

Knowledge Acquisition Using Multiple Domain Experts
in the Design and Development of an
Expert System for Disaster Recovery Planning

By

Frank W. Nasuti

Formal proposal for dissertation submitted in partial fulfillment of the requirements for the
degree of Doctor of Philosophy

School of Computer and Information Sciences
Nova Southeastern University

2000

An Abstract of a Dissertation Submitted to Nova Southeastern University in Partial
Fulfillment of the Requirements for the Degree of Doctor of Philosophy

Knowledge Acquisition Using Multiple Domain Experts
in the Design and Development of an
Expert System for Disaster Recovery Planning

By
Frank W. Nasuti
2000

The increasing dependence of organizations on data processing to perform the basic functions of corporate America, together with recent disasters such as earthquakes, tornadoes and hurricanes have awakened management to the realization that they require Disaster Recovery Plans (DRP) and Business Resumption Services (BRS). To address these needs, organizations frequently consult with outsiders to help them develop disaster recovery and business resumption plans. Although consultants and vendors specializing in disaster recovery planning are available, their number is limited and the quality of their services may be questionable. In addition, the information gathering process by consultants is a time consuming process and in most cases requires the use of multiple vendor experts, as well as various resources within the customer's organization. This research proposes, as a solution to address these deficiencies, the design and development of an expert system to assist in the determination of the needs of an organization for disaster recovery and business interruption services, as well as the evaluation of existing plans. This research will design an expert system to develop a disaster recovery plan. It will include the knowledge acquisition processes necessary to elicit information from multiple domain experts. The specific goals of this research are: (1) knowledge acquisition specific to the problems of using multiple domain experts; (2) design and development of a prototype expert system for disaster recovery planning; and (3) validation of the prototype expert system.

Table of Contents

Abstract	ii
List of Tables	iv
List of Figures	v

Chapters

1. Introduction	1
Statement of the Problem to Be Investigated and Goal to be Achieved	1
Relevance and Significance	3
Barriers and Issues	5
Elements, Theories, and Research Questions to Be Investigated	8
Limitations and Delimitations of the Study	9
Definition of Terms	10
Summary	12
2. Review of the Literature	15
Historical Overview of the Theory and Research Literature Specific to the Topic	15
Summary of What Is Known and Unknown About the Topic	36
The Contribution this Study Will Make to the Field	37
3. Methodology	38
Introduction	38
Expert System Project Design	40
Application Description & System Objectives	40
System Justification	41
User Requirements	42
Inputs/Outputs	43
Cost/Benefits	45
Design Specifications	48
Design Methodology	49
Approach	50
Milestones	55
Resources	56
Summary	57
4. Discussion of Expectations	58
Anticipated Benefits & Projected Outcomes	58
Practical Applications of the Findings	58
Constraints and Limitations of the Study	59
Recommendations for Additional Studies	59
Contributions to the Field of Study and Advancement of Knowledge	60

Annotated Bibliography	66
-------------------------------	-----------

Reference List	89
-----------------------	-----------

List of Tables

Tables

1. The Knowledge Acquisition Process 61
2. Validation & Verification of Knowledge Acquisition & Expert Systems 62
3. Milestones 63

List of Figures

Figures

1. Four Phase Approach to Expert System Design & Development 64
2. Flowchart of Routines 65

Chapter 1

Introduction

Statement of Problem to be Investigated and Goal to be Achieved

Disaster Recovery Planning (DRP) and Business Resumption Services (BRS) are both time consuming and costly processes. This, as well as the feeling that “we have never had a problem,” has made companies hesitate to implement plans to resume business activities to prepare for a catastrophic event. However, several factors have caused companies to change their attitude towards their need for DRP and BRS. Natural disasters such as earthquakes, tornadoes and hurricanes have awakened management to the realization that it could happen to them (Cerullo & McDuffie, 1994; Hiles, 1992; Garcia-Molina & Polyzois, 1990). The increasing dependence of organizations on data processing, the fear of lawsuits by shareholders, recent accounting regulations (AICPA, 1988), and Federal regulations that call for the implementation of DRP and BRS plans (Garcia-Molina & Polyzois, 1990) have caused management to address the need to implement DRP and BRS plans.

Business opportunities in both the consulting and services areas for DRP and BRS have increased tremendously (Jacobs & Weiner, 1997; Rudolph, 1990). However, various factors may hinder the ability to provide the services needed by companies.

Three of these are:

1. The lack of experienced consultants to evaluate organizational needs for DRP and BRS, or to evaluate and review existing plans;

2. The need for multiple skills ranging from auditing, assessment, and communications; and
3. The need for expertise in specific areas including hardware, software, telecommunications, general business and industry specific concerns.

Although consultants and vendors specializing in disaster recovery planning are available, their number is limited (Jacobs & Weiner, 1997). The information gathering process between the consultant and the client is a time consuming process (Andrews, 1994) and in most cases requires the use of multiple vendor experts, as well as various resources within the customer's organization (Money & Harrald, 1995). Money and Harrald noted that knowledge acquisition for disaster recovery planning is complicated because of the knowledge must be elicited from multiple sources.

This research proposes, as a solution to address these deficiencies, the design and development of an expert system to assist in the determination of the needs of an organization for disaster recovery and business interruption services, as well as the evaluation of existing plans. This research will attempt to design the expert system described above. The development methodology will include the knowledge acquisition processes necessary to elicit information from multiple domain experts. The specific goals of this research are:

1. Knowledge acquisition specific to the problems of using multiple domain experts;
2. Design and development of a prototype expert system for disaster recovery planning; and
3. Validation of the prototype expert system.

Relevance and Significance

Disaster Recovery Planning

Disaster Recovery Planning and Business Resumption Services are critical in all organizations. However, the ability to resume the minimum functions and operations necessary to ensure continuing operations after a disaster is questionable in all but the smallest entities. Research has shown that over 70% of organizations that experience a catastrophic event close within two years (Andrews, 1994).

Recent events have demonstrated the need for disaster recovery capabilities by all organizations relying on data processing capabilities. Disaster, as it relates to data processing, is defined as an event which causes the loss of computing service, or of a significant part of it, or of systems, communications or applications, for a length of time which prevents the impacted organization from achieving its mission or which imperils the business (Hiles, 1992). Disasters such as Hurricane Andrew, the World Trade Center bombing, the Northridge Earthquake, and Hurricane Hugo (Cerullo & McDuffie, 1994; Griswold, Lightle, & Lovelady, 1990) left many companies in California and Florida helpless. More recent disasters, such as Hurricane Floyd and the El Nino weather incidents, have made many companies realize that both DRP and BRS are necessary. This concern is critical in organizations where sizeable computer processing is present, whether as a service or support function. Added to this are various federal regulations that require key groups such as the banking industry to implement and test a disaster recovery plan (Garcia-Molina & Polyzois, 1990). The result of this is the need for the implementation of disaster recovery planning for data processing operations, as well as the extension of this plan to other critical business functions and operations (Cerullo &

McDuffie, 1994). Evidence strongly supports the concept that companies can survive a disaster if they: (1) plan for the possibility of a disaster, (2) formulate strategies for recovering critical business functions, (3) implement technology to aid the recovery of automated functions and systems, and (4) train employees to implement the strategies (Rudolph, 1990).

Knowledge Acquisition from Multiple Domain Experts

Turban and Aronson (1998) define expertise as the set of capabilities that underlines the performance of human experts, including extensive domain knowledge, heuristics rules that simplify and improve approaches to problem-solving, metaknowledge and metacognition, and compiled forms of behavior that afford great economy in skilled performance. Expertise can be expressed in textbooks, case studies, and documentation, but typically is possessed by what is termed “domain experts.” In many cases, the expertise for a domain is resident in a single expert; however, many expert systems are complicated and the expertise needed must be elicited from multiple domain experts. The elicitation of knowledge from multiple experts brings to surface many complications that do not surface in single domain expert systems. Knowledge acquisition from multiple domain experts is compounded by problems that arise from the obstacles and risks of trying to coordinate human interactions and integrate multiple knowledge (LaSalle & Medsker, 1991).

The key to overcoming these obstacles is a strong knowledge engineer, who has skills in knowledge acquisition from multiple domain experts, who is familiar with knowledge acquisition tools and techniques, and is knowledgeable about the specific domain itself.

The knowledge engineer should also be able to effectively elicit the needed knowledge from the domain experts, while resolving all conflicts in the process.

Barriers and Issues

Specific barriers and issues need to be addressed in both the disaster recovery/business resumption areas and the expert system areas. The latter needs to be broken down to address the barriers and issues specific to knowledge acquisition, knowledge acquisition using multiple experts, the design and development of the expert system and finally the validation of the expert system, which will allow for the validation of the knowledge acquisition process.

The barriers and issues relating to disaster recovery/business resumption areas are based on people and logistics issues. DRP & BRS involve multiple personnel, both at the consultation and customer/client level, each possessing specific knowledge and expertise. DRP involves recovery sites, hardware, software, network and telecommunications, software and data backup, maintenance, testing and other areas relating directly to the data center and computer operations. BRS includes areas not specific to data processing, but more to business continuity, such as voice communications, customers, office space, supplies, and support functions.

Barriers & Issues Specific to DRP/BRS

Contingency planning service is the ability to provide a multi consultant approach to the customer's needs for various disaster recovery planning and business resumption services in the event of a disaster. This concept requires the consulting services of

multiple domain experts, including at a minimum, a DRP/BRS consultant, a telecommunications consultant, a recovery site expert, and an industry expert. This multi consultant/expert approach presents many barriers. These include:

1. The need for each consultant/expert to meet with each customer, to evaluate the customer's environment and needs;
2. Consistency of the consultant/expert at each meeting;
3. The availability of all consultants/experts at the same time or when needed;
4. The need to meet with several different customer personnel; and
5. The need for speedy turnaround time from meetings/consultations to providing the written proposal to the customer.

Many consulting jobs require one or more consultants to meet with the client during the information gathering stage, where the objective is to learn as much as possible about the customer so as to be able to submit a proposal. This is complicated for DRP/BRS services because a single person may not know the detailed information needed by the consultants. Typically, the consultants must meet with anywhere from five to ten key employees to gather all of the specific information they need to understand the client's data processing and business environments.

Additional problems and issues in developing a disaster recovery plan relate to the difficulty of gathering the essential information (knowledge acquisition or elicitation) to be detailed in the body of the disaster recovery plan (Jacobs & Weiner, 1997). This is complicated by the lack of identified or available experts (which is viewed in the industry as a critical lack of qualified personnel), the available expert/consultant being viewed as

costly, or the expert/consultant not being readily available at a particular location, when needed.

Other barriers are the problems inherent in knowledge acquisition, and the complications of this difficult process when using multiple domain experts.

Barriers & Issues Specific to Knowledge Acquisition & Multiple Domain Experts

Key issues in the development of any expert system relate to domain experts. First is the identification of each domain expert. The knowledge engineer must ensure that each expert is truly an expert in his/her domain. If the expert selected is not a true expert in the domain or does not possess the right expertise needed, the knowledge elicited may be incomplete or incorrect, which may leave the expert system susceptible to failure. Second is the availability of the expert during not only the knowledge acquisition process, but throughout the expert system design process, particularly the verification and validation phase. Third is the willingness of the domain expert to truthfully and fully contribute their knowledge and expertise. Although individuals may agree to serve as experts, some may hold back information, either consciously or unconsciously. Various reasons for such behavior include the fear of being replaced, mistrust, or other psychological reasons. Fourth is the threat of the domain expert dropping out. This is catastrophic to the process, particularly at the knowledge acquisition and verification/validation phases. Experts may drop out for various reasons, including conflict with their work, time constraints, or loss of interest. Replacing an expert during the process will not only cause delays, but may result in inconsistencies, redundancies, and incompleteness. Fifth is the theory of paradox of expertise. Liebowitz (1993)

suggests that a major reason why knowledge acquisition is a difficult process relates to the Paradox of Expertise which states that the more competent domain experts become, the less able they are to describe the knowledge they use to solve problems. This is all further complicated when multiple experts are needed

Barriers and Issues Specific to Verification and Validation

The major outcome of this research will be the development of a functional expert system. The beta testing of the expert system will allow for the verification and validation of the system, as well as a validation of the knowledge acquisition process. Issues addressed will include ease of use (user-friendliness), the completeness of the system (its ability to elicit all of the information needed during the consultation process to ensure a quality proposal by the consultants), consistency, and elimination of redundancies. Barriers specific to verification and validation are those commonly found in most beta testing environments. Several of those identified are (1) an adequate number of beta test sites and the ability to expand the number, if necessary, (2) whether the beta sites are adequate representations of the users of the expert system, (3) the completeness or thoroughness of the testing, and (4) the consistency of the testing methodology at all beta sites.

