ICIS 1987 Proceedings

International Conference on Information Systems (ICIS)

1987

SKILLS AND TECHNIQUES FOR KNOWLEDGE ACQUISITION: A SURVEY, ASSESSMENT, AND FUTURE DIRECTIONS

Jack W. Fellers
Indiana University

Recommended Citation

Fellers, Jack W., "SKILLS AND TECHNIQUES FOR KNOWLEDGE ACQUISITION: A SURVEY, ASSESSMENT, AND FUTURE DIRECTIONS" (1987). *ICIS 1987 Proceedings*. 9.
http://aisel.aisnet.org/icis1987/9

SKILLS AND TECHNIQUES FOR KNOWLEDGE ACQUISITION: A SURVEY, ASSESSMENT, AND FUTURE DIRECTIONS

Jack W. Fellers
Operations and Systems Management
Indiana University

ABSTRACT

In recent years there has been a tremendous increase in the development of expert systems in organizations. This increased development is straining the already limited supply of qualified expert system developers. These expert system developers have come to be known as knowledge engineers, and their job as knowledge engineering. The process of knowledge engineering is divided into two tasks: knowledge acquisition and expert system construction. Knowledge acquisition has been defined as "The process of extracting, structuring, and organizing knowledge from several sources, usually human experts, so it can be used in a program" (Waterman 1986, p. 392). This process of knowledge acquisition has been identified as the "bottleneck" that currently constrains the development of expert systems.

This paper summarizes what is known about the skills required and the techniques utilized in the knowledge acquisition process. Due to the similarities that exist between expert systems and traditional systems development, the literature pertaining to traditional information requirements determination and to systems analysts will be utilized to guide this exploration. Case study reports of actual expert system development projects and the practitioner literature will also be referenced.

Given the lack of research in this area, future research directions are suggested to aid in developing a better understanding of the knowledge acquisition process. Pursuing these research questions should lead to the identification of the skills and techniques necessary to successfully perform knowledge acquisition. Once these skills have been identified, selection and training programs can be developed to help reduce the shortage of qualified knowledge engineers and, ultimately, facilitate the increased development of expert systems in organizations.

1.0 INTRODUCTION

In recent years there has been a tremendous increase in the development of expert systems within the information systems domain in organizations. This push by organizations to increase their commitment to the development of expert systems is stretching the already limited number of qualified expert system developers. The severe shortage of these developers has been cited as one of the constricting factors to the further development of expert systems (McDermott 1983; Hayes-Roth 1983; Williams 1986). The individuals who develop expert systems are known as knowledge engineers, and their job as knowledge engineering. The knowledge engineer's job has been defined as the process of "mining, molding, assembling, and refining an expert's knowledge" (Hayes-Roth 1983, p. 43).

Dibble and Bostrom (1987) separate the knowledge engineering process into two main functions: knowledge acquisition and expert systems construction. Knowledge acquisition has been defined by Waterman (1986) as "The process of extracting, structuring, and organizing knowledge from several sources, usually human experts, so it can be used in a program" (p. 392). The task of knowledge acquisition has long been cited as the "bottleneck" that slows down the expert systems development process (Feigenbaum and McCorduck 1983; Buchanan et al. 1983). With increased emphasis on expert systems development in organizations, the capability to identify the skills and techniques required to better facilitate knowledge acquisition should lead to an increased understanding of knowledge engineering and hopefully to improved performance and productivity of knowledge engineers. The identification of skills should also provide a starting point in the effort to reduce the critical shortage of qualified knowledge engineers.

This paper summarizes what is known about the skills required and the techniques utilized in the knowledge acquisition process. While there are noted differences between expert systems development and traditional information systems development, there are also similarities. Therefore, to guide this exploration, the existing research pertaining to information requirements determination and that which pertains to the attributes of effective systems analysts will be utilized. Case study reports of actual expert systems development, along with the practitioner literature from this highly applied field, will also be referenced. Finally, given the lack of research performed in the knowledge acquisition area, research questions will be raised in order to stimulate debate and serious inquiry into this increasingly important topic.

2.0 KNOWLEDGE ENGINEERING MODEL

A model of the knowledge engineering process (Figure 1) has been developed (Fellers 1987) that separates the tasks performed by the knowledge engineer into the two phases described earlier: knowledge acquisition and expert systems construction. The knowledge acquisition process involves one or more knowledge engineers interacting with one or more domain experts. (Unless otherwise specified, knowledge engineer will mean one or more knowledge engineers, and expert will mean one or more domain experts.) Each of these participants brings a certain set of attributes to this interaction with the goal of developing a shared representation, or model, of the expert's problem-solving processes.

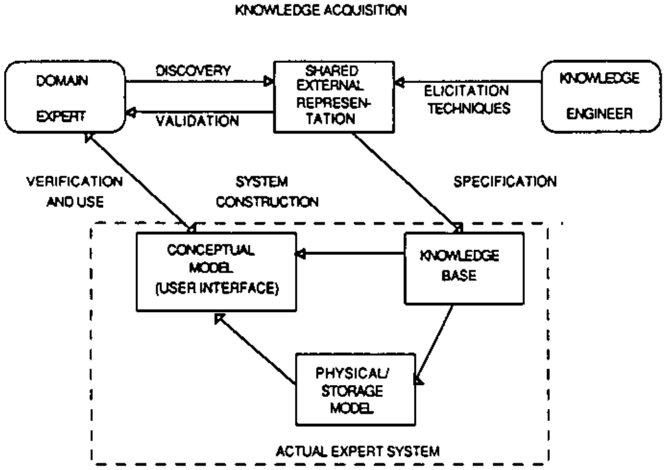

Figure 1. Knowledge Engineering Process

To develop this shared representation, the knowledge engineer must capture information about that process by using elicitation techniques (such as interviewing) along what has been labeled the "discovery" path. As the knowledge engineer develops this representation, the expert must provide a degree of corroboration via the "validation" path. The desired end result of this process is a shared and accurate external representation of the expert's problem-solving process. The knowledge engineer will utilize modeling techniques and representation schemes, such as cognitive maps or rules, to develop this shared external representation.

This external representation may, or may not, be in the same form as the eventual representation chosen to implement the knowledge base. If the knowledge engineer choses to use rules to develop this external representation, these rules may be implemented directly into a knowledge base representation. If some other method is chosen for this external representation, such as cognitive maps, an additional step is needed to implement the cognitive maps into a knowledge base representation.

