prototype is used as a single data point. As noted in Winston [25, p. 2]:

> Computer programs exhibit unlimited patience, they require no feeding and they do not bite. Moreover, it is usually simple to deprive a computer program of some piece of knowledge in order to test how important that piece really is. It is impossible to work with animal brains with the same precision.

6.2. Experimental Use of Developed Prototypes

In a similar manner, if the researcher has access to information about a number of expert system prototypes then additional inferences can be made, based on commonalities between the systems. For example, this approach has been used by O'Leary [26] in the analysis of validation of expert systems. In this situation, each prototype becomes a single data point.

6.3 Theory Testing vs. Theory Generation

The use of expert system prototypes as a database can be used either for Theory Testing or Theory Generation. As noted in Winston [25, p. 2]:

> Computer models force precision. Implementing a theory uncovers conceptual mistakes and oversights that ordinarily escape even the most miticulous researchers. Major roadblocks often appear that were not recognized as problems even before beginning the cycle of thinking and experimenting.

Expert system prototypes can be used to *test theories* about the knowledge base. For example, given an expert system prototype, the user can test theories by inquiry into the knowledge base or experiments with the knowledge base. For example, Meservy [18] developed an expert system prototype, in part, to test some hypotheses of auditor behavior.

Alternatively, expert system prototypes can be used to *generate theories.* Some of the first expert system prototypes were developed to solve problems for which there was no established theory. In effect, the development of the knowledge base was an assemblage of knowledge not related by any theory (e.g., Davis [8]). Thus, interaction by the researcher with the knowledge base allows the possibility of the development of theory.

However, as noted by McDermott [27], it is more difficult to develop an expert system if there is no existing theory to guide the development efforts. Accordingly, it is likely to be more difficult to develop a system for theory generation.

7. WHAT KINDS OF RESEARCH ISSUES CAN AN EXPERT SYSTEM PROTOTYPE BE USED TO INVESTIGATE?

Expert system prototypes can be used as a tool or vehicle to understanding particular problems, understanding the process of solving a problem, developing a proof of concept that a particular idea can be implemented as a computer program or analyzing the use of expert systems prototyping as a research tool.

7.1. Understanding Particular Problems Using Prototype Expert Systems

Expert system prototypes may provide better methods of decision making than other tools. For example, if a prototype expert system could be developed to improve bankruptcy prediction, then that model would likely be of interest to those doing research in bankruptcy prediction.

In addition, expert system prototypes summarize some of the knowledge in a given area. This can lead to a structuring of previously unstructured knowledge. Thus, the development of a prototype may move unstructured or poorly structured problems into a more structured framework. In those situations where there is no central storage location of the knowledge or where there is no structure to the knowledge, simply the development of the prototype may provide a research opportunity.

Expert system prototypes also can be used to assess the "fineness" of the knowledge about the problem. Fineness is a characteristic of information structure. Marschak and Radner [28] say that

> . . . given two information structures, n_1 and n_2, that n_1 is as fine as n_2 if n_1 is a subpartition of n_2; That is, if every subset in n_1 is in contained in some set in n_2. (Thus, n_1 tells us all that n_2 can tell, and possibly more besides.) If n_1 and n_2 are distinct, and n_1 is as fine as n_2, then we shall say that n_1 is finer than n_2.

Few other tools that the researcher can use would even contain or have the ability to capture differences in the fineness in the knowledge base. However, because the knowledge base captures symbolic information, it is likely to capture the fineness of the information. Accordingly, the knowledge base of an expert system prototype can be used to analyze the comparative fineness of information in the knowledge base.

Fineness may reflect the understanding of the information in the knowledge base. If the knowledge is not equally fine throughout the knowledge base, then this may lead to efforts to better "understand" the part that is not as fine. Alternatively, the existence of different levels of fineness in the knowledge base may reflect the state of nature of the knowledge. Some knowledge may come in "bigger chunks" (e.g., Newell and Simon [17]).