Elements, Theories, and Research Questions To Be Investigated

There are several major elements of this research. The first element is the need for disaster recovery/business resumption plans and the need for an expert system to assist consultants in developing these services for customers. The second element is the

knowledge acquisition process, during which the knowledge or expertise needed by the expert system is elicited and which is complicated by the need for multiple domain experts. The third element is the design and development of the expert system. The fourth element is the validation of the expert system through beta testing, which will further allow for the validation of the knowledge acquisition process, the fifth element of this research project.

Several research questions will also be addressed. One question will be whether the selection of a knowledge engineer knowledgeable in the domain areas will overcome many of the obstacles faced by knowledge engineers lacking such expertise. A second research question will be to identify knowledge engineering techniques that are successful in overcoming the traditional obstacles encountered by knowledge engineers in the past. A third research question is whether an expert system for disaster recovery planning could be developed that is user-friendly and can be utilized by the user as a training tool for DRP. A fourth question is whether the data gathered from the use of the expert system, commonly known as the consultation, will be complete and accurate to allow the consultant to submit a proposal for services to the customer in a timely manner. This will be resolved through the beta testing of the expert system to validate the expert system.

Limitation and Delimitations of the Study

The major outcome of this research will be the development of a functional expert system. The beta testing of the expert system will allow for the verification and validation of the system, as well as of the knowledge acquisition process. Failure to

develop the functional expert system will hinder the researcher's ability to test many of the research issues discussed in the next section and elsewhere in this paper.

In addition, the expert system will utilize volunteer organizations to beta test the software through simulated consultations. One or two such organizations will be used and may not represent a good sample of the diversity of organizations that may be customers for the use of the expert system.

The domain experts selected are volunteers and have no vested interest in the success of the project other than their personal integrity to see the project through to the end. Any failures regarding availability, interest, completion of assignments and testing may result in failure of this project, including the knowledge acquisition process and the verification and validation of the expert system.

Definition of Terms

Business Resumption Services (BSR) is a term that is used to describe all critical business functions including computer/data processing, telecommunications, and support functions such as accounting and customer services.

Backward Chaining is an inference method where the system starts with what it wants to prove.

Circular Rules-a set of rules is circular if the chaining of these rules in the set forms a cycle.

Conflicting Rules-Two rules are conflicting if they succeed, in the same situation, but with conflicting conclusions.

Consultation is a term that is used to describe the interaction that takes place between an expert system and a "user" seeking advice.

Disaster (as it relates to data processing) is defined as an event, which causes the loss of computing services, or of a significant part of it, or of systems, communications or applications, for a length of time which prevents the impacted organization from achieving its mission or which imperils the business.

Disaster Recovery Planning (DRP) is a term that is used to describe activities that cover computer/data-processing activities.

Domain is a subject matter area or problem-solving task.

Domain Expert is defined as an articulate, knowledgeable person with a reputation for producing good solutions to problems in a particular field.

Expert System is defined as computer programs, comprising of both hardware and software that mimics an expert's (or experts) thought process to solve complex problems in a field or domain.

Expert System (ES) = Knowledge-Based System (KBS).

Expert System Shell is a software package containing a generic inference engine, a user interface, and a collection of other tools that enable users to develop and use expert systems.

Expertise is defined as the set of capabilities that underlines the performance of human experts, including extensive domain knowledge, heuristics rules that simplify and improve approaches to problem-solving, metaknowledge and metacognition, and compiled forms of behavior that afford great economy in skilled performance.

Forward Chaining is an inference method where the If-portion of rules are matched against facts to establish new facts.

Knowledge is the information an expert system must have to behave intelligently.

Knowledge Acquisition is defined as the process of extracting, structuring and organizing knowledge from several sources.

Knowledge Engineer (KE) is usually a computer scientist experienced in applied artificial intelligence methods who designs and builds the expert system.

Knowledge Engineering is the name given to the construction of knowledge-based systems.

Knowledge Representation is the method used to encode and store facts and relationships in a knowledge base.

Metaknowledge is knowledge in an expert system about how the system operates or reasons, such as knowledge about the use and control of domain knowledge.

Redundant Rules-Two rules are redundant if they succeed in the same situation and have the same conclusions.

Subsumed Rules- One rule is subsumed by another if the two rules have the same conclusions, but one contains additional constraints on the situations in which it will succeed.

Unnecessary If Conditions-Two rules contain unnecessary IF conditions if the rules have the same conclusions, an IF condition in one rule is in conflict with an IF condition in the other rule, and all other IF conditions in the two rules are equivalent.

Summary

Disaster Recovery Planning and Business Resumption Services are both time consuming and costly processes. However, recent disasters such as earthquakes, tornadoes and hurricanes have awakened management to the realization that such disasters could also strike them. This, together with the increasing dependence of organizations on data processing to perform basic business functions, have caused management to address the need to develop DRP and BRS plans. To address this need, many DRP vendors, as well as consultants, are being called on to advise customers on their DRP and BRS needs, as well as selling them services such as hot sites and backup services. Although consultants and vendors specializing in disaster recovery planning are available, their number is limited and the quality of their services may be questionable.

Several factors are hindering the ability to provide the services needed, including a lack of experienced consultants, the need for multiple skills sets for the consultants, and time and logistics obstacles. In addition, information gathering by consultants is a time consuming process and in many cases requires the use of multiple consultants, as well as various resources within the client's organization.

Expert systems are available to address shortages in expertise needed in specific domains. This development research project will attempt to design and develop an expert

system to assist the consultant in disaster recovery planning. The specific goals of this research include knowledge acquisition specific to the problems of using multiple domain experts, design and development of a prototype expert system for disaster recovery planning, and validation of the prototype expert system. The barriers and issues surrounding this research project are multiple and include those specific to disaster recovery planning and to knowledge acquisition and multiple domain experts. Two major elements of this research are the need for disaster recovery/business resumption plans and the need for an expert system to assist consultants in developing these services for customers. A third element is the knowledge acquisition process, during which the knowledge or expertise needed by the expert system is elicited and which is further complicated by the need for multiple domain experts. The fourth element is the design and development of the expert system utilizing an expert system shell. The fifth element is the validation of the expert system through beta testing, which will further allow for the validation of the knowledge acquisition process.

Several major research questions will also be addressed. First research question is whether the selection of a knowledge engineer knowledgeable in the domain areas will overcome many of the obstacles faced by knowledge engineers lacking such expertise. The second question is to identify knowledge engineering techniques that are successful in overcoming the traditional obstacles encountered by knowledge engineers in the past. A third research question is whether an expert system for disaster recovery planning that is user-friendly and can be utilized by the user as a training tool for DRP could be developed. A fourth question is whether the data gathered from the use of the expert

system will be complete and accurate and allow the consultant to submit a proposal for services to the customer in a timely manner.

Chapter 2

Review of the Literature

Historical Overview of the Theory and Research Literature Specific to the Topic

In completing the literature review for this dissertation, several areas need to be researched and investigated. These include disaster recovery planning, expert systems, knowledge acquisition, knowledge acquisition techniques, multiple domain experts, knowledge representation tools including flowcharts, dependency diagrams and decision tables, expert system shells, and the verification and validation process.

Disaster Recovery Planning

The catalyst for much of the research in disaster recovery seems to be initiated by disasters that hit major metropolitan areas. Hurricanes that hit the east coast and the not-too-infrequent earthquakes on the west coast seem to draw out the necessity for good, working disaster recovery and business resumption plans. These papers cover topics such as the effects of disasters and the consequences of the lack of a disaster recovery plan, the need for disaster recovery planning, and surveys on the existence of plans.

Garcia-Molina and Polyzois (1990) discuss issues in disaster recovery, including the motivation for remote backups and their applicability to various systems, the design, implementation and evaluation of remote backup mechanisms, and some of the criteria for evaluating and comparing various remote backup mechanisms. They noted consistency is of major importance for remote backups and if the backup does not preserve consistency, this inconsistency may lead to delays in transaction processing or

even to system crashes. These type of delays or crashes, in midst of a disaster event prompting the need for such backups, can be catastrophic to the disaster recovery plan.

Rudolph (1990) researched the requirements for restoring telecommunication services following a disaster. This paper was prompted by the Hinsdale incident, during which a fire at a central office of US West in the Chicago suburb left a half million customers, including many businesses, without service for several weeks. He noted that reality dictates that proactive, preplanned solution for restoring communications capability must be part of any contingency plan for companies that rely on telecommunications for critical business functions.

Hiles (1992) noted that the first step in any disaster recovery plan is to identify what risks exist and how significant they are within the specific environment of the organization. His list includes fire, flood, lightning strikes, malicious damage, power failure, impact, loss of supplied services or special consumables, rodent damage, insect infestation, contamination, telecommunications failure, industrial action, theft of equipment or data, loss of data, hardware or software failure, and viruses. He notes that this risk review should include a Critical Component Failure Analysis, which examines each component in delivering the computing service, identifies the threats to each component, determines what resilience exists in the event of its failure, and describes the impact of its failure. In addition, Hiles notes that in order to fully understand the impact of loss of service, a Business Impact Analysis must be performed to establish, in dollar and intangible terms, the value of each application to the business. Hiles, as well as others, outline the contents of disaster recovery plans to include immediate reaction procedures, restoration of the computing infrastructure, restoration of applications,

resumption of business processing under emergency arrangements and restoration of the permanent computer services. The plan must document, in exact detail, the function, personnel and equipment that must be re-established after a disaster (Jacob & Weiner, 1997). They also identified eleven steps to create a disaster recovery plan.

Jacob and Weiner (1997) noted that some of the problems encountered in completing a disaster recovery plan relate to the difficulty of gathering or eliciting the essential information to be detailed in the body of the plan. They also note that CPA's and other consultants interpret Statement on Auditing Standards (SAS) No. 60 as justification for bringing specific disaster planning issues into the audit process. SAS 60 requires the auditor to communicate to the audit committee of public companies any significant deficiency in the design or functioning of the internal controls structure which could adversely affect an organization's ability to record, process, summarize and report financial data, and provide communications for corrective action.

Cerullo and McDuffie (1994) surveyed 71 companies in Charleston, South Carolina area after Hurricane Hugo. They noted that out of 41 respondents, 56% had no disaster recovery plan and that 87% of these companies had computer downtime of 1-15 days, 9% had downtime of 16-60 days, and 4% had downtime of 2-4 months. Of the 23 firms without a DR plan, only 5 were able to process all critical accounting applications. A survey by Starr (1997) revealed that only 34 % of surveyed small businesses in Texas have either a formal or informal disaster recovery plan and that 44% of these have never been tested.

Money and Harrald (1995) discuss the techniques used to integrate the information and analysis provided by experts (i.e., knowledge acquisition) into the disaster

preparedness planning process. They noted that knowledge acquisition for DR planning is complicated because the knowledge must be derived from many different sources. The focus of their work is on the use of multiple experts as knowledge engineers and domain experts that can provide for the mix of knowledge that is required in a complex situation such as disaster recovery planning, which can cause differences in opinion because of the lack of knowledge about specific areas of a problems, statistical uncertainties, and use of different lines of reasoning. Their research illustrated how a strategic planning process can utilize multiple techniques for collecting and combining expert opinions to integrate the inputs and observations of experts in the disaster strategic planning process.