Once this external representation has been developed, the systems construction phase begins. This phase includes the selection of the appropriate knowledge base representation and involves the mapping of the external representation previously discussed to this knowledge base representation. Once the system is constructed, the expert will utilize the user-system interface to test and use the system along the verification path. The focus of this paper is on the knowledge acquisition phase of the knowledge engineering process. Issues that relate to representation selection problems and expert validation along the user-system interface are presented in Fellers (1987). The next section covers the desired attributes for the knowledge engineer, followed by a discussion of elicitation techniques, the issues pertaining to the development of the shared external representation, and, finally, specific research questions are provided for each of these areas.

3.0 KNOWLEDGE ENGINEER ATTRIBUTES

Literature from three areas will be discussed in this section. The first area is the related research that has been performed regarding the skills, abilities, knowledge, and problem-solving behaviors of systems analysts. The second area is the very limited research that has been performed in the expert systems area. The third area is the practitioner literature from this highly applied field that discusses the preferred attributes of knowledge engineers.

3.1 Research Pertaining to Systems Analysts

A useful beginning in analyzing the desired attributes of a knowledge engineer is to draw upon the research that pertains to systems analysts. There are a number of parallels between the job of systems analyst and that of knowledge engineer. In fact, one practicing knowledge engineer (Rolandi 1986) stated that knowledge engineering is more systems analysis than it is not. He pointed out that the tools and techniques of knowledge engineers and systems analysts are largely the same; the difference lies in the object of analysis. While the systems analyst "studies the flow of information through [usually] a clerical process...the knowledge engineer studies the flow of knowledge through the decision-making process of a human expert" (Rolandi 1986, p. 59). The early work done in the systems analysis area was to try to identify the key skill components necessary to effectively perform the tasks required of a systems analyst. This research revolved mainly around the technical skill requirements such as knowledge of programming languages and techniques, just as the early developers of expert systems were required to have strong programming knowledge and skills to construct expert systems due to the low-level of tools available. Key research pertaining to systems analysts will be presented in order to identify the skills found to be critical for effective systems analyst performance.

Arvey and Hoyle (1974) developed a behaviorally based ratings scale for systems analysts by identifying the major dimensions of their job behaviors. Research by Henry (1974) and Benbasat, Dexter and Mantha (1980) found that technical skills were important, but that there was another set of skills that were also important for systems analysts. These skills were labeled "people-oriented," or generalist, and focused more on such areas as communications, organizational understanding, interpersonal relations, and modeling techniques. Another survey of MIS skills (Cheney and Lyons 1980) practically ignored "people-oriented" skills and again focused more on the technical systems skills required for systems analysts.

Vitalari performed a field experiment using protocol analysis to contrast the abilities of high-rated versus low-rated systems analysts (rated by supervisor). Vitalari and Dickson (1983) reported the first results of this experiment in an examination of the problem-solving behaviors of the systems analysts in two categories: mental behaviors and problem-solving modes. Based upon their exploratory results, they provided a prognosis for effective analyst performance:

- o **Analogical Reasoning:** Effective analysts use information from the environment to classify problems and relate them to previous experiences. If a match is found, the analyst draws upon that previous experience to partially structure the current problem, search for additional information and, in some cases, employ previous solutions.
- o **Planning, Goal Setting, and Strategy Formulations:** Effective analysts should set high level but measurable goals to map out the relevant subproblems and structure the overall task. It is conceivable that analysts deal with a hierarchy of goals that allow them to deal at different levels of detail.
- o **Hypothesis Management:** Effective analysts should develop and manage hypotheses of the problem-solving process to effectively reject low probability hypotheses while retaining and pursuing valid hypotheses.
- o **Operative Knowledge for Application of Heuristic Knowledge:** Effective analysts should apply the use of heuristics to aid in the facilitation of the problem-solving process.
- o **Problem Facilitation:** Effective analysts should understand the importance of the character and quality of the interpersonal relationship between the analyst and the user. In a sense, the entire

problem-solving process depends on the quality of this relationship.

In a further analysis of the experimental data, Vitalari (1985) explored the content of the emerging structure of the systems analyst's knowledge base. He identified six basic areas of knowledge that should exist in an analyst's knowledge base: core systems analysis domain knowledge, high-rated domain knowledge, application domain knowledge, functional domain knowledge, organizational specific knowledge, and knowledge of methods and techniques. Sadek, Hull and Tomeski (1983) presented a case study that supports Vitalari's identification of the importance of domain and organizational knowledge. They discussed the transfer of knowledge/skills between jobs and the importance of understanding such organizational issues as policies, structure, and idiosyncrasies. They pointed out that it takes time to learn these facets of an organization and that systems analysts must take time to adapt to new organizations. A field experiment in the area of software design that was performed by Adelson and Soloway (1985) also supported the importance of domain experience and knowledge from a system design perspective.

Bostrom (1984) was among the first to emphasize the importance that effective patterns of communication play in the development of shared, accurate, and complete system specifications. He explained and advocated the use of the Precision Model, a general communications model, that helps to easily and quickly develop shared "maps" or models between individuals. He also presented the results of an action research project involving the use of the Precision Model to aid in the development of a large system which had been unsuccessfully attempted three times over a 13 year period. By utilizing the Precision Model, the system was successfully implemented as both users and developers had an improved ability to "get the requirements right." Besides being able to reduce the number and length of development meetings, he found that developers were better able to develop and maintain rapport with users, and that team members felt more productive and satisfied when the development meetings concluded. Bostrom indicated that this research was exploratory in nature and lacked the rigor to make any causal statements. However, he believes that it adds support to the importance of communications in the development process and that the model provides guidelines that can easily be followed.

Bostrom's exhortation of the importance of effective communications in the information requirements determination process was supported by a survey of systems analysts and users (Cronan and Means 1984), where both groups stressed the importance of communication skills in the development process. An exploratory study conducted by White and Leifer (1986) also supports the importance of communications skills for systems developers and their importance to the success of a systems development project. This awareness of the importance of effective communication in the systems development process has led to guidelines for effective communication (e.g., Cronan 1984) and frameworks explaining the communications process and the development of research agendas (Guinan and Bostrom 1986).