7.2 Understanding Human Intelligence and Expertise

Alternatively, developing a prototype can result in a better understanding of how humans make decisions in general or in particular types of situations. Further it may provide insight into the notion of expertise.

As noted by Winston [25, p. 1], "Making computers intelligent helps us understand intelligence." In particular, as noted by Winston [25, p. 2],

> Computer metaphors aid thinking. Work with computers has led to rich new language for talking about how to do things and how to describe things. Metaphorical and analogical use of concepts involved enables more powerful thinking about thinking.

Intelligence includes such issues as memory structures, causal relationships and problem formulation. Expert systems prototypes allow researching such issues (e.g., Biggs and Selfridge [29]).

In addition, part of the research on expertise lies in understanding the amount or kind of knowledge or processing of knowledge necessary to develop good decisions. As noted by Winston [25, p. 4]

> Computer implementations quantify task requirements. Once a program performs a task, upper bound statements can be made about how much information processing the task requires.

Since expert system prototypes summarize knowledge gathered from an expert and they are computer programs, they allow research into such evaluations of task quantification.

7.3 Proof of Concept

Alternatively, an expert system can be used to provide a proof of concept. That is, although it may be asserted that a particular technique can be used in an expert system, until it actually is programmed there is uncertainty as to feasibility.

Theoretical developments in artificial intelligence concepts require proof of concept in order to ensure that they can be adapted to particular domains.

This may include research in the areas of uncertainty representation and in automated knowledge acquisition by, e.g., reading text. For example, the expert system GC/X (Biggs and Selfridge [29]), tests the use of some concepts required to read and organize new knowledge in accounting systems.

7.4 Analyzing the Use of Expert System Prototypes as a Research Tool

Expert systems prototypes also can be used to examine various facets of using prototypes as a research tool. Although this section reviews two particular applications, other research questions are summarized in section 9.

In the case of information fineness, if the information is not equally fine throughout the model then this may affect the parameterization of the model, which also is likely to be a research issue. For example, in a rule-based system which employs weights on the rules to represent uncertainty or strength of belief, should two rules that reflect information that are not equally fine but have the same likelihood of occurring, have the same weight?

Alternatively, Langley, et al., [30, p. 44] argue that "In recent years there have been several challeges to the idea that the process of theory generation is mysterious and inexplicable." They argue that there are "rules" of discovery. This opens the process of discovery to, e.g, other expert systems. Langley et al. [30] present analysis of this issue and discussion of prototypes that can discover theory.

8. LIMITATIONS OF EXPERT SYSTEM PROTOTYPING

Although prototyping can be a useful research tool, there are some limitations of prototypes that should be recognized. This section discusses some of these limitations, in addition to those mentioned earlier.

First, *a working prototype is not prima facie evidence of a successful representation of the expert* (Winston [25]). Although the model "works" it may not be a very good representation of expertise and decision making capabilities that the expert used.

Second, *problems can be overlooked, buried or shirked* (Haugeland [31]). In order to get a working model certain simplifications must be made. If critical problems are avoided then the prototype may be representative of only a small portion of the decision making capabilities of the expert.

Third, *domains may be unrealistically or artificially restricted or unrealistically prestructured* (Haugeland [31]). If the domain or problem is too narrowly stated then the system may not be important or may not solve the problem of interest.

Fourth, *prototype solutions may be suboptimal* (Henderson and Ingraham [2]). The prototype may provide an inappropriate, incomplete or less than an optimal solution to the problem it was designed to solve. If the system is suboptimal then studying the system may not yield what the researcher is interested in.

Fifth, when developing a prototype *it may be very difficult to know when to stop building the prototype.* There are few or no stopping rules for the overall quality of the prototype.

Sixth, *there are few ways to determine if a model is parameterized correctly.* In least squares analysis, the parameterization that minimizes the sum of the squares of the errors is used as the basis to choose the coefficients on the variables. There is no such analogous method for use in expert systems prototypes.