The need for DRP as a proactive approach is well documented in the literature. The consequences of failing to restore data processing services and its related business functions can be devastating to a company. Early studies at the University of Minnesota (1978) and University of Texas (1987) revealed that many businesses could not survive beyond five days without data processing capabilities. Hiles (1992) noted that within the first two days of a disaster, business efficiency can decline by 70 percent and that some financial applications in particular, can "go critical" within hours, if not minutes. A 1993 industry study commissioned by the then Digital Computer Company determined that 90% of companies that experience a catastrophic loss of data and equipment and do not have a disaster recovery plan are out of business within two years (Jacobs & Weiner, 1997). Rudolph (1990) uses the term "maximum time to belly up" (MTTBU) to describe the time frames that contingency planning managers have to restore data processing applications, telecommunications, and other essential services. Andrews (1994) noted that the literature states that it takes 12 to 18 months to develop an effective disaster

recovery plan and that once developed, it will require frequent maintenance to keep it current. He also noted that it is difficult to acquire the needed information from the knowledge of personnel within the data processing departments, as well as from others throughout the company.

Expert Systems

Expert systems, also called knowledge-based systems (KBS), are computer systems, which embody knowledge about a specific problem domain and can thus be used to apply the knowledge to solve problems from that problem domain (Smith, 1996).

Wolfgram, Dear & Galbraith (1987) define expert systems as computer programs, comprising both hardware and software, that mimic an expert's thought process to solve complex problems in a field or domain. They noted that suitable applications for expert systems include interpreting and identifying, predicting, diagnosing, designing, planning, monitoring, debugging and testing, instructing and training, and controlling. Baldwin-Morgan & Stone (1995) noted that the best problem domains for expert systems are those that are small, important, and have human experts and test cases available for development and validation.

Prerau (in Braden, B., Kanter, J. & Kopcsó, D., 1989) discusses a model for an expert systems solution in the area of diagnosis. This model describes a situation where (1) the problem is sufficiently complex but narrow in scope, (2) there is at least one recognized expert who is willing to act as a source of information and is articulate enough to be effective, (3) it takes the expert minutes to a few hours to solve the problem, (4) the expertise is scarce, (5) there appears to be a logical process to diagnose the problem that

does not require a great amount of intuition, and (6) there is a high payoff in problem resolution. Baldwin-Morgan & Stone (1995) add the need to preserve the expert's expertise. The principal conclusion of a paper by Hayes-Roth & Jacobstein (1994) was that knowledge based systems would be increasingly deployed as (1) assistants to human operators, (2) autonomous decision-making components of complex systems, (3) generators, critics and evaluators of configured information structures such as designs, plans and schedules, and (4) monitors of implementation and execution efforts that aim to activate or realize such encoded specifications.

Benefits of Expert Systems

Smith (1996) describes five benefits of expert systems and notes that expert systems (1) can make knowledge and expertise much more accessible than would otherwise be possible, (2) can be much cheaper than hiring the services of a real human expert, (3) can be used to preserve knowledge which would otherwise be lost over time, (4) are not prone to human error in the application of their knowledge, and (5) can be used to facilitate communication between humans and hence improve their own knowledge. Sangster (1996) adds the benefits of consistent and permanently reproducible performance, and a capability to provide expert level support even in situations where no human expert is present. Other reasons include enhancing product/service quality, gaining more insight into decision making, better control of complex systems, distributing scarce expertise, training less experienced employees, and monitoring vast amounts of information (Tsai, Necco, & Wei, 1994). Byrd (1992) has given evidence that expert systems are mostly used as decision makers for lesser skilled domain

personnel and as advisors to assist the expert and that they provide personnel with ways to make more consistent, timely and accurate decisions, in hope of improving completeness.

While expert systems have been very successfully adapted throughout industry and commerce, their main area of development has been in accounting, finance and manufacturing (Hayes-Roth & Jacobstein, 1994). This proliferation is mainly due to the ability of the technology to address three key areas: improved quality and dependability of work, upgraded customer service and improved productivity, and capturing the expertise of human experts (Awad & Lindgren, 1992).

Problems Areas Associated With Expert Systems

Tsai et al. (1994) list eleven major problem areas associated with expert systems. Of these, the three areas most encountered are the difficulty of integrating within the existing data processing environment (47%), strong resistance to change (43%), and the difficulty in finding experienced knowledge engineers (40%). They cite that two major weaknesses in expert system development are the lack of management support (23%) and the lack of user involvement (19%). They noted that reasons for not using expert systems include problems in finding knowledge engineers (20%) and the difficulty in finding adequate domain experts (16%), and that knowledge acquisition was the most difficult task in expert system development, followed by verification and validation. Gorney & Coleman (1991) note that many organizations do not follow any development standards for expert systems and that this failure is cited as one of the reasons for expert system failure.

Others reasons listed include poor idea selection, poor planning, inadequate funding, overly optimistic scheduling and technology problems.

Knowledge

Hayes-Roth & Jacobstein (1994) define knowledge as those kinds of data that can improve the efficiency and effectiveness of a problem solver. They describe three major types of knowledge: (1) facts that express valid propositions; (2) beliefs that express plausible propositions; and (3) heuristics or rules of thumb that express methods of applying judgment in situations for which valid algorithms generally do not exist. They defined "expert knowledge" as knowledge used by problem solvers to find an acceptable solution that meets or exceeds requirements with a reasonable expenditure of resources. Specifically, they noted that expert knowledge helps problem solvers improve their efficiency by marshalling relevant facts, avoiding common errors, making critical distinctions between problem types, pruning useless paths of investigation ordering search, eliminating redundancy, reducing ambiguities, eliminating noise in data, exploiting knowledge from complementary disciplines, and analyzing problems from different perspectives or levels of abstraction.

Knowledge Acquisition

The process of capturing knowledge is defined as the collection, organization, evaluation, and incorporation of knowledge within a working expert system (Lichti, 1993). In designing expert systems, the process of eliciting information has been termed knowledge acquisition. According to Hoffman (1987), knowledge acquisition, also

known as knowledge elicitation, involves extracting problem-solving expertise from knowledge sources, which are usually domain experts. Waterman (1985) defines knowledge acquisition as the process of extracting, structuring and organizing knowledge from several sources, usually human domain experts, so it can be used in a program. Smith (1996) noted that knowledge acquisition involves the elicitation of data from the expert, interpretation of the data to deduce the underlying knowledge and creation of a model of the expert's knowledge in terms of the most appropriate knowledge representation. This knowledge acquisition process involves one or more knowledge engineers interacting with one or more domain experts, each of which 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 (Fellers, 1987). The knowledge engineer must: (1) familiarize himself with the domain of the expert, (2) clearly identify the areas of the domain that needs modeling, (3) and represent this knowledge in a form that can be computerized (knowledge representation) (Smith, 1996).

Sharp's (1994) discussion of knowledge acquisition notes that it involves (1) problem definition, (2) identifying a suitable knowledge source and gathering suitable knowledge, (3) modeling the knowledge in a form that can be used by the knowledge based systems, and (4) implementing and refining the knowledge base.

The acquisition of knowledge is a major and critical phase in the development of expert systems. Knowledge acquisition is considered by many to be the most difficult and precarious stage in the knowledge engineering process (Smith, 1996). Tsai et al. (1994) noted that this was because knowledge acquisition involves communications between people with completely different backgrounds, human experts and knowledge

engineers, who must formulate the concepts, relations and control mechanisms needed for the expert system. Hoffman (1987), Rook & Croghan (1989), Byrd (1992), Liebowitz (1993), Hwang (1994), Jeng, Lieng, & Hong (1996), amongst others, stated that knowledge acquisition has often been described as the bottleneck in knowledge based systems development. Keyes (1989) noted that the process of gathering knowledge is a fuzzy process. Byrd (1992) has further described knowledge acquisition as a separate and distinct process from knowledge engineering.

Knowledge Acquisition Vs. Requirements Analysis

Sharp (1994) noted that there are many parallels between knowledge acquisition and requirements analysis. He stated that the goals of the two are common and can be characterized as gathering information and modeling it in a form suitable for implementation. Byrd, Cossick & Zmuk (1992) noted that knowledge acquisition is nothing but an extension of requirements analysis, which involves end users and systems analysts interacting in an effort to recognize and specify the data and information needed to develop an information system. Their research showed that many of the entities and processes in knowledge acquisition and requirements analysis are almost identical. A widely supported view of a knowledge acquisition life cycle, as described by Buchanan (in Fujihara, Simmons, Ellis & Shannon, 1997), involves identification, conceptualization, formalization, implementation, and testing. During the identification stage, the most important characteristics of the problem domain are defined. Identifying structure protocols is performed in the conceptualization phase. During the formalization stage, the knowledge acquired is mapped into representation and then coded during

implementation. The final testing phase is where the knowledge base is validated. These five stages represent an iterative refinement process, which must be performed by the knowledge engineer.

Although most systems analysts do receive training in requirements analysis, little or no training is available at the college or professional level in knowledge acquisition (Kemp & Purvis, 1996). Liebowitz (1993) strongly suggests that knowledge engineers need a stronger foundation in the fundamentals, methodologies, techniques and tools for gathering knowledge. It has been well documented (Smith, Ross, Awad, Green & MacIntyre, 1994; Awad & Lindgren, 1992) that there is a lack of standardization of training and qualifications for knowledge engineers. These same authors conducted surveys in both the US and UK, which established similarities and differences between knowledge engineers and systems analysts, examined the skills and personality attributes required in each discipline, and looked at the methods and technology in use in KBS development. Both of these surveys also demonstrated that information gathering skills and personality attributes are extremely important to both knowledge engineers and systems analysts and that diplomacy seemed to be the skill that knowledge engineers needed to work on. They also noted that the seven most important skills for knowledge engineers are (from the most important) knowledge representation, fact-finding, human skills, verbal skills, analysis, creativity and management.

Why Expert Systems Fail

O'Neil (1989) attempted through surveys to demonstrate why the vast majority of expert systems fail. Five reasons noted by him and others are:

1. The lack of user participation in design (Rees, 1996);
2. The lack of structure and organization of knowledge acquisition (McGraw & Harbison-Briggs, 1989);
3. Communication problems between the knowledge engineer and the domain expert (O'Neil, 1989; Jeng, Lieng & Hong, 1996);
4. Failure in identifying the right candidates for knowledge acquisition (Stein, 1993); and
5. Failure of verification and validation (O'Leary, 1990; O'Keefe, Balci & Smith, 1990; Vanthienen & Robben, 1993).

In addition, the knowledge engineer may have little knowledge about the domain and may not understand the jargon used by the domain expert (Jeng, Lieng & Hong, 1996). To address this, knowledge engineers must be willing to repeatedly ask the expert what he or she knows, be willing to be continually corrected by the expert, and be capable of eliciting information that is second nature to the expert (Braden, Kanter & Kapcsó, 1989).

Rook and Croghan (1989) suggest a systems engineering conceptual framework to address the knowledge acquisition bottleneck in the knowledge based system development rather than focusing on specific knowledge acquisition methods and techniques. They noted that a framework that “structures” knowledge acquisition steps within the context of the expert system development phase is critical to overcoming the difficulties of knowledge acquisition. Specifically, they outlined a structured knowledge acquisition process in order to: (1) specify the goals of knowledge acquisition as they relate to the expert system development cycle phases; (2) identify the specific steps or tasks involved in each phase; (3) specify the specific goals of each knowledge acquisition

step; (4) identify constraints or barriers to successful knowledge acquisition activities; and (5) provide a knowledge acquisition task structure that will facilitate the selection of appropriate knowledge acquisition methods.