Guinan (1986) conducted a field experiment in order to rigorously test some of the behaviors outlined in the Precision Model: she utilized content analysis to evaluate the communications behavior of effective systems analysts. As Vitalari had done before, she assigned the systems analysts to two groups, high-rated and low-rated, based on supervisory performance ratings, and then evaluated their performance in terms of outcome measures and communications behaviors. High-rated analysts were found to be significantly better at the following outcome measures: achieving shared meaning, establishing and maintaining rapport, and overall satisfaction with the interaction (users and analysts). High-rated analysts were found to significantly exhibit more of the following effective communications behaviors than low-rated analysts: meta-communication, use of pointers, outcome frames, backtrack frames, and reframing. High-rated analysts were also found to be better (but not significantly so) at the use of metaphors, relevancy challenges, and as-if frames. Overall, good communications skills were found to be extremely important. An overview of these outcome measures and communications behaviors is:

Outcome Measures

- o *Shared Meaning:* a mutual understanding between the two parties engaged in communication.
- o *Rapport:* the harmony, accordance and congruity developed in relationships.

Communications Behaviors

- *Meta-communications:* communicating about communication.
- *Pointers:* verbal cues designed to elicit responses in specific terms.
- *Frames:* conceptual windows which give our models of the world depth and scale.
  - *outcome frame:* a person defines the goals of a meeting or an oral interaction; or to get a person to focus on what they want.
  - *backtrack frame:* a person reviews the progress made in a meeting/interaction.
  - *reframing:* looking at the problem from a different point of view.
  - *as-if frame:* gives the ability to expand the frame beyond the available information; e.g., an individual is encouraged to act as if any needed information existed.
- *Metaphor:* a figure of speech by which a thing is spoken of being like that which it resembles, not fundamentally, but is a marked characteristic.
- *Relevancy Challenges:* when an individual challenges or questions a statement in terms of its relevance to the desired outcome.

The results reported by both Vitalari and Dickson (1983) and Guinan (1986) support the Thomas and Carroll (1981) assertion of the importance of goal setting in the development of systems requirements. Thomas and Carroll also emphasized the importance of such communications behaviors as meta-communications and the use of metaphors in the information requirements determination process. The importance of nonverbal communication has also been stressed in information requirements determination since over 90% of the total message being sent is believed to be nonverbal in nature (e.g., intonations and body language) (Jenkins and Johnson 1977).

3.2 Expert Systems Research

Littman (1986) presented the preliminary results of a study of the knowledge engineering behavior of six individuals with extensive artificial intelligence (AI) experience who were given the task of designing an AI program. While the individuals worked on the program, they were videotaped and encouraged to "talk aloud" during the process and these verbal protocols were recorded. An analysis of the protocols produced seven themes that appear to be central to the behavior of the designers: 1) the importance of the knowledge engineer's goal structure, 2) the importance of world knowledge, 3) the selection of a general representation schema, 4) causal simulation of the domain, 5) the identification of heuristics, 6) model testing and progressive refinement, and 7) focusing on a "touchstone." Many of these behaviors support the results of Vitalari and Guinan that were presented earlier. They emphasize in particular the importance of goal structures and the development of specific goals (outcome frame) to facilitate the problem-solving process.

Littman's preliminary analysis resulted in four major implications: 1) the means by which expert knowledge engineers extract mental models of the domain experts is itself heavily knowledge based; 2) a great deal of the behavior of experienced knowledge engineers appears to be based on the heuristic classification of problem types; 3) the process of empirically studying the methods that knowledge engineers use to extract mental models of domain experts is a potentially useful enterprise; and 4) the videotape protocol methodology appears to provide a potentially useful tool for studying the process of designing knowledge engineered programs.

Bimson and Burris (1987), in a case study discussing the job of a knowledge engineer, described it as a creative business that "is more heuristic than algorithmic, more technique than technology" (p. 460). They stated that individuals must learn the techniques of knowledge engineering by practice. They also believe that the techniques required in knowledge engineering do not possess a "natural set of bounds...we have found knowledge engineering to be an art, a creative business of gathering, analyzing, and organizing knowledge using a variety of techniques which interact in complex ways" (p. 469).

3.3 Expert Systems Practitioner Literature

The first knowledge engineers were often computer scientists who became knowledge engineers due to their strong technical abilities to construct expert systems. While the early expert systems development efforts required the use of low-level development tools, with the advent of high-level development tools, the requirements for strong technical skills have now diminished somewhat (Dibble and Bostrom 1987). At this point the emphasis is starting to shift away from technical skills and move towards the skills required to perform knowledge acquisition. In fact, Williams (1986) believes that experienced programmers are not always the best choice for knowledge engineers. He believes that the emphasis should not be on specialized programming skills, but rather on the ability to think rationally and communicate well. He further stated that the "traditional systems analyst more closely resembles our knowledge-engineering ideal" (p. 67). Sviokla (1986) reiterated this point saying that: "Knowledge Engineering requires a special mix of skills: part consultant, part apprentice, and part programmer" (Part 2, p. 7).

Harmon and King (1985) support many of Vitalari's findings as applied to knowledge engineers. They stressed the importance of the knowledge engineer having a solid understanding of the application or problem domain and the language used by the expert. They also discuss the ability of the knowledge engineer's drawing on past experiences and being able to utilize these experiences in the selection of tools and techniques to be used in the development of the current expert system. Chorafas (1987) also supports the importance of understanding the problem domain and stated that knowledge engineers need to understand the concepts and jargon used in this domain. He believes that the expert will talk more freely with someone who understands what they are saying. Chorafas also subscribed to Vitalari's notion that the knowledge engineer must keep current on new tools, techniques, and methodologies; this is even more critical in the dynamic, rapidly changing area of expert systems tools.

Rolandi (1986) discussed a number of characteristics that he believes are important for a knowledge engineer. He believes that knowledge engineers need more than computer skills. These skills should be complemented by a healthy exposure to the liberal arts so that the knowledge engineers will be broadly educated and well-informed. He also supports Vitalari's notion of the importance of having an understanding of the application domain. Rolandi pointed out that many of the critical skills needed by knowledge engineers are mainly social in nature. To be successful, the knowledge engineer must be generally intelligent, patient, tolerant, amiable, have a good sense of humor, and must be an effective communicator. Diplomacy can be an important asset as the knowledge engineer must be sensitive to the feelings, pride, and prestige of the experts. Finally, he believes knowledge engineers need advanced, socially sophisticated verbal skills.