Seventh, *expert system prototypes may include a researcher bias.* As noted earlier, the choice of first principles, or the choice of using an alternative approach when first principles exist, may reflect the researchers' bias. By employing research methods, and the appropriate controls, these biases can be mitigated to a certain extent.

9. SOME RESEARCH QUESTIONS INTO USING EXPERT SYSTEM PROTOTYPES AS A RESEARCH TOOL

This paper is only a first layer in the analysis of the use of expert system prototypes as research tools. Based on the analysis in this paper, a number of research issues about the use of expert system prototypes as research tools have been addressed. This section summarizes some of the questions elicited above.

First, what are the experimental controls that can be implemented in prototypes? Should we do a pre and a post test on the expert to see if the model development process has changed the expert's knowledge?

Second, what types of experimental controls are unique to expert systems prototyping?

Third, how do we measure the fineness of information in a knowledge base? What is the impact of different levels of fineness on the parameterization of an expert system's weights?

Fourth, what relationship is there between the use of knowledge in an expert system prototype and that of a human?

Fifth, in some research tools there are established measures of "goodness." What are some measures of goodness that can be used to evaluate the parameters in expert systems?

Sixth, what mechanisms are available to ascertain when an expert system prototype is "good enough" to stop building the model and use it to investigate the research question?

Seventh, to what extent is expert systems prototyping a cost beneficial research tool to employ? That is, how easy is it for researchers to learn expert systems prototyping as a research tool.

Eighth, which forms of knowledge acquisition are "preferred."

At this point, there has been little presented in the way of solving these issues, and many open questions remain. However, the purpose of this paper has been to elicit and summarize these issues.

10. SUMMARY

This paper has argued in favor of expert systems prototyping as a research tool. Prototyping can be used in those situations where the research question requires the investigation of nonnumeric, symbolic information. Expert systems prototyping can be an intermediate step in developing a more structured representation of the problem. Further, expert systems can be used in those situations where an alternative research tool is desired to those that may be theory laden.

The development of the knowledge for an expert system can be done in a number of established manners (e.g., case study) each of which has its own set of controls. In addition, there are other controls that may be somewhat unique to expert systems prototyping.

Once an expert system prototype has been established it can be used as a database for either Theory Generation or Theory Testing. Further, expert system prototypes can be used to develop alternative solution methodologies, study human intelligence and expertise and develop proof of concept.

Finally, the paper summarizes some of the limitations of expert system prototypes and develops some of the research questions associated with the use of expert system prototypes as research tools.

ACKNOWLEDGEMENT

The author would like to thank Alan Rowe and Efraim Turban for their comments on an earlier version of this paper. Of course, any limitations remain due to the author.

REFERENCES

[1] Earl, M., "Prototyping Systems for Accounting and Control," *Accounting, Organizations and Society,* Vol. 3, No. 2, pp. 161-170.

[2] Henderson, J. and R. Ingraham, "Prototyping for DSS: A Critical Appraisal," in Ginzberg, et al. [31], pp. 79-96 (1982).

[3] Hayes-Roth, F., D. Waterman, and D. Lenat, *Building Expert Systems*. Reading, Mass.:Addison-Wesley, 1983).

[4] Jenkins, A. and J. Fellers, "An Annotated Bibliography on Prototyping," Graduate School of Business, Indiana University, Institute for Research on the Management of Informations, Working Paper W613, (1986).

[5] Piaget, J., *To Understand is to Invent*. New York:Grossman, 1973.

[6] Howe, J., "Learning Through Model Building," in D. Michie, *Expert Systems in the Micro-electronic Age*. Edinburgh University Press, 1979.

[7] Watkins, P. and D. O'Leary, "Knowledge Acquisition for Small Scale Expert Systems from Consulting Experts: A Field Study Approach," in *Proceedings of the First European Workshop on Knowledge Acquisition*, September, 1987.

[8] Davis, R., "Reasoning From First Principles," *International Journal of Man-Machine Studies,* Vol 19, (1983), pp. 9-16.