Knowledge Acquisition Techniques

Key to the success of the design and development of any expert system is the selection of the correct or most appropriate technique for knowledge acquisition. As noted by Niederman (1996), much research is available discussing the various techniques. Choices include interviewing, observations (Olson & Rueter, 1987) and protocol analysis (Wolfgram, Dear & Galbraith, 1987). Fellers' (1987) discussion included structured and unstructured interviews, constrained-processing tasks, including both simulation of familiar tasks and the use of scenarios, and the use of "tough cases." Byrd, Cossick and Zmud's (1992) list includes observation of the expert in action, unstructured elicitation which corresponds to unstructured interviews, mapping domain knowledge as a combined effort between the knowledge engineer and expert to develop a cognitive map, formal analysis of the domain, and structured elicitation corresponding to structured interviewing. Keyes (1989) noted that no single technique is best. Byrd, Cossnick & Zmud (1992) and Fellers (1987) agree that it is best to use several techniques in some combination, depending on what works in a particular environment (Niederman, 1996). The interview process remains the most frequently used technique for extracting knowledge from human experts (Smith, 1996). A survey by Smith, Rose & Awad (1994) revealed that 77 percent of knowledge base systems have been developed using interview techniques. Fujihara et al. (1997) noted that interviewing continues to be the primary

method of acquiring expert knowledge, requires little equipment and can yield a considerable amount of knowledge, if the knowledge engineer is skilled. Their research noted that it would be useful to have a computerized knowledge and conceptualization tool to assist the knowledge engineer in extracting and structuring knowledge from the interview data. Their research revealed such a tool could retrieve knowledge in about half of the time for the manual process and that the number of knowledge components retrieved from knowledge acquisition is about four times greater. Jeng & Hong (1996) noted that automated approaches using interactive induction methods have been attempted but suffer from the unavailability of sufficient case studies to predict a problem solving behavior.

Multiple Domain Experts

The key parties in expert systems design and knowledge acquisition are the knowledge engineer and the domain expert(s). A domain expert is defined as an articulate, knowledgeable person with a reputation for producing good solutions to problems in a particular field (Waterman, 1985). Methods to identify candidates for knowledge acquisition have been discussed (Stein, 1993). Key to the identification is whether the expert really has the knowledge needed to meet the goals and objectives of the expert system. The selection of domain experts may depend on several criteria as discussed in Liou (1999). Selection should be based on expertise, experience or reputation of the individual(s). In addition, the individual(s) should possess business and personal attributes that lend themselves to the knowledge acquisition process. Thirdly, the individual(s) must be available during both the knowledge acquisition process and

expert system development process, including the validation phase of the project. Stein (1993) discussed the use of network analysis as a method for selecting candidates for knowledge acquisition, especially in situations where the knowledge is specific to the organization.

Most early expert systems were developed using a single domain expert (O'Neil & Morris, 1989). Few involved multiple experts and the problems of knowledge acquisition that occur with the use of multiple domain experts. They noted that the main reason for this seemed to be that it was easier to elicit knowledge and to avoid contentious issues and conflicting opinions. However, knowledge engineers soon realized that expertise might not be resident in the knowledge of a single domain expert. As each expert's expertise may be limited to his own domain, if an expert's domain does not cover the problem area, incorrect decisions may surface. Also, over confidence or ignorance may lead to errors. Liou (1999) also cited the difficulty of merging each individual expert's knowledge into one group structure that provides the underlying problem solving expertise of the expert system, and the difficulty of the generation of group knowledge that does not reside in any one individual expert but resolves as a result of group interaction. The use of multiple domain experts can help eliminate the elicitation of undesired or bad knowledge caused by these problems (Hwang, 1994). Multiple experts can provide the mix of knowledge that is required in a complex structure, such as disaster recovery planning, and provide coverage for the many problems and solutions (Money & Harrald, 1995). Liou (1999) noted that incorporating the expertise of a team of experts provides for positive effects on the resulting expert system. These positive effects are: (1) it assures that the knowledge base can be complete; (2) it improves the likelihood of

obtaining specialized knowledge in subdomains of the problem; (3) it increases the quality, i.e., reliability and consensus among experts, of the acquired knowledge; (4) it assures that the facts that are included in the knowledge base are important ones; (5) it enhances understanding of the domain knowledge through discussion, debate and exchange of the hypotheses between members of the expert team; and (6) it encourages interactions among the experts and creates a synergy such that the acquired group knowledge is greater than the sum of the individuals' knowledge. Liou also noted several techniques for collaborative knowledge acquisition, including brainstorming, nominal group technique, Delphi technique, focus group interviews, voting, group repertory grid analysis, and group support systems. Niederman (1996) argues for the desirability of investigating the expertise that facilitators use to lead groups as a tool for knowledge acquisition among multiple experts.

Knowledge acquisition involving multiple domain experts is fraught with the problems of dealing with a single expert, compounded with the obstacles and risks of trying to coordinate human interactions and integrate multiple knowledge bases (LaSalle & Medsker, 1991). The problems of using multiple domain experts include the issues of conflict between experts, and the failure of the knowledge engineer to express the relationship between multiple views in requirements specifications which may overlap, complement and contradict each other (Nuseibeh, Kramer & Finkelstein, 1994). Hwang (1994) and Liou (1999) noted that one of the most difficult problems of knowledge acquisition is to integrate domain knowledge from multiple experts, especially when inconsistencies and conflicts need to be resolved. Hwang also noted that experts of most application domains seem to be very busy and that it may be impossible to get them

together. Added to this is the distant physical location of the domain experts and travel costs to get them all in one place. Hwang also noted additional problems in trying to elicit expertise from multiple domain experts, including the difficulty of integrating knowledge from the various experts, differences of opinion, and the difficulty in conflict resolution due to the insistence of the experts that "they are right." Reasons for differences of opinion can also be due to one expert requiring more data than another expert, the expert's lack of interest in the area, the differences in weighting the importance of the item, or the fact that the expert may not consider it important.

O'Neil & Morris (1989) noted ways to avoid conflict and dissent: (1) asking the experts to provide documentary back-up evidence; (2) using probabilities and uncertainty factors to express degrees of agreement or disagreement amongst the experts; and (3) creating systems modularly so that different experts could be specifically called on for particular sections and consultations.

In attempting to integrate multiple requirements specifications, overlaps must be identified and expressed, complementary participants made to interact and cooperate and contradictions resolved (Nuseibeh et al., 1994).

Flowcharts/Dependency Diagrams/Decision Tables

Three key tools in the design of expert systems utilizing rule-based systems are flowcharts, dependency diagrams and decision tables. These tools were prominent in the early years of expert system development (O'Neil & Morris, 1989) and continued to be used for fuzzy logic applications in the nineties (Chen, 1993).

Flowcharts can be used to represent the flow of the logic including the order of asking for data needed during the expert system's consultation process and the branching of the flow based on the data received or calculated by the application. Graphics based flowchart-editing tools such as Flowtool (Watkins, Dimopoulos, Neville, & Li, 1993) allow the knowledge engineer to quickly create and edit the flowcharts.

Dependency diagrams and decision tables are used to represent the decision-making process based on all of the data inputted or calculated. Decision tables are easily updated and converted into rules. Both of these tools can be used for design purposes and can be further utilized as verification and validation tools.

Expert System Shells

During the early years of expert system development most expert systems were developed using one of several available programming languages like Prolog and LISP that were appropriate for expert systems development. These early expert systems were well received but took years to develop. The developers of MYCIN, an expert system designed to provide physicians with advice on bacteremia and meningitis infections, soon realized that most expert systems utilized the same logic. They discovered that the key differences were the knowledge and data required by the inference engine. By stripping away the knowledge and data from MYCIN, the resulting shell (EMYCIN) could be used to develop new expert systems.

Expert system shells are software packages containing a generic inference engine, a user interface, and a collection of other tools that enable users to develop and use expert systems. Using the shell's tools, a knowledge engineer can develop new knowledge

bases and can structure, add, delete, and modify the knowledge contained in them (Stylionou, Smith, & Madey, 1995).

Expert system shells have been available since the late eighties and early nineties. Most were developed in the DOS environment and those that survived were modified for use in Windows. However, the loss of interest in these shells left many good ones, such as VP-Expert, dropped by the developers or not supported. A 1992 survey by the public accounting firm of Touche Ross conducted for the Department of Trade and Industry (in Smith, 1996) noted that shells were used to develop 42% of the surveyed expert systems. Durkin (1996) noted that more than 40% of surveyed experts systems were developed using shells and that more than 60% were PC-based. Preece & Moseley (1992) conducted a study comparing use of a shell to programming in Prolog and concluded that the shell was significantly more effective overall, in terms of rate and speed of development, efficiency in testing and debugging.

Which expert system shell to use is a key decision to be made by the knowledge engineer. Using the wrong shell can result in project failure. There is agreement in the literature that different application types require different expert system shell capabilities. In their research, Stylionou, Smith & Madey (1993) defined taxonomy of expert system application types with similar capability requirements into categories with each category called a generic application category. In subsequent research Stylionou et al. (1995) noted that the selection of a shell depends on factors such as application and project characteristics, capabilities of the shell, user sophistication, ease of integration with existing software and hardware, and vendor support. The shell's inference engine is typically guided through one of two directions, forward chaining or backward chaining.

Although backward chaining has been viewed as the most efficient method, expert systems for design applications should utilize forward-chaining since its desired goals are not known in advance or are too numerous to list and it is useful in creating data driven processes and expert systems with unknown goals (Song, Strum & Medsker, 1991).

Verification and Validation Process

Many expert systems in the past have failed to include any verification and validation phase and as a result have failed. Similar to the testing phase of conventional structured application design, the verification and validation process are used to ensure that the system does what it is supposed to do. Verification and validation has become a key issue for expert systems used in real-world applications that require a high degree of reliability (Motoda, Mizoguchi, Boose & Gaines, 1991). The verification and validation process must include all of the players involved, including domain experts, knowledge engineers and users.

Although verification and validation are related, they are different processes and often confused with each other. Validation refers to building the right system; that is, substantiating that the system performs with an acceptable level of accuracy. Verification refers to building the system “right”; that is, substantiating that a system correctly implements its specifications (O’Keefe, Balci, et al., 1990). While verification tests for the consistency and completeness of the expert system, validation analyzes the knowledge base and decision-making capabilities of the expert system for content validity, the level of expertise and the reliability of the expert system. O’Keefe, Balci, et al. noted the major problems in validating an expert system’s performance are: (1) what

to validate; (2) what to validate against; (3) what to validate with; (4) when to validate; (5) how to control the cost of validation; (6) how to control the bias; and (7) how to cope with multiple results. They noted that the alternatives of “what to validate” include intermediate results, or the final results or conclusions, or the reasoning of the logic, or any combination of the three. The choices for “what to validate against” include known results or expected performance. As noted by O’Leary (1990), validation is important to the decision-making success of an expert system and to the continued use of the expert system. An expert system that does not make the correct decisions will lead to a loss of confidence in it. Therefore, expert systems must be validated prior to implementation in the field for real world use. O’Leary also noted that validation requires: (1) ascertaining what the system knows, does not know, or knows incorrectly; (2) ascertaining the level of expertise of the system; (3) determining if the system is based on a theory for decision-making in the particular domain; and (4) determining the reliability of the system. Nguyen, Perkins, Laffey, & Pecora (1990) noted particular problems in the knowledge base when checking for consistency and completeness. They listed five problems for consistency: redundant rules, conflicting rules, subsumed rules, unnecessary IF conditions, and circular rules. In their research they indicated that any one of four situations can be indicative of gaps or missing rules in the knowledge base, including unreferenced attribute values, dead-end goals, unreachable conclusions and dead-end IF conditions.

Tools utilized in the verification and validation process include those specific to expert systems, as well as some of the tools used in conventional application design. However expert systems have different characteristics than conventional software and

require different verification and validation methodologies ((Motoda et al., 1991). These include the explanation facility of expert system shells, flowcharts, decision trees and decision tables. The latter has been used in a vast majority of cases to provide for extensive verification and validation assistance. It easily allows the knowledge engineer to check for contradictions, inconsistencies, incompleteness and redundancy in the problem specification (Vanthienen & Robben, 1993; Nguyen et al., 1990).