Many authors in the knowledge engineering area lend support to Bostrom and Guinan's advocacy of the importance of communications skills. While the importance of communications skills in general is often discussed, Finch et al. (1987) provided a specific example of using such meta-communications behaviors as "go ahead" or "any more thoughts" in the facilitation of responses from the expert. They, along with others, also emphasized the extreme importance of the knowledge engineer's ability to establish and maintain rapport with the expert. These illustrations support Guinan's findings of the importance of such communications behaviors. Many of these notions of the importance of the relationship that is developed between the knowledge engineer and expert also support Vitalari's contention of the significance of problem facilitation and the development of a good working relationship by the knowledge engineer. This may be even more important in the expert systems domain due to the high degree of interaction that takes place as the knowledge engineer attempts to elicit the expert's problem-solving processes.

The knowledge acquisition process has traditionally been thought of as an interaction between one knowledge engineer and one expert. However, this is not always the case. For a number of reasons, the use of multiple knowledge engineers is not uncommon. The team approach of using multiple knowledge engineers can change the way the expert system is developed and therefore the skills required by the knowledge engineers on the team. Smith (1984) discussed the experience of using a team approach in expert system development and having different individuals support different tasks. One member of the team interacted with the expert and encoded the domain knowledge. This person did not necessarily have to be able to construct the expert systems, but did have to become familiar with the domain area and be able to use the existing expert system framework to continue the evolution of the system. The other team member needed to have a detailed understanding of the design and implementation process of expert systems development in order to construct the system.

Sviokla summed up the issue of knowledge engineer skills by stating:

> The appropriate background for a knowledge engineer is not standard. Even though expert systems are on the leading edge of computer software, computer science Ph.D.s may not always make the best knowledge engineers. Taylor, an experienced practitioner, suggests that the best knowledge engineers are an eclectic lot who come from varied backgrounds like English, philosophy, and art. [Sviokla 1986, Part 2, p. 9]

Clearly the skill set required for a knowledge engineer has moved beyond just the technical expert systems construction skills. Good communications ability appears to be an essential element of this skill set. A broad background consisting of varied abilities such as patience, good conceptual ability, intelligence, and diplomacy are also considered essential. Other factors of importance include an understanding of the organization in general and the application domain in particular. One final skill that is viewed as critical is experience: many authors indicate the best way to learn is by practice. While much work still needs to be done in the specific identification of the particular skills that are pertinent to knowledge engineers, the research performed by Vitalari and Guinan can be used as a starting point. Once those skills have been identified, the next step can be determining how best to proceed in order to reduce the existing shortage of knowledge engineers. Specific questions to answer include the identification and training of potential knowledge engineers. One possible source of potential knowledge engineers is the current pool of systems analysts: it may be found that these individuals are good candidates to move into knowledge engineering positions.

4.0 ELICITATION TECHNIQUES

Although it is often stated that knowledge acquisition is the "bottleneck" in the expert systems development process, very few authors actually deal with this issue. Most literature in this area is more concerned with how the knowledge is actually implemented in some form of knowledge base representation rather than how it was elicited. Hoffman (1987) takes this problem one step further by stating that: "In short, apparently little or no systematic research has been conducted on the question of how to elicit an expert's knowledge and inference strategies" (p. 54).

The goal of knowledge acquisition is to document the knowledge of the expert in order to build the knowledge base. The most commonly used approach is the technique of interviewing the expert. These extensive interviews often last months and may even take place over several years. During the interviewing process, the knowledge engineer actively questions the expert who is consciously focusing on the knowledge that is being used in the problem-solving process. In an unstructured interview, the expert is often performing a "familiar" task, one that he/she performs on a frequent basis, while the knowledge engineer asks "more-or-less spontaneous questions" (Hoffman, 1987). Chofaras (1987) pointed out that one of the roles of the knowledge engineer is that of devil's advocate: "the knowledge engineer should raise conceptual difficulties and let the expert react" (p. 105). Hoffman (1987) also points out that most knowledge engineers rely exclusively on the unstructured interview method.

The structured interview often takes place after the initial knowledge base has been established and is often used to refine it. One such approach may involve the use of limited-information tasks, which restrict the amount of information available to the expert (Hoffman 1987). The goal is to force the expert into relying on reasoning skills and knowledge. By doing this, it is hoped that additional evidence about how an expert performs a task will be gained, particularly the strategies that the expert uses.

Interviews are useful for providing a wealth of information about a given domain, but additional methods that force the expert to focus more on their actual problem-solving processes are needed in order to develop a more accurate and complete knowledge base. One way to force the expert to do this is to utilize constrained-processing tasks. These tasks attempt to restrict or change the reasoning strategy utilized by the expert. Hoffman (1987) discussed two approaches: the method of simulated familiar tasks and the method of scenarios.

Simulated familiar tasks utilize archival data to perform a familiar task. Prerau (1987) discusses utilization of this method for knowledge acquisition for the COMPASS system (Central Office Maintenance Printout Analysis and Suggestion System). Prerau discussed two types of simulation: hand simulation and computer simulation. With hand simulation, each small step of the expert's reasoning could be examined, while the computer simulation was most useful after a large amount of knowledge had been implemented.

When using scenarios, the expert will often draw on analogies of previous situtations or cases. Scenarios can be thought of as a "what if" type of approach in that they can force an expert to concentrate on a specific task or problem. Bimson and Burris (1987) advocate the use of scenarios stating that they "force the expert to focus on the problem solving and allow the knowledge engineer to infer the knowledge used from the expert's description of the actual event" (p. 462). They found scenarios to be a critical component of their knowledge gathering activity. They also pointed out that scenarios are different from case studies in two ways: 1) typically, case studies are a matter of record and scenarios are not, and 2) scenarios tend to focus more deeply on specific, isolated problems, whereas case studies cover the entire spectrum of problem solving in that domain (p. 464). They recommended analyzing actual projects and provided some guidelines to help in project selection.

Another approach to determine the refined or subtle aspects of the expert's reasoning is the use of tough cases. These cases may also be of a limited information or constrained processing nature. Hoffman (1987) states that "Subtle or refined aspects of an expert's reasoning are often manifested when an expert encounters a tough case, a case with unusual, unfamiliar, or challenging features" (p. 57). Buchanan et al. (1983) supported the notion of having the expert work through some problems and then "go back through each solution in detail to determine the apparent reasoning strategy, the justification for each problem-solving step, and the knowledge brought to bear on the problem" (p. 154). Rolandi (1986) also believes that the emphasis should be on case analysis and that a large number of cases should be used. Prerau (1987) also advocates the use of cases. He recommends using test cases to elicit the initial knowledge and then use a large number of cases in order to expand and modify that knowledge. Smith (1984) suggested that, by working on problems that the expert actually wants to solve, the commitment and interest of the expert will be increased.