[9] Wilensky, R., "Knowledge Representation -- A Critique and a Proposal," in Kolodner and Reisbeck [22] (1986).

[10] Dudycha, L. and J. Naylor, "Characteristics of the Human Inference Process in Complex Choice Behavior," *Organizational Behavior and Human Performance,* Vol. 1 (1966) pp. 110-128.

[11] Burton, A., N. Shadbolt, A. Hedgecock, and G. Rugg, "A Formal Evaluation of Knowledge Elicitation Techniques for Expert Systems," in *Proceedings of the First European Workshop on Knowledge Acquisition*, September, 1987.

[12] Simon, J. L., *Basic Research Methods in Social Science*. New York: Random House, 1978).

[13] Biggs, S., W. Messier, and J. Hansen, "A Descriptive Analysis of Computer Audit Specialists' Decision Making Behavior in Advanced Computer Environments," *Auditing,* (Spring, 1987) Vol. 6, No. 2, pp. 1-21.

[14] Yin, R. K., "The Case Study Crisis: Some Answers," *Administrative Science Quarterly* (March 1981) Vol 26, pp. 58-65.

[15] Lethan, H. and H. Jacobsen, "ESKORT -- An Expert System for Auditing VAT Accounts," *Proceedings of Expert Systems and Their Applications,* Avignon, France 1987.

[16] Holsti, O., *Content Analysis for the Social Sciences and Humanities.* Reading, Mass.:Addison-Wesley, 1969.

[17] Newell, A. and H. Simon, *Human Problem Solving.* Englewood Cliffs, N. J.:Prentice-Hall, 1972.

[18] Meservy, R., "Auditing Internal Controls: A Computational Model of the Review Process," Unpublished Ph. D. Dissertation, The University of Minnesota, 1985.

[19] Frederiksen, C. H., "Cognitive Models and Discourse Analysis," in Cooper, C. R. and Greenbaum, S. (Eds.), *Studying Writing Linguistic Approaches* Beverly Hills, Ca.:Sage Publications, 1986, pp. 227-267.

[20] Rowe, A., R. Mason, and K. Dickel. *Strategic Management and Business Policy.* Reading, Mass.:Addison-Wesley, 1985.

[21] Bouwman, M., "Human Diagnostic Reasoning by Computer," *Management Science* (June 1983) pp. 653-672.

[22] Shpilberg, D., L. Graham, and H. Schatz. "ExperTAX: An Expert System for Corporate Tax Planning," *Expert Systems,* (July, 1986).

[23] Kolodner, J. and C. Reisbeck, *Experience, Memory and Reasoning.* Hillsdale, N. J.:Lawerence Erlbaum Associates, 1986.

[24] Cook, T. and D. Campbell. *Quasi-Experimentation.* Skokie, Il.:Rand McNally, 1979.

[25] Winston, P., *Artificial Intelligence.* Reading, Mass.:Addison-Wesley, 1984.

[26] O'Leary, D., "Validating Expert Systems," *Decision Sciences* (Summer 1987) pp. 468-486.

[27] McDermott, J., "Comments in, Knowledge-Based Expert Systems," The Second Artificial Intelligence Satellite Symposium, Texas Instruments (1986).

[28] Marschak, and Radnor, *Economic Theory of Teams*. New Haven:Yale University, 1972.

[29] Biggs, S. and M. Selfridge, "GCX: An Expert System for the Auditor's Going-Concern Judgment," Unpublished Presentation, American Accounting Association, New York, (1986).

[30] Langley, P., H. Simon, G. Bradshaw, and J. Zytkow, *Scientific Discovery.* Cambridge, Mass:The MIT Press, 1987.

[31] Haugeland, J., Unpublished Presentation, Forum on Artificial Intelligence in Management, Virginia, (1986).

[32] Ginzberg, M., W. Reitman, and E. Stohr, *Decision Support Systems.* New York:North-Holland, 1982.