Summary of What Is Known and Unknown About the Topic

In the area of disaster recovery planning and business resumption services much is known and expressed in the literature. The vulnerability of data processing and related functions to natural and man-made disasters is unquestioned. Research seems to evolve soon after such events with the successes and consequences of data processing restoration surveyed. Also known is the value of expert systems to mimic the expertise of domain experts. Literature from the late eighties and early nineties is abundant. The use of expert systems, its successes and failures, benefits, and problems associated with design and development has been researched. As a subset of this literature, research in knowledge acquisition as a critical step in the expert system development cycle has been developed. It is well documented that knowledge acquisition is not only the most critical phase, but also considered by many as the most difficult and precarious stage in the knowledge engineering process and often described as the bottleneck. This is complicated when multiple domain experts are needed. The literature has attempted to develop solutions to these difficulties including automated knowledge acquisition tools and new techniques for knowledge acquisition. Another subset of the literature on expert

systems is the discussion of verification and validation as a tool to enhance the success of expert systems by ensuring that the expert system makes the "right" decisions, in a consistent basis that is complete.

Although the literature has addressed knowledge acquisition issues concerning single domain experts, there is less research into the complications of multiple domain expert systems. Unknown is the effectiveness of using knowledge engineers with excellent (or high degree) of domain knowledge to control all domain experts. Methods to resolve multiple domain experts conflicts needs to be expanded and tested to determine which are effective under different situations. The concept of segregating domain experts to avoid conflict and its effectiveness needs to be researched. The effectiveness of having knowledge engineers make decisions to avoid or resolve conflicts as a control in using multiple domain experts needs to be further researched.

The Contribution this Study Will Make to the Field

This study will attempt to contribute to the field in several ways. First it will add to the literature on expert systems, specifically the design and development of an expert system using multiple domain experts and knowledge acquisition. It will attempt to develop alternative methods to address the difficulties with knowledge acquisition and multiple domain experts. To a lesser degree, this research will add to the literature on disaster recovery planning and business resumption services through the design and development of the DRP expert system. Thirdly, it will add to the literature in accounting and consulting by providing accountants and consultants with a tool to provide consulting services for disaster recovery planning.

Chapter 3

Methodology

Introduction

Four research questions as discussed in the Elements, Theories & Research Questions section of this paper will be addressed. The methodologies to be utilized to answer these questions are as follows:

Question one will attempt to answer the question whether a knowledge engineer experienced in the domain area can overcome many of the obstacles faced by knowledge engineers who have lacked experience in the domain area. Past research has demonstrated that many expert systems have failed because they do not do what they needed to do. Many point out that this may be due in part to the inability of the knowledge engineer to elicit all of the knowledge needed from the domain expert(s) and that this failure may be due to the lack of domain knowledge of the knowledge engineer (Jeng, Lieng & Hong, 1996). To address this question the knowledge engineer selected is an expert in disaster recovery planning and has developed and audited disaster recovery plans. This expertise will allow him to direct and question the domain experts and assist in the resolution of conflicts between domain experts. It is felt that his expertise will result in a good working expert system for DRP that is operational, complete, consistent, meets its objectives, and avoids redundancies. Although no formal methodology will be utilized to test this approach, notes will be taken to document problems encountered during the knowledge acquisition process and a hypothesis may be made that any obvious lack of problems are due to the use of the domain expertise of the knowledge engineer.

Question two will attempt to identify knowledge engineering techniques that are successful in overcoming the traditional obstacles encountered by knowledge engineers in the past. The use of available methods of conflict resolution and alternatives such as non-group decision-making will be considered, as well as group versus individual meetings.

Question three will attempt to answer the question whether an expert system for DRP can be developed that is user-friendly and can be utilized by the user as a training tool for DRP. The design and development of the expert system will follow the SDLC methodology as noted in the expert system design section of this chapter. The details of the design specifications of the expert system, its user requirements, hardware requirements, and feasibility are discussed. The knowledge engineer will help to ensure that the system design is easy to use and user friendly. This will be achieved in part through the use of the design tool, VP-Expert, which allows for the development of expert systems that are easy to use. The use of multiple-choice questions will also facilitate the ease of use. The use of the explanation facility of VP-Expert will help achieve the objective of providing an imbedded training tool for DRP by providing the user with the ability to ask the system why a question is being asked or why the information is needed. During beta testing in the validation phase, users will be advised of this feature and asked to use it and comment on its use as a training tool. Extensive testing of this feature can be the subject of future research.

Question four will attempt to answer the question whether the data that is gathered from the use of the expert system is complete and accurate and allows the consultant to

submit a proposal in a timely manner. This will be answered by beta testing during the validation phase. At least two volunteer organizations will undergo a consultation session using the expert system. The results of the consultations will then be presented to the domain experts, who will be asked to determine if they have all of the information they need to submit a timely proposal.

EXPERT SYSTEM PROJECT DESIGN

The design and development of the expert system will utilize methodologies consistent with the design and development of conventional applications, modified for areas specific for expert systems such as knowledge acquisition, verification and validation. Research in disaster recovery planning, expert systems, knowledge acquisition, knowledge acquisition tools, and multiple domain experts will be used to provide a basis for the knowledge needed by the knowledge engineer to successfully complete the design and development of the expert system.

APPLICATION DESCRIPTION & SYSTEM OBJECTIVES

Timing, logistics and cost problems are considered detrimental to the ability to provide DRP and BRS services. The ideal solution would be to take the personnel constraints out of the process. This may be accomplished by simulating the expertise of each consultant into an expert system. The solution is to provide a single salesperson or consultant with the tool (expert system) needed to obtain the necessary information to assess the customer's needs and provide a timely proposal.

The purpose of the expert system, called the Disaster Recovery Plan (DRP) Analyzer, is to assist the consulting team to evaluate potential customers for services offered. The customer information obtained, coupled with the expertise of the consulting team members, will be used by the team to prepare recommendations and a "quick price" quote, addressing the customer's disaster planning and recovery needs. The specific objectives to meet the above are:

1. Examine (audit) the customer's disaster preparedness plans, if any, and identify critical business functions, i.e., data processing, telecommunications, business operations, and human resources, that need to be adequately protected from the impact of a disaster.
2. Identify which contingency planning service, if any, should be proposed to the customer, as well as alert the consultants to other selling opportunities.
3. Provide a printout covering the information the consultants require in order to prepare a "quick quote price proposal."
4. Provide a "valued" experience to the customer and raise the customer's awareness of disaster vulnerabilities and educate them on the benefits received from a professional contingency planning vendor for DRP and BIS services.
5. Provide a learning experience to the user, customer, and others through the 'how' and 'why' command option during expert system consultations.

SYSTEM JUSTIFICATION

The justification for the development of this expert system is based on time and cost. Without an expert system, it would be required that the consulting team visit each

customer in addition to the sales member. This would require a minimum of four consultants in addition to the sales member. Timeliness of a turnaround proposal to the customer is also critical. The development and use of the proposed expert system will:

- Reduce the need for group sales/consulting teams;
- Reduce personnel and travel costs;
- Ensure consistency and completeness in obtaining information needed; and
- Provide a timely proposal to the customer.

USER REQUIREMENTS

The goal of this design project is to develop an expert system, called Disaster Recovery Analyzer, for use by consultants (user) in selling disaster recovery/business resumption services. The concept of this contingency planning service is based on the ability to provide a multi consultant solution to the customer's needs for various business resumption services. This one stop customer solution is based on the concept of ensuring coordination of all disaster recovery solutions. This multi consultant approach presents several problems, including the:

1. Need for each consultant to be available for meeting with each customer;
2. Consistency of the expertise of each consultant at the meetings;
3. Availability of all consultants at the same time and when needed;
4. Need to meet with several different customer personnel; and
5. Need for speedy turnaround time from customer consultations to providing the proposal.

The timing, logistics and cost problems of the above are considered detrimental to the success of this process. A strategic solution is to take the personnel constraints out of the process, which would best be accomplished by simulating the expertise of each consultant representative into an expert system. This strategic solution will provide a single consultant team member with the tool needed to obtain the information needed to access the customer's requirements and provide a timely proposal to the customer. This system is considered a strategic system in that without the expert system the assessment of the customer's needs for disaster recovery/business resumption services would be costly.

INPUTS / OUTPUTS

Input into the expert system will be via real time consultation with the customer. The expert system will be designed into five modules, each addressing a general or specific area. The modules are designed to facilitate input (answers) from the appropriate customer MIS representative, who may be different for each module. Input will be answers to the questions in each module (see Figure 2).

A general description of each module, the type of input needed to achieve the objective of the module and the appropriate customer representative include:

1. INTRODUCTION MODULE:

The Chief Information Officer (CIO) will typically answer these questions. They are simple biographical types of questions.

2. GENERAL PREPAREDNESS MODULE:

The CIO or the Disaster Recovery Manager (DRM) will answer these questions.

The purpose of this module is to determine the company's preparedness and plans of action.

3. DATA PROCESSING MODULE:

The CIO or the DRM will answer these questions. The purpose of this module is to gather information identifying hardware, software, critical applications, and other data processing areas.

4. PLAN TEST/MAINTENANCE MODULE:

The CIO or DRM will answer these questions. The purpose of this module is to gather information on the company's plan maintenance practices including team members, testing, distribution and updating.

5. TELECOMMUNICATIONS MODULES:

The Network Manager will answer these questions. The purpose of this module will be to identify the company's telecommunications system including voice, data, configuration, lines, trunks, and vendors.

The input (answers) to the above questions will provide all consulting team members with the information necessary to prepare a "quick" quote for proposal, the ultimate objective for the expert system. Also, via the opportunity clauses, it will provide the team with other opportunities for selling and/or servicing the client. This information will be provided to the team member via the output described below.

Output from the expert system will include (see Figure 2):

1. A data file for each module, which will contain the inputted responses to each question.
2. A print file for each module which will produce a report, including questions, answers, and opportunity clauses.
3. Hardcopy documentation of the consultation, including each question, the customer's answer inputted during the consultation, and an "Opportunity Clause" which the system will generate as a business opportunity to be considered as a result of the answer to the question.

COST/BENEFITS

In justifying the development of any application, cost and benefit information is needed. Although many of the benefits of an expert system are intangible, a cost benefit analysis must be completed.

Costs

Initial costs are based on the estimated fee to design and develop the expert system. In addition, one time hardware costs for laptop computers and printers are estimated. These costs include estimated startup costs, software development costs, and hardware costs¹. On-going costs are based on an assumed fee for annual software maintenance and costs for operating the team. These include annual operating costs, such as annual software maintenance, sales personnel expense, sales travel expense and overhead.

Benefits

The benefits associated with the expert system can be measured in both tangible and intangible terms that are in both monetary and non-monetary terms.

Key non-monetary benefits are:

- The use of a one-person sales team is anticipated with the use of the expert system, whereas, without the expert system, all consultants will be needed at each customer sales meetings.
- Ability to provide proposals in a timely basis. It is estimated that without the use of the expert system it would take, at a minimum, approximately two weeks from the time the information is obtained to the delivery of the proposal to the client. This time lag may be unacceptable. It may cause loss of business opportunities. With the expert system it is anticipated that the proposal will take no more than one week, cutting the turnaround time in half.
- It is anticipated that within a short time, refinement and enhancements in the expert system (through the addition of a proposal module) would allow for immediate proposals within ten-percent accuracy of estimated costs.
- Provide sales personnel with a training tool. It is anticipated that the expert system will have a help function (explanation tool), which will allow for explanation of why the questions were asked or why the information is needed. This will educate sales personnel as well as allow the sales member to explain to the customer why the information is needed.

¹ Based on current prices, it is estimated each laptop will cost \$3000 (including installed software) for a total cost of \$66,000. In addition, 1 high-speed color ink jet printer will be required for printing reports and proposals. Estimated cost=\$1,000.

Monetary benefits are expressed in savings at the personnel and travel areas. Estimated savings per customer visit and estimated annual savings as listed below.