In one of the very few research studies in this area, Grover (1983) described an exploratory experiment where four different approaches to interviewing an expert were evaluated: 1) forward scenario simulation (archetype acquisition or "walk throughs"), 2) goal decomposition (20 questions), 3) procedural simulation (protocol analysis), and 4) pure reclassification (frame analysis). Grover reports that the first and fourth techniques turned out to be the most useful.

With most of these methods, the knowledge engineer takes very thorough notes in order to document the knowledge the expert shares. To supplement this note taking process, it is often recommended that they record the information provided by using such techniques as protocol analysis. The expert is asked to "think aloud" as they work on the problems presented to them. The thinking aloud, or verbal protocols, are usually recorded by the knowledge engineer so that an in-depth study of the problem-solving processes of the expert can be performed. Rolandi (1986) advocates the recording of verbal protocols for all cases. While indicating that recording the protocols is important, he believes that audio recordings are enough, since video recording may make the expert nervous. Another important point Rolandi made is that the knowledge engineer should try to replicate the normal decision-making circumstances of the domain expert. Prerau (1987) stated that they initially audio taped their sessions, but after a while the expert would slow down enough that he could explain each step of his analysis. This is one of the keys to the successful transfer of knowledge form the expert to the knowledge engineer: the ability of the expert to explain each step of the problem-solving process and the ability of the knowledge engineer to elicit and capture this information.

Hoffman (1987) performed an evaluation of these techniques based upon his experience in developing expert systems. He concluded that all experts, domains, and expert system development projects differ, and that some methods will work for some projects, while others will not. Prerau (1987) supports this notion and urges knowledge engineers to modify their development strategies to fit the situation and the people involved. Hoffman goes on to criticize authors reporting the results of expert systems development projects stating that they ignore reporting how the knowledge was acquired, but typically "jump right into a discussion of systems architecture" (p. 62). He calls for developers to report their knowledge acquisition methods and the efficiency of the methods.

In addition to the expert in an area, there are other sources of knowledge that can be captured. Prerau (1987) advocates the use of reference books and other documents as the basis for an initial knowledge base. Waterman (1986), in a discussion of the knowledge acquisition process, described other indirect sources of information, such as textbooks, reports, databases, empirical data, and personal experience. He also provided a summary of techniques that can be used for extracting knowledge from the expert that consist of on-site observation, problem discussion, problem description, problem analysis, system refinement, system examination, and system validation (p. 158).

Even though there are a number of elicitation techniques that are commonly accepted and used, efforts are continually being made to incorporate techniques from other disciplines in an attempt to improve this process. Examples of these efforts include Discourse Analysis (Belkin, Brooks and Daniels 1986) and Psychological Scaling (Cooke and McDonald 1986).

There are a number of other issues that arise when considering elicitation techniques. There has been significant recent work in the area of automated knowledge acquisition, but this topic is beyond the scope of this paper. Fellers (1987) provides a discussion of some examples and issues that are pertinent in this area. Another issue that has begun to generate debate is the use of multiple knowledge engineers and/or experts in the development of expert systems.

Rolandi (1986) believes that there should never be more than two knowledge engineers involved in the interviewing process because it can lead to a "chaotic questioning scheme." He has also decided that one knowledge engineer is not optimal either in that one knowledge engineer may get stuck on an invalid line of reasoning, or engage in excessive interpretation of the expert's behavior. There are potential problems of the knowledge engineer's imposing his/her personal interpretation on the expert's problem-solving strategy. Rolandi warns that the knowledge engineer may lose objectivity and try to design the system to prove his/her interpretation of how the expert solves problems. To help alleviate some of these problems, Rolandi recommends using a two-person interviewing team. In this approach, one knowledge engineer acts as the interviewer and the other acts in a quality control capacity. This second knowledge engineer is to control the quality of the communication that goes on between the first knowledge engineer and the expert, as well as "to head off or diffuse miscommunications and other unavoidable frustrations that prolonged discourse which might otherwise lead to problems between the knowledge engineer and the expert" (p. 60). Again, many of these issues reiterate the importance of the effective communications skills that knowledge engineers must have in order to be successful in the elicitation process.

Another issue is that of the elicitation process with multiple experts. Many authors, such as Mittal and Dym (1985), advocated the use of multiple experts in this process in order to better understand the kinds of expertise utilized in the domain. Smith (1984) worked with only one expert on a project, but plans, in the future, to work with multiple experts with differing backgrounds in order to get multiple points of view. McDermott (1983) described using seven experts on one project, since no one expert possessed all the knowledge necessary to develop the system. Prerau (1985) and Sviokla (1986) both recommended that, when dealing with more than one expert, there should be a designated chief or primary expert.

Eliciting knowledge from multiple experts opens up the issue of interpretation of the alternative points of view and methods of problem solving. Chorafas (1987) recommended group discussions as one possibility. However, he believes that just recording the group discussions is insufficient for gaining an understanding of the contents and reasoning in the area. He provided a few novel suggestions for ways of improving elicitation from groups: one is to use teleconferencing, which could lead to facilitating distributed types of meetings; the other is holding a "knowledge competition." Although controversial, he believes that this may work in some situations. Elicitation of knowledge from multiple experts has been, and will continue to be, a requirement of many systems.

There are a number of different elicitation techniques that have been described and/or prescribed for knowledge acquisition. While many of these techniques appear to be useful, it is not known precisely which technique, or combinations of techniques, are the most effective, and under what circumstances. It is not known which techniques are best at eliciting the different kinds of knowledge needed for the development of the knowledge base. It is not known which skills are required of the knowledge engineer in order to utilize each technique. There are also questions to answer when multiple knowledge engineers or experts are used as to which techniques are most effective and how and when they should be utilized. Another issue is dealing with the multiple viewpoints and integrating this knowledge from several different sources.

5.0 EXTERNAL REPRESENTATIONS

The goal of the interaction between the expert and the knowledge engineer is the development of a shared representation, or model, of the expert's problem-solving process. The different elicitation techniques previously described are utilized in order to facilitate the development of this representation. The end result of this process is the development of an external representation that depicts this shared model and can then be used to develop the actual knowledge base.