Savings per Customer Visit:

Personnel Costs \$4,000 (based on \$1,000 salary savings per consultant)

Travel Costs \$6,000 (based on \$1,500 travel savings per consultant)

Annual Savings based on One Customer Visit per Week:

Personnel Costs \$200,000 ($\$4000 * 50$ visits/year)

Travel Costs \$300,000 ($\$6000 * 50$ visits/year)

Total Savings \$500,000

Operational Feasibility:

Based on presently available development tools such as VP Expert, the expert system is operationally feasible. The runtime version of the software will allow for easy implementation of the software. Research has determined the runtime version of the package should be fairly easy to learn and that any needed training for the team members will be minimal.

Technical Feasibility:

Based on present available development tools such as VP Expert, the expert system is technically feasible. The shell development tool will allow for the design and development of the expert system. The runtime version of the software will allow for easy implementation of the software. Research has determined the runtime version of the

package should be fairly easy to learn and that any needed training for the team members will be minimal. The platform for operating the expert system is based on existing technology. Once a working prototype is completed and beta tested, it is estimated that conversion to a Windows based system can be easily completed.

DESIGN SPECIFICATIONS

System Architecture

Hardware

Specifications require that all consultants utilize a laptop computer to run the expert system consultation at the customer location. Minimal specifications for each machine are:

- 500 MHz processor
- 64 MB RAM (expandable)
- 6 GB hard drive
- K56 fax/modem
- Sound card (16 bit stereo sound)
- 14-15" active matrix color display
- Microsoft Windows 98 or Windows 2000
- Microsoft Office 2000
- Carrying case
- Extended life battery
- Extra battery
- 3 year warranty

Software

The expert system will be developed utilizing VP EXPERT, an expert system development tool. The expert system will be copyrighted under the copyright laws for software. The expert system name will also be copyrighted with appropriate trademarks.

DESIGN METHODOLOGY

A. Establish Precise User Requirements

1. Meet with potential users to determine their requirements. Included will be information needs, report needs, formatting, screen input, and output design. Additionally, minimum hardware requirements will be established.
2. Meet with consultants to determine information needs. Specifically:
 - Telecommunications consultant/expert,
 - DRP consultant/expert (for General Preparedness section),
 - Site vendor consultant/expert, and
 - Industry consultant/expert.
3. Interview each consultant to determine what information each requires in order to provide a quote for his appropriate section of the plan. Questionnaires will be used for high-level information. Detailed interviews will be conducted.

This review will include:

- Confirmation that the information is truly needed;
- Validation that the wording of the question is clear to the customer;
- Establish why the question is being asked (explanation); and

- Discuss the order of the questioning.

The consultants/domain experts will be asked to develop the list of questions they need answered. Duplicate questions will be eliminated. Multiple-choice answers will be used whenever possible, with an optional choice of “other” if needed. Questions will be categorized according to the following five modules:

1. Introduction.
 2. General Preparedness.
 3. Data Processing.
 4. Plan Test/Maintenance.
 5. Telecommunications.
- B. Flowchart rules to ensure correct order of questioning and branching.
- C. Prior to programming, verification and validation will be performed with each domain expert to ensure accuracy of questions and appropriateness of answer types (i.e., multiple choice or fill in). All questions and flowcharts will be validated and approved by the domain experts prior to programming.
- D. Coding rules.
- E. Final testing with the domain experts.
- G. Beta testing with volunteer customers.

APPROACH

This expert system design and development will utilize a phased approach. Approaches and methodologies will be taken from the disaster recovery literature, the

knowledge acquisition literature, and the expert system literature. Our approach will include four phases (see Figure 1) as discussed below:

Phase 1:

This initial phase will entail the knowledge acquisition process from three domain experts to elicit the information needed (requirements specifications) by the various disaster recovery vendors to determine the data processing environment of the customer. This knowledge acquisition process will be based on a combination of the methods of knowledge acquisition, including interviewing techniques, documentation such as textbooks and available literature, review of examples, and case studies.

Phase 2:

This phase will include the design of the dependency diagrams and decision tables needed by the expert system to determine the appropriate disaster recovery plan services for the customer. Flowcharting tools will be used to document the flow and branching of questions during the expert system consultation. Both forward and backward chaining will be considered.

Phase 3:

During this phase the prototype expert system will be developed. Initial design plans call for the use of VP Expert, an expert system shell development tool with the final version compiled for use as a standalone application.

Phase 4:

During this final phase, the prototype expert system developed in phase 3 will be validated. Attempts will be made to validate the decision-making capability of the expert systems by comparing its conclusions with those of domain experts during real-life walkthroughs and/or simulations and/or case studies. Methodologies utilized for the validation phase are described below in step 8.

The specific steps are:

1. Research in the area of knowledge acquisition, including knowledge acquisition techniques, knowledge acquisition methodology, and knowledge acquisition validation and verification methods.

Specifically, the literature will be reviewed for general information on knowledge acquisition, including problems in the elicitation process, theories, comparisons with requirements analysis, elicitation methods and case studies of real-life situations and simulations. In addition:

- Knowledge acquisition techniques will be identified, including the advantages and disadvantages of each. The use of these techniques, singularly or in combination, will be evaluated to determine the best or most appropriate method(s) to use in the prototype.
- Knowledge acquisition methodologies will be identified and evaluated. A model will be developed that will be used in the design and development of the prototype expert system.

- Knowledge acquisition validation and verification methods will be identified and evaluated to determine which methods will be used in the validation and verification of the prototype. Methods for use during design, as well as the final prototype, will be researched.

2. Research in the area of multiple domain experts and identification of associated problems.

Although many existing expert systems have been developed using a single domain expert, the proposed expert system will utilize multiple domain experts. It has been noted that the use of multiple domain experts has presented more complications than experienced with single domain experts. The problems associated with multiple domain experts will be identified and measures taken to avoid or minimize their affects during the application development.

3. Research in the area of disaster recovery planning.

Specifically, analysis of past disasters will address the need for the expert system. Case studies will be reviewed to insure the accuracy of the disaster recovery plan recommended during the consultation process of the expert system.

4. Research in expert systems in general, including development methodology and tools.

Specifically, identification of development methodologies and models recommended in the literature to determine what expert system methodology will be used in the design of the prototype.

5. Identification and selection of domain experts.

Key to the success of the expert system will be the accuracy, content, consistency and completeness of the knowledge acquisition process. Such success will depend on the knowledge engineer's ability to identify all available domain experts and ensure that those selected have the appropriate expertise and knowledge. needed to design a working prototype.

6. Knowledge acquisition using domain experts.

During this step, the knowledge acquisition process will take place as outlined in Table 1.

7. Design of the expert system for disaster recovery planning.

During this step, the expert system will be developed using the expert system shell development tool.

8. Validation and verification of the knowledge acquisition process and the expert system prototype.

During this step, the prototype expert system will be validated and verified using several of the methods described in Table 2. At the completion of the expert system development, attempts will be made to validate the results of the conclusions of the expert system against real conditions. The validation plan will include one or more of the following options: (1) validating against domain experts in real-life situations; (2) validating against simulations; and (3) validating against case studies. Choice will depend on availability of beta test volunteers and available case studies.

Milestones

The literature will be reviewed to determine the appropriate expert system / knowledge based system design methodology. The milestones and estimated time to complete are noted in Table 3.

Resources

Three domain experts and two software tools will be used to design and document the expert system. These resources are described below.

Domain Experts

Domain expert # 1 is a Certified Public Accountant and senior managing director for an independent consulting firm and a former Director in the Information Technology Group of a Big Five accounting firm. As an EDP auditor and consultant, he has extensive experience in the auditing of contingency planning and disaster recovery plans and is presently involved in designing crisis management plans for major organizations. He will serve as the domain expert in the hardware, software, maintenance and recovery team areas.

Domain expert # 2 is an independent consultant specializing in security and crisis management. He was the Vice President for Security for a large retail electronics firm. He is a retired federal law enforcement agent responsible for telecommunications security. He has over twenty-five years experience in telecommunications, including the

design and implementation of backup planning in telecommunications. He will serve as the domain expert for telecommunications and contribute in other areas of his expertise.

Domain expert # 3 is managing director for a firm specializing in computer security and integrity controls. He holds a doctorate in Computer Science and is an expert in computer security, data processing controls, telecommunications and disaster recovery. He has over 25 years experience. He will serve as the domain expert in design of DRP systems, recovery site selection and plan maintenance and will assist in the telecommunications area.

Tools

VP Expert, an expert system development tool, will be used for designing the expert system. ABC Flowcharter will be used to flowchart the design of the expert system and verify and validate the logic of the expert system.

Summary

The design and development of the expert system will utilize methodologies consistent with the design and development of conventional applications, modified for areas specific for expert systems such as knowledge acquisition, verification and validation. Research in disaster recovery planning, expert systems, knowledge acquisition, knowledge acquisition tools, multiple domain experts and validation and verification will be accomplished to provide a basis for the knowledge needed by the knowledge engineer to successfully complete the development of the expert system. The expert system will be PC based and designed utilizing an expert system shell. Three

domain experts will be utilized for the knowledge acquisition and validation phases. The project will consist of four phases: knowledge acquisition, knowledge representation, design and development and validation. Once the design and coding of the expert system is completed, it will be compiled. The validation phase will be performed by the domain experts and two volunteer organizations.

Chapter 4

Discussion of Expectations

Anticipated Benefits & Projected Outcomes

The successful completion of the knowledge acquisition process will result in the creation of a working expert system for designing and developing disaster recovery plans. Accountants, consultants, and others can utilize this tool. The expert system will enable the users to elicit, in a speedy and efficient manner, the information that is needed to submit a timely proposal to the customer. In addition, the use of the explanation function of the expert system shell will provide both the user and a customer with an explanation of why the information is needed, providing a training tool for disaster recovery planning. During the process of the development of the expert system, it is anticipated that knowledge acquisition techniques dealing with multiple domain experts will be explored, confirming past research results and hopefully recommending new techniques that can overcome the many problems raised by research.

Practical Applications of the Findings

The need for disaster recovery plans by organizations dependent on computer services is unquestioned. The development of this expert system will assist accountants, consultants and other users by providing them with a tool to assist in providing these services to their customers. The expert system can be used by individuals who are not experts, filling the void of available experienced DRP consultants, and reducing the costs

and time frame. This tool could assist the organization's accountants and other consultants who were otherwise not qualified to perform disaster recovery consulting. The use of the expert system can now provide customers with quick estimates of what they need and allow the consultants to provide them with the estimated cost.

The outcome of the knowledge acquisition and knowledge engineering processes, and the successful methods for overcoming past problems experienced when using multiple domain experts will allow researchers and designers to develop expert systems in a more efficient and effective matter, with a higher probability of success.

Constraints and Limitations of the Study

The three domain experts utilized for the knowledge acquisition process will be volunteers and will have limited time to provide for this project. Validation of the expert system will be limited to the three domain experts and two volunteer organizations, one representing a \$1 billion company with extensive data processing needs, the other representing a small non-profit organization with a mini computer, multiple server environment.

Recommendations for Additional Studies

Additional testing of the expert system can be performed to validate that it is complete, consistent and provides the information needed to provide the proposal. User friendliness of the system and effectiveness of the explanation function as a training tool can also be validated. The findings in the area of knowledge acquisition can be validated by either duplicating this project or utilizing it in similar studies.

Contributions to the Field of Study and Advancement of Knowledge

The expert system will be developed as a usable tool for accountants and consultants to assist their customers in designing and developing a disaster recovery plan. The tool itself, as well as the effectiveness of the knowledge acquisition process, will provide future researchers with a new basis for additional research and the advancement of the knowledge in the areas of disaster recovery, expert systems development, knowledge acquisition, knowledge engineering, and multiple domain experts.

Table 1. The Knowledge Acquisition Process

1. Select domain experts.
2. Meet with all domain experts as a group to explain the purpose of the expert system, the knowledge acquisition process, constraints, timelines and procedures.
3. Each domain expert develops an initial list of questions in their area of expertise.
4. The knowledge engineer reviews #3 and returns to the domain experts for follow-up action.
5. Each domain expert is asked to write explanations for why their questions were being asked and/or why the information is needed. (This will be used as a validation method and used in the explanation facility of the expert system).
6. Knowledge engineer reviews above and eliminates any duplicate questions.
7. Final questions are modularized by the knowledge engineer and distributed to the domain experts.
8. Knowledge engineer meets with the domain experts as a group to discuss the latest version.
9. Knowledge engineer makes final revisions and flowcharts the expert system.
10. Knowledge engineer meets with domain experts and individually or as a group for final approval of final version of questions, explanations and flowcharts. Revisions made immediately for implementation.

Table 2. Validation & Verification of Knowledge Acquisition & Expert Systems.

TOOL/TECHNIQUE	REASON FOR USE
Explanation Facilities	<ul style="list-style-type: none"> • Used to validate during knowledge acquisition phase the need for question and data requested. • Used to assist client during consultation. • Used to provide client with reason question is being asked. • Can be used as a training tool.
Flowcharts/Decision Trees	<ul style="list-style-type: none"> • Used to verify flows and branching of questions and expert system during consultation.
Decision Tables	<ul style="list-style-type: none"> • Used to validate and verify decision-making based on rules.
Walkthroughs	<ul style="list-style-type: none"> • Used to validate results by comparing decisions making or conclusions by expert systems to those of experts during actual consultation..
Simulations	<ul style="list-style-type: none"> • Used to validate results by comparing decisions making or conclusions by expert systems to those of experts during simulated walkthroughs.
Case Studies	<ul style="list-style-type: none"> • Used to validate results by comparing decisions making or conclusions by expert systems to those of experts using case studies.

Table 3. Milestones

Task	Estimated Time to Complete
1. Selection of domain experts	1-2 weeks
2. Initial briefing to all domain experts	1 day
3. Knowledge acquisition process begins. Each domain expert completes version 1 of general questions & questions specific for their area of expertise.	2 weeks
4. Knowledge engineer (KE) reviews all questions, categories and eliminates duplicates	1 week
5. Domain experts review product in step 4, make comments, add or change questions. (=version 2)	1 week
6. KE reviews all questions, categories and eliminates duplicates for version 2	1 week
7. KE finalizes above & develops flowcharts	1 week
8. Domain experts review final version and flowcharts for signoff before programming	1 week
9. KE codes programs for expert system using shell	4 weeks
10. Expert system is debugged and tested & revisions made	1 week
11. Expert system is tested by domain experts & revisions made	2 weeks
12. Final product approval	1 week


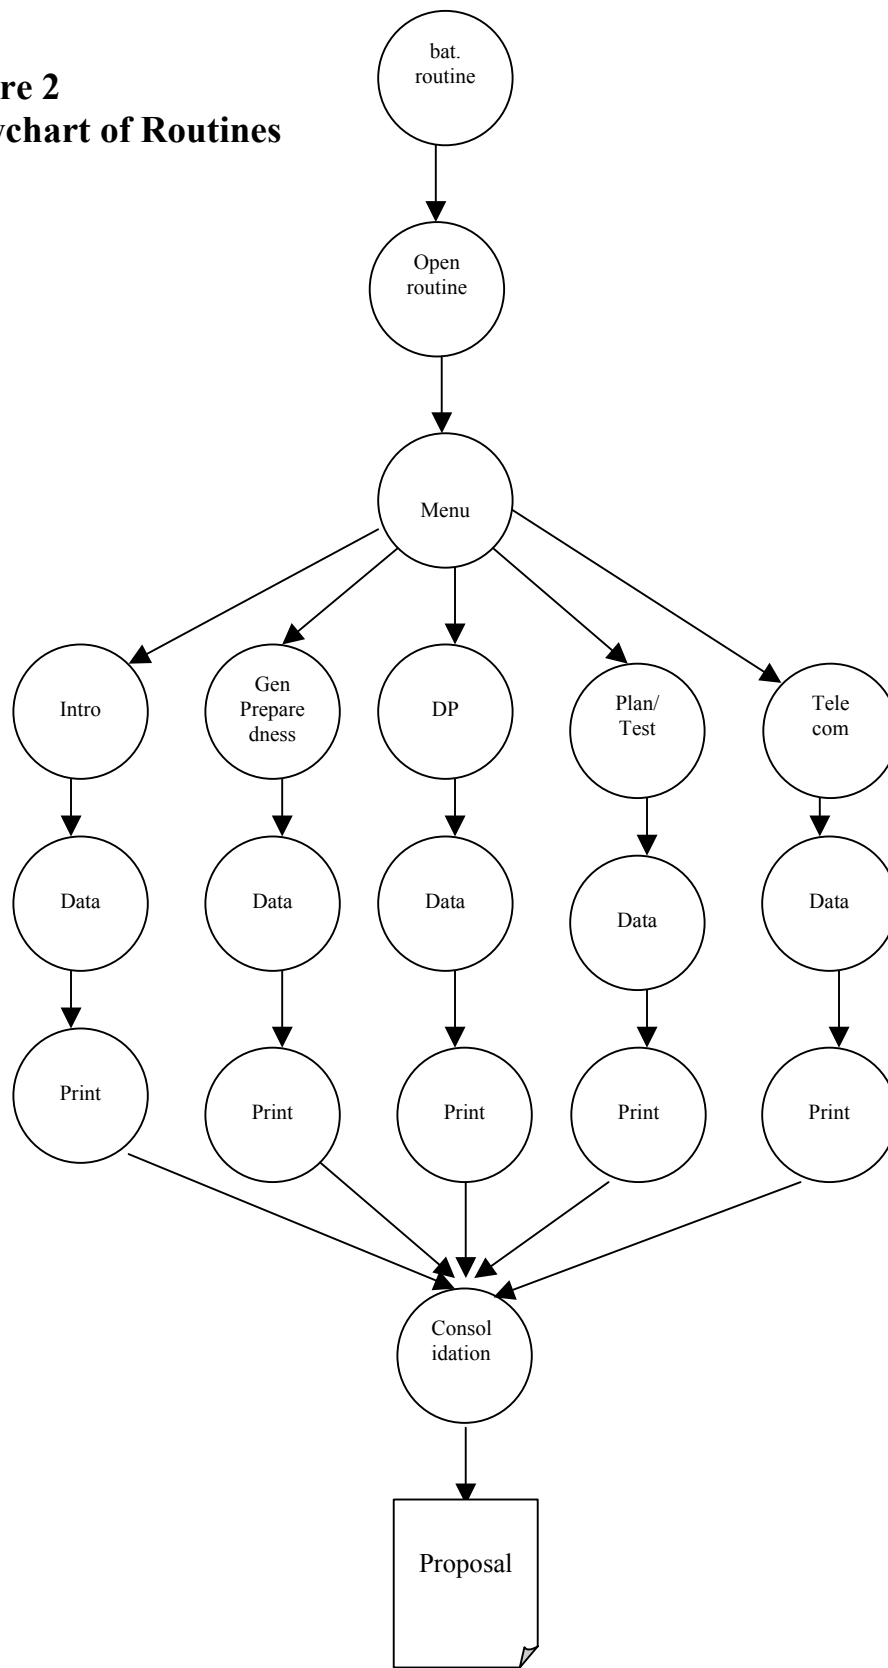
Figure 1**Four Phase Approach to Expert System Design & Development**

Figure 2
Flowchart of Routines



Annotated Bibliography

Aasgaard, D. et al. (1978). An evaluation of data processing machine room low and selected recovery strategies. *Working paper MISRC-79-04, Management Information Systems Research Center, University of Minnesota*, 82 pages.

Study referenced in University of Texas study. One of the earliest studies in DRP, performed by MBA studies. Frequently quoted by both researchers and DRP vendors. Noted that many businesses could not survive without data processing capabilities beyond five days. Specifically, the maximum downtime by industry was 2 days for financial, 3.3 days for distribution, 4.9 days for manufacturing and 5.6 days for insurance.

American Institute of Certified Public Accountants (AICPA). 1988. The auditor's consideration of an entity's ability to continue as a going concern. *Statement on Auditing Standards No. 59*. New York, NY: AICPA

Statement is an auditing standard to be followed by CPA's in auditing the financial statements of an entity. Generally accepted accounting principles are based on the going-concern concept, which means that the entity is expected to continue in operation and meet its obligations as they become due, without substantial disposition of assets outside the ordinary course of business, restructuring debt, externally forced revisions of its operations, etc. Hence an opinion that financial statements are in conformity with GAAP means that continued existence may be presumed for a reasonable time not to exceed one year beyond the date of the financial statements. Today, this going-concern concept is applied to the entity's ability to resume data processing services (disaster recovery planning and business resumption planning), especially in entity's where data processing ability is critical to operations.

Andrews, R. A. (1994). An ounce of prevention: Guidelines for preparing a disaster recovery plan. *Proceedings of the IEEE 1994 National Aerospace and Electronics Conference, 2*, 802-806

The effects of disasters can render an organization helpless. Over 70% of organizations that experience a serious emergency close within two years. Virtually all high-tech organizations are computer-dependent and most cannot perform day-to-day operating activities without their computer systems. In defense software development efforts, the loss of months or years of development work could end a program. The most difficult task of the recovery stage is resuming computer operations. The computer system(s) of an organization must be reconfigured quickly if the organization is to avoid long-term damage. Not only must new hardware and system software be acquired, but also the entire system must be configured to the original "pre-disaster" state. This rebuilding

process includes restoring the required communication links and restoring the application software to each designated machine. This assumes that the systems are backed-up regularly, and the media are stored offsite. The objective of this paper is to heighten the awareness of the need for every business organization, or laboratory project, to have a disaster recovery (or business resumption) plan. Experience has shown that organizations which have contingency plans in place prior to an emergency have a much greater chance for survival than those organizations which only prepare the plan after a disaster has occurred. The focus is on preparing for emergencies involving an organization's computer systems. Steps for putting together a recovery plan, including issues that must be addressed in coordinating such a plan, are discussed

Awad, E. & Lindgren Jr., J. (1992). Skills and personality attributes of the knowledge engineer: an empirical study. *International Association of Knowledge Engineers '92 Proceedings*.

Paper describes a 1992 survey in the US in order to closely scrutinize the skills and personality of the knowledge engineer. The survey was based on a similar one done in the UK in 1994. The results established that certain skills and personality attributes are commonly held to be important for knowledge engineers.

Baldwin-Morgan, A. & Stone, M. (1995). A matrix model of expert system impact. *Expert Systems with Applications*, 9(4), 599-608.

Paper describes the growing use of expert systems for use in accounting and business and attempts to answer the question of what are the impacts of expert systems on organizations that use them. Authors describe a 2-dimensional matrix model of the impact of expert systems on organizations, including the nature of expert systems and the motivations of accounting organizations to develop expert systems, description of the matrix model, application of the model to existing expert systems, and end with strategies for future research. Matrix model identifies and classifies potential impacts at the industry, organization, task and individual user level. The system impacts the efficiency, effectiveness, expertise, education and environment of organizations using them.

Key points of article include:

- Best problem domains are those that are small, important, have human experts and test cases available for development and validation.
- Following conditions will make development of expert systems worthwhile: (1) shortage of human experts, (2) need to preserve experts' expertise, (3) high cost of expert advice or erroneous decisions, (4) critical requirement for expert advice and (5) routine, detail-dependent decision making (Hu, 1987).

Braden, B., Kanter, J. & Kopcsó, D. (1989). Developing an expert systems strategy. *MIS Quarterly*, December, 1989, 459-467.

Authors discuss development of an expert system and conclude that important aspects of the design and development include stressing people skills, needing an expert willing to have his expertise cloned, and having a KE able to translate the expert's experience into decision rules.

Key point: Availability of KEs to design and implement expert systems. KEs must be willing to repeatedly ask the expert to explain what he or she knows, be willing to be continually corrected by the expert, and be capable of eliciting information that is second nature to the expert.