5.1 Representation Research Model

Juhn and Naumann (1985) have developed a representation research model that describes the role of representation in the systems analysis and specification process. They view the purpose of the representation as being a communication media between users and analysts. They identify four paths in the model: 1) discovery, 2) validation, 3) specification, and 4) verification. Since these paths are applicable to knowledge engineering, they have been incorporated into the knowledge engineering process model shown in Figure 1. The discovery path entails the elicitation of domain knowledge from the expert in order to build the shared representation. The validation path is where the expert corroborates the representation that has been developed by the knowledge engineer. The specification refers to the (external) representation that is ready to be used by the system builders in construction of the system. The verification path allows for the constructed system to be used and tested by the expert to see whether it meets his/her standards.

Two lab experiments have been performed utilizing the representation research model. Juhn and Naumann (1985) evaluated the effectiveness of data representation characteristics on user validation. They compared two semantic (or graphical) representations with two relation-based models. The results showed that semantic (or graphical) representation provided a more effective communication of the data model in terms of user comprehension than the relation-based model. Larsen and Naumann (1986) performed an experiment to determine whether an abstract or concrete model is best for the process of requirements discovery. Their results indicated that the concrete model was superior to the abstract model in the discovery process. Juhn and Naumann (1985) point out that the uses of these representations are conflicting: what may be a good representation for one path may not be for another.

Although the previous studies cited pertain to the database area, this line of research is very important to the development of external representations for expert systems. The model developed by Juhn and Naumann is applicable to the expert systems field and provides yet another rich area where research is needed. In order to better facilitate the development of external representations there are a number of other types of representations that are now being used in the expert systems domain. Some have been used successfully in traditional information requirements determination, some have their roots in other areas. Examples of the use and development of alternative external representations will be discussed.

5.2 Developing External Representations

Among the different external representations that can be used by knowledge engineers to assist in the expert system development process, some are procedural, such as production rules, while others are declarative in nature, such as semantic networks (S-nets). These two representations are commonly used due to their ultimate role as actual knowledge base representation mechanisms. This is also a reason why they often make poor choices for external representations. In order to facilitate the validation process, the expert must be able to understand the representation. When rules or S-nets are used, this may not be the case. Prerau (1987) recommends using some form of quasi-English if-then rules to facilitate this understanding. He stated that "An expert should be able to understand this method of knowledge representation more easily than other AI paradigms and after some exposure might be able to relate knowledge to the knowledge engineers by utilizing this paradigm" (p. 47).

Bimson and Burris (1987) described the process they utilized to gather knowledge and data in the development of a Software Project Management System. The route they chose to take down the discovery path involved the questioning of multiple experts about: 1) management practices, 2) management concepts and relations, and 3) actual scenarios. During the process of knowledge acquisition, they began to determine what type of representation would best suit this problem domain. Their search led them to semantic inheritance networks (SI-nets) to model the declarative component of the knowledge base. Standard production rules were chosen to provide a deductive component to enable reasoning over the declarative component. What was missing in their description of this process was any form of validation by the expert. There was no indication of any expert involvement once the interviews were concluded; i.e., no measure of validation on the representation being developed by the knowledge engineers.

In another description of an expert systems development process, Finch et al. (1987) utilized cognitive mapping and information display boards to elicit and structure the expert's knowledge. Cognitive mapping, which has been successfully used to aid in traditional information requirements determination (Montazemi and Conrath 1986), was used to discover the domain attributes and relationships, while information display boards were used to examine the expert's decision processes. They described a four-interview process in which the expert was heavily involved as the knowledge engineers elicited and developed the cognitive maps and information display boards. This example highlights the roles played by the expert and knowledge engineer, along with the importance of both the discovery and validation paths, in the development of the external representation.

Boose and Bradshaw (1987) discussed the use of NeoETS (an extension of the expertise transfer system) to elicit problem-solving knowledge and store it in ratings grids. The columns in the grids are problem solutions (elements) and the rows represent solution traits (constructs). Traits are determined when the expert is asked to discriminate among a group of elements and provide a rating to indicate where the elements fall on the rating scale. Ideas from George Kelly's Personal Construct Psychology were utilized in the interviewing techniques used to construct the grid. Similar techniques have been used in the information requirements determination area (Grudnitski 1984). NeoETS has the capability to assist the expert in analyzing the grid ratings to help in the refinement of its problem-solving capability. This system has the capability to cover both the discovery and validation paths in its operations.

Pracht (1987) proposed a visual modeling system that: 1) is based on visual modeling techniques to provide an easy and natural interface for managers, 2) supports the creative tasks of modeling the structure of complex problems, 3) provides the capability for modeling the behavior of a system, and 4) serves as the basis for incorporating heuristic rules appropriate for automated inference procedures. Pracht discussed the importance of images and the role they play in creative problem solving and stated that

> Images represent the user's mental model of a problem domain. If the visual constructs of these images are given precise meaning, the user's knowledge can be captured and stored in a knowledge base. If the user is provided with a way to work with a visual representation of the model or knowledge, then the power of the computer can be applied to knowledge structuring and acquisition in a manner that more closely matches the natural thought processes. [Pracht 1987, p. 480]

Such a system would also provide facilities for discovery and validation.

The use of such visual or graphical aids to assist in the discovery, understanding, and modeling of an expert's knowledge has long been utilized in the operations research (OR) area (O'Keefe 1985). O'Keefe pointed out that such graphical techniques as activity-cycle diagrams and decision trees have been used to "capture the essential features of a system and portray this to a manager" (p. 127). He stated that visual simulation can "provide a playback of the model representing the expert's conceptualized world" (p. 128). It is believed that knowledge engineers could benefit from the techniques utilized and the experience gained by those in OR for the development of representations. Elam and Henderson (1983), in a discussion of decision support system design, also pointed out the attractiveness of graphical representations given that: 1) the expressive power of graphs is sufficient to encode any fact or concept that is encodable in any other representation and 2) graphical structures themselves serve as a guide for knowledge retrieval and processing.

The examples covered here demonstrate some of the variety of external representations and development strategies that exist. In some approaches, the knowledge engineer will elicit requirements and then develop the external representation without additional expert intervention, while others utilize repeated involvement on the part of the expert. Some even propose automation to facilitate this process. The support for graphical representations provides some encouragement for Juhn and Naumann's (1985) contention of the advantage of semantic (or graphical) representations for the validation process. Even with this support there are still many unanswered questions. One such question is whether there is one best type of representation to use, and, if so, what is it? Another pertinent issue is which elicitation techniques are required to develop a particular representation? Other issues revolve around the research model pertaining to the determination of the appropriate representation for each path if, as Juhn and Naumann suspect, the differing purposes of the paths will require different representations.

6.0 RESEARCH QUESTIONS

Some specific research questions are put forth for each of the areas covered in the paper. For research questions governing the overall management of the expert systems development process and how these areas are linked together, see Dibble and Bostrom (1987). The questions are organized by topic and include section headings and numbers for reference to the body of the paper.

Participant Attribites (3.0)

1. What is the desired skill set for a knowledge engineer (problem-solving behaviors, communications behaviors, modeling skills, and technical skills)?

2. Once the skill set has been identified, can it be used to identify potential candidates for the job of knowledge engineer? How can it be used to help train individuals for the job of knowledge engineer?

3. Upon identification of this skill set, are existing systems analysts good candidates for the job of knowledge engineer? Could current systems analysts be trained to move into knowledge engineering positions? If systems analysts are not good candidates, what other individuals would make viable knowledge engineering candidates?

4. What effect would using two-person knowledge engineer teams have on this skill set? Does dividing the knowledge engineering tasks between two people change the requirements of the job? Is the knowledge engineer team approach a short-term solution to the bottleneck or does it have potential in the long run?

Elicitation Techniques (4.0)

1. Is there one best elicitation technique for knowledge acquisition? If not, what is the best combination of techniques? Which techniques are best used under which circumstances? What skills are required in order to utilize each of the techniques?

2. What is the effect of using two knowledge engineers on the elicitation process? Is the "quality control" approach as described by Rolandi (1986) superior to one knowledge engineer, or multiple knowledge engineers? If so, under what circumstances?

3. What are the best methods for eliciting knowledge from multiple experts? How can their responses best be integrated? What is the best manner for handling conflict and disagreement among multiple experts?

External Representations (5.0)

1. Is there one best form of external representation? A combination of representations? If not, which representations work best and under which circumstances?

2. Which elicitation techniques are required in order to develop each type of representation?

3. Does one representation adequately serve the purposes of both the discovery and validation paths? If not, which representation is optimal for each path?

7.0 CONCLUDING REMARKS

The knowledge acquisition process still appears to be the least understood and the most critical in expert systems development. A number of issues pertaining to the appropriate skills and techniques in the knowledge acquisition area have been addressed. By identifying these issues and drawing upon the related MIS literature, some common ground has been established. In addition to the discussion of these issues, research questions have been outlined to suggest ways of improving our understanding in each of these areas. A shortage of qualified knowledge engineers, coupled with the increased demand for expert systems development and the very difficult process of knowledge acquisition, makes this area critical for further study.

In order for the types of questions put forth here to be answered, a coordinated research effort must take place. Since the research topics in the area of expert systems cross the boundaries of many disciplines, yet fit entirely into none, coordinated research efforts among researchers working in such areas as AI, computer science, OR, cognitive psychology, and MIS must take place. It is through such cooperative and cummulative efforts that real gains can be made.

Sviokla (1986) quoted Clippenger as saying that "Expert system development is like walking around in a dark room; you don't know where the walls are until you hit them" (Part 2, p. 8). It is hoped that this paper has shed a little light on some of the key issues affecting knowledge acquisition and the development of expert systems. Hopefully the issues outlined here and the research questions provided will stimulate serious inquiry into how these factors can be addressed in order to improve the development of expert systems within organizations.

REFERENCES

Adelson, B., and Soloway, E. "The Role of Domain Expert Software Design." *IEEE Transactions on Software Engineering,* SE-11:11, November 1985, pp. 1351-1360.

Arvey, R. D., and Hoyle, J. C. "A Guttman Approach to the Development of Behaviorally Based Rating Scales for Systems Analysts and Programmer/Analysts." *Journal of Applied Psychology,* 59:1, 1974, pp. 61-68.

Belkin, N. J.; Brooks, H. M.; and Daniels, P. J. "Knowledge Elicitation and Discourse Analysis." *Knowledge Acquisition for Knowledge-Based Systems Workshop,* Banff, Alberta, Canada, November 2-7, 1986, pp. 3-0 - 3-16.

Benbasat, I.; Dexter, A. S.; and Mantha, R. W. "Impact of Oganizational Maturity on Information System Skill Needs." *MIS Quarterly,* 4:1, March 1980, pp. 21-34.

Bimson, K. D., and Burris, L. B. "The Craft of Engineering Knowledge: Some Practical Insights." *Twentieth Annual Hawaii International Conference on System Sciences,* 1987, pp. 460-469.

Boose, J. H., and Bradshaw, J. M. "A Knowledge Acquisition Workbench for Eliciting Decision Knowledge." *Twentieth Annual Hawaii International Conference of System Sciences,* 1987, pp. 450-459.

Bostrom, R. P. "Development of Computer-Based Information Systems: A Communication Perspective." *Computer Personnel,* 9:4, August 1984.

Buchanan, B. G.; Barstow, D.; Betchel, R.; Bennet, J.; Clancy, W.; Kulikowski, C.; Mitchell, T.; and Waterman, D. "Constructing an Expert System." In F. Hayes-Roth, D. Waterman, and D. B. Lenat (eds.), *Building Expert Systems,* Addison-Wesley Publishing Company, Reading, Massachusetts, 1983, pp. 127-168.

Cheney, P. H., and Lyons, N. R. "Information Systems Skill Requirements: A Survey." *MIS Quarterly,* 4:1, March 1980, pp. 35-43.

Chorafas, D. N. *Applying Expert Systems in Business.* McGraw-Hill Book Company, New York, 1987.

Cooke, N. M., and McDonald, J. E. "The Application of Psychological Scaling Techniques to Knowledge Elicitation for Knowledge-Based Systems." *Knowledge Acquisition for Knowledge-Based Systems Workshop,* Banff, Alberta, Canada, November 2-7, 1986, pp. 10-0 - 10-16.

Cronan, T. P. "Application Systems Development: A Communication Model for Business Users and DP Personnel." *Data Base,* 16:1, Fall 1984, pp. 21-26.

Cronan, T. P., and Means, T. L. "System Development: An Empirical Study of User Communications." *Data Base,* 15:3, Spring 1984, pp. 25-33.

Dibble, D., and Bostrom, R. P. "Managing Expert Systems Projects: Factors Critical for Successful Implementation." *ACM SIGBDP-SIGCPR Conference,* Coral Gables, FL, March 3-4, 1987, pp. 96-128.

Elam, J. J., and Henderson, J. C. "Knowledge Engineering Concepts for Decision Support System Design and Implementation." *Information and Management,* 6:2, 1983, 109-114.

Feigenbaum, E. A., and McCorduck, P. *The Fifth Generation.* Addison-Wesley Publishing Company, Reading, Massachusetts, 1983.

Fellers, J. W. "Key Factors in Knowledge Acquisition." *Computer Personnel,* 11:1, May 1987, pp. 10-24.

Finch, L.; Landry, J.; Monarchi, D.; and Tegarden, D. "A Knowledge Acquisition Methodology Using Cognitive Mapping and Information Display Boards." *Twentieth Annual Hawaii International Conference on Systems Sciences,* 1987, pp. 470-477.

Grover, M. D. "A Pragmatic Knowledge Acquisition Methodolgy." *International Joint Conference on AI,* Karlsruhe, West Germany, 1983, pp. 436-438.

Grudnitski, G. "Eliciting Decision-Makers' Information Requirements: Application of the Rep Test Methodology." *Journal of Management Information Systems,* 1:1, Summer 1984, pp. 11-32.

Guinan, P. J. "Specialist-Generalist Communication Competence: An Investigation of the Communications Behavior of Information Systems Developers." Unpublished Doctoral Dissertation, Indiana University, 1986.

Guinan, P. J., and Bostrom, R. P. "Development of Computer Based Information Systems: A Communications Framework." *Data Base,* 17:3, Spring 1986, pp. 3-16.

Harmon, P., and King, D. *Expert Systems.* John Wiley & Sons, Inc., New York, 1985.

Hayes-Roth, F. "The Industrialization of Knowledge Engineering." *Symposium on Artificial Intelligence Applications in Business,* New York University, New York, May 1983.

Henry, R. M. "Skills Possessed and Skills Useful for MIS Practitioners -- A Research Report." *National Computer Conference,* 1974, pp. 889-895.

Hoffman, R. H. "The Problems of Extracting the Knowledge of Experts from the Perspective of Experimental Psychology." *AI Magazine,* 8:2, Summer 1987, pp. 53-67.

Jenkins, A. M., and Johnson, R. D. "What the Information Analyst Should Know About Body Language." *MIS Quarterly,* 1:3, September 1977, pp. 33-47.

Juhn, S. H., and Naumann, J. D. "The Effectiveness of Data Representation Characteristics on User Validation." In *Sixth International Conference on Information Systems,* Indianapolis, Indiana, December 16-18, 1985, pp. 212-226.

Larsen, T. J., and Naumann, J. D. "Systems Analysis Representation for Discovery: Abstract Versus Concrete." Working Paper, Department of Management Science, University of Minnesota, May 26, 1986.

Littman, D. C. "Modelling Human Expertise in Knowledge Engineering: Some Preliminary Results." *Knowledge Acquisition for Knowledge-Based Systems Workshop,* Banff, Alberta, Canada, November 2-7, 1986, pp. 28-0 -28-13.

McDermott, J. "Building Expert Systems." *Symposium on Artificial Intelligence Applications in Business,* New York University, New York, May 1983.

Mittal, S., and Dym, C. L. "Knowledge Acquisition from Multiple Experts." *The AI Magazine,* 6:2, Summer 1985, pp. 32-36.

Montazemi, A. R., and Conrath, D. W. "The Use of Cognitive Mapping for Information Requirements Analysis." *MIS Quarterly,* 10:1, March 1986, pp. 45-56.

O'Keefe, R. M. "Expert Systems and Operational Research -- Mutual Benefits." *Journal of the Operational Research Society,* 36:2, 1985, pp. 125-129.

Pracht, W. E. "A Visual Modeling Approach for DSS Knowledge Acquisition and Organization."

Twentieth Annual Hawaii International Conference on System Sciences, 1987, pp. 478-484.

Prerau, D. S. "Selection of an Appropriate Domain for an Expert System." *The AI Magazine*, 6:2, Summer 1985, pp. 26-30.

Prerau, D. S. "Knowledge Acquisition in the Development of a Large Expert System." *AI Magazine*, 8:3, Summer 1987, pp. 43-51.

Rolandi, W. G. "Knowledge Engineering in Practice." *AI Expert*, 1:4, December 1986, pp. 58-62.

Sadek, K. E.; Hull, R. W.; and Tomeski, A. E. "Information Systems Professional's Job Transitions: Its Influence on Information Systems Design." *Journal of Systems Management*, 34:8, August 1983, pp. 21-29.

Smith, R. G. "On the Development of Commercial Expert Systems." *The AI Magazine*, 5:3, Fall 1984, pp. 61-73.

Sviokla, J. J. "Business Implications of Knowledge-Based Systems, Part I." *Data Base*, 17:4, Summer 1986, pp. 5-19.

Sviokla, J. J. "Business Implications of Knowledge-Based Systems, Part II." *Data Base*, 18:1, Fall 1986, pp. 5-16.

Thomas, J. C., and Carroll, J. M. "Human Factors in Communication." *IBM Systems Journal*, 20:2, 1983, pp. 237-263.

Vitalari, N. P. "Knowledge as a Basis for Expertise in Systems Analysis: An Empirical Study." *MIS Quarterly*, 9:3, September 1985, pp. 221-241.

Vitalari, N. P., and Dickson, G. W. "Problem Solving for Effective Systems Analysis: An Experimental Exploration." *Communications of the ACM*, 26:11, November 1983, pp. 948-956.

Waterman, D. A. *A Guide to Expert Systems*. Addison-Wesley Publishing Company, Reading, Massachusetts, 1986.

White, K. B., and Leifer, R. "Information Systems Development Success: Perspectives from Project Team Participants." *MIS Quarterly*, 10:3, September 1986, pp. 214-223.

Williams, C. "Expert Systems, Knowledge Engineering, and AI Tools -- An Overview." *IEEE Expert*, 1:4, Winter 1986, pp. 66-70.