Byrd, T. (1992). Implementation and use of expert systems in organizations: Perceptions of knowledge engineers, *Journal of Management Information Systems*, 8(4), 97-.

Author discusses implementation, use and effects of expert systems. Survey results show that KEs believe that competitive advantage was the most important motivation for using expert systems. Improved productivity was the major benefit. Paper has given evidence that:

- (1) KA is indeed the 'bottleneck' of ES implementation and a separate and distinct process from knowledge engineering. Some difficulties associated with KA may account for this-(a) finding a willing and able expert can be difficult; and (b) How many experts should be used; and
- (2) current ESs are mostly used as decision makers for lesser skilled domain personnel and, to a lesser extent, as advisors to experts (Provides personnel with ways to make more consistent, timely, and accurate decisions in hopes of improving completeness.)

Byrd, T., Cossick, K. & Zmud, R. (1992). A synthesis of research on requirements analysis and knowledge acquisition techniques. *MIS Quarterly*, 16(1), 117-138.

Article compares techniques for requirements analysis used in conventional applications and knowledge acquisition used in expert systems. Authors note that each is the most critical step in their respective systems development process. Article presented an initial categorization scheme to facilitate the merging of research across these two areas and divides the techniques into groups by method of elicitation. Article also demonstrates that techniques in RA and KA share similarities in characteristics and purpose. With expert systems adopting some of the features of information systems, and vice versa, it seems imperative for knowledge engineers and analysts to be aware of a variety of elicitation techniques, regardless of their origin, to be able to develop advanced systems.

Chen, S. (1993). A knowledge acquisition scheme for ruled-based systems. *Proceedings IEEE Region 10 Conference on Computer, Communication, Control and Power Engineering*, 2, 621-625.

Article discusses KA schemes for rule-based systems called fuzzy Petri nets to represent the fuzzy production rules of a rule-based system, where the domain expert or KE can

add new knowledge or update the existing knowledge by inserting new fuzzy production rules into the file or by updating the fuzzy production rules in the file.

Cerullo, M. J., & McDuffie, R. S. (1994). Planning for disaster. *CPA Journal*, 64(6), 34-37

Authors are accounting professors and discuss disaster recovery planning. They note that recent disasters (World Trade Center, Hurricanes Andrew & Hugo, San Francisco earthquake) are calls to industry for the need of DRP. Includes: areas of concern for recovery (facilities, communications); risk analysis; matters to consider (e.g., DR manager, voice & telecommunications, vital records, etc.); benefits. Article concludes with discussion of SAS 60 (Communication of Internal Controls Structure Related Matters in an Audit), which notes that the lack of a DRP may be a reported condition. Key element of article is reference to survey conducted of 71 companies in Charleston, SC after Hurricane Hugo. Results: 41 responded; 44% had DRP before hurricane; 56% did not; of the latter, downtime was 1-15 days for 20 companies, 16-60 days for 2 companies, and 2-4 months for 1 company.

Christensen, S. & Schkade, L. (1987). Financial and functional impacts of computer outages on business. *Working paper CRIS-87-01, Center for Research on Information Systems, University of Texas at Arlington*, 16 pages.

Early study in need for DRP. Survey revealed that 85% of organizations are heavily dependent upon computer systems, that 75% of organizations would have reached critical or total loss of functioning within 2 week of the loss of DP, 63% of respondents had a DRP and only 61% of these were tested, only 50% of the companies that had experienced a computer outage had a DRP.

Cullen, J. & Bryman, A. (1988). The knowledge acquisition bottleneck: time for reassessment? *Expert Systems*, 5(3), 216-225.

Authors question the notion that knowledge acquisition is a major bottleneck in the development of expert systems. Their research included a survey to prove their hypothesis. The survey also provided some excellent data on expert systems, including:

- The majority of expert systems (47%) were developed using shells;
- The majority of expert systems (40%) were developed for use on PCs;
- Major knowledge acquisition strategies included: fast prototyping (38%), evolutionary (25%), ad hoc (18%), expert driven (17%), and machine induction (4%).
- According to evidence, knowledge acquisition actually takes up a relatively small proportion of the development time taken to produce a system;
- Most common knowledge acquisition method is the unstructured interview, followed by documentation;
- 27% used three or more techniques, 24 %=2 techniques.

- They conclude that expert systems development failures can be attributed to three factors: (1) external factors usually associated with organizational problems; (2) inappropriate acquisition strategies and elicitation techniques, which are incompatible with the knowledge structures embedded in the application domain; and (3) the unavailability of appropriate representation formalisms to suit the complexity of the knowledge structures or functions embedded in the application domain.

Durkin, J. (1996). Expert systems: a view of the field. *IEEE Expert*, 56-63.

Author researched current status of the expert systems in answer to reports of decline of expert systems. Survey results indicated impressive growth, as researchers developed systems to tackle difficult but commercially rewarding systems. Survey was based on extensive review of magazine articles, conference proceedings, books and information provided by software vendors. Survey on covered approximately 2500 systems, which the author feels represents only 20 percent of all developed expert systems. Survey results included: (1) expert systems accomplish generic task on the basis of problem types and that many employ more than one activity; (2) majority of expert systems were developed in application serious of business, manufacturing's and medicine. Science was last of 23 areas. (3) majority of expert system application by problem types work diagnoses, interpretation and prescription; (4) more than 60 % are PC based; (5) more than 40% were developed with shells.

Fujihara, H., Simmons, D., Ellis, N. & Shannon, R. (1997). Knowledge conceptualization tool. He *IEEE Transactions on Knowledge and Data Engineering*, 9(2), 209-220.

Paper presents a knowledge conceptualization tool (KCT) in which the knowledge engineer can effectively structure and formalize knowledge components, so that the resulting knowledge base is accurate and complete. The KCT uses information retrieval techniques to facilitate conceptualization, which is one of the human intensive activities of knowledge acquisition. Their analysis of the results of the use of KCT showed that the process time to retrieve components (e.g. facts, rules, protocols and uncertainty) is about half that of the manual process, and the number of knowledge components retrieved from knowledge acquisition is four times more than that retrieved through a manual process. Authors note use and limitations of interviewing and suggest KCT as an alternative semiautomatic tool.

Garcia-Molina, H. & Polyzois, C. (1990). Issues in disaster recovery. *COMPCON Spring '90. Intellectual Leverage. Digest of Papers. Thirty-fifth IEEE Computer Society International Conference*, 573-77.

Articles discusses several issues in disaster recovery including motivation for remote backups and their applicability to various systems, issues involved in the design,

implementation, and evaluation of remote backup mechanisms, and some criteria for evaluating and comparing various remote backup mechanisms.

Authors define a disaster as total loss of processing power and access to data at one site for a period of time and possible permanent loss of data at that site. Disasters include natural (earthquakes, power outages), malicious acts, hardware failure, operator errors, etc. They noted that a remote backup can ensure continuous operation even in the presence of extensive failures that may render an entire site non-operational for which local replication may be inadequate.

Also noted that some federal regulations call for the existence or development of a disaster recovery plan (e.g. banks).

Gorney, D. & Coleman, K. (1991). Expert systems development standards. *Expert Systems with Applications*, 2, 239-243.

Paper discusses current state of expert system development standards. Authors note that many organizations do not follow any development standards for expert systems and that this failure is cited as one of the reasons for expert system failure. Reasons commonly cited for failure also include poor idea selection, improper planning, inadequate funding, overly optimistic scheduling and technology problems. Cited are various reasons for the lack of standards, including, (1) corporate beliefs that we know all there is... so why follow someone else's standards, (2) cultural differences around the world, and (3) the rapid advancement of the technology and the varying levels of technology. Authors also describe KBS development methodology used at Deloitte & Touche (Big 5 CPA firm).

Griswold, J. S., Lightle, T. L., & Loverlady, J. G. (1990). Hurricane Hugo: effect on state government communications. *IEEE Communications*, 28(6), 12-17

Article discusses the impact of Hugo on state government communications in South Carolina and activities in the wake of the disaster. The state of South Carolina had been chosen, along with California, as a model state for development of the National Communications System/Telecommunications Services Priority (NCS/TSP) model plan for emergency restoration of telecommunications services. This effort is being undertaken in conjunction with the National Communications System Organization in Washington, DC. They are assisting the state in development of the telecommunications services priority (TSP) plan, to be used in the future for coordinating restoration of critical federal, state, and local telecommunication services. A private sector partnership has been initiated with the telecommunications and utility companies in South Carolina to provide fiber-optic alternate routing and other communications backup capability to each of the emergency preparedness offices in South Carolina counties in future disasters. An exchange of disaster recovery plans with other states has been undertaken. Information regarding each of these efforts is provided.

Hayes-Roth, F. & Jacobstein, N. (1994). The state of knowledge-based systems. *Communications of the ACM*, 37(3), 26-39.

Authors have produced an excellent survey of the technology, successes, and failures of knowledge base systems. Based on their experience and the survey, they discuss the essential value of knowledge processing; industrial and commercial applications, with real world examples; key trends in applications of KBS technology and experience; perceived failures, difficulties, and disappointments; and the state of the act technology components. The main conclusions of the paper include that KBS has attained a permanent and secure role in industry; KBSs are integrated with other information technology; and KBS has a rapid and wide impact and diffusion in the world economy.

Hiles, A. (1992). Surviving a computer disaster. *Engineering Management Journal*, 2(6), 271-274.

Most organizations are so reliant on their computing service that loss of real-time systems could have devastating effects. The first step to improving this situation is to identify what risks exist and how significant they are within the specific environment of the organization. Disaster recovery plans can then be formulated in order to restore permanent computing service.

Disaster is defined as an event, which causes the loss of the computing service, or of a significant part of it, or of systems, communications or applications, for a length of time which prevents the impacted organization from achieving its mission or which imperils the business. Risk include fire, flood, lightning strikes, malicious damage, power failures, impact, subsidence, loss of supplied services or special consumables, rodent damage, insect infestation, contamination, telecommunications failure, industrial action,.... A risk review will identify the key threats and includes a critical component failure analysis and business interruption analysis. Elements of a DRP include: immediate restoration procedures, restoration of computing infrastructure and of applications, resumption of business processing under emergency arrangement, and restoration of the permanent computing service.

Hoffman, R. (1987). The problem of extracting knowledge of experts from the perspective of experimental psychology. *AI Magazine*, 8, 53-64

Author discusses his research on generating methods for extracting the knowledge of experts from the experimental psychology side of expert systems. He notes, as others have, that knowledge acquisition is a significant bottleneck in the system development process. His research offers: (1) a working classification of methods for extracting knowledge (table 1); (2) some ideas about the types of data that the methods yield; (3) a set of criteria by which the methods can be compared relative to the need of the system development (table 5); (4) some salient advantages and disadvantages of the various

methods table 4); (5) comparison of the results of four of the methods in terms of efficiency (table 7); and (6) some steps for extracting and characterizing knowledge prior to the construction of an expert system (table 8). The discussion highlights certain issues, including the contrast between the empirical approach taken by experimental psychologists and the formalism-oriented approach that is generally taken by cognitive scientists.

Hwang, G. (1994). Knowledge elicitation and integration from multiple experts. *Journal of Information Science and Engineering*, 10, 99-109.

Article discusses knowledge acquisition from multiple experts. Author notes that knowledge acquisition is known to be a critical bottleneck for building expert systems and that this is complicated when multiple experts are needed due to conflicts and inconsistencies. The paper discusses the use of a KA method called MERGE (multiple expert repertory grid elicitation) to cope with these problems by integration of the expertise.

Jacobs, J., & Weiner, S. (1997). The CPA's role in disaster recovery planning. *CPA Journal*, 67(11), 20-25

Article discusses opportunities for consulting services by CPAs. Notes SAS 60 as justification for bringing DR issues into the audit process. Includes: Types of companies that used DRP; Disaster planning alternatives; Steps in the formation and implementation of DRP (11 steps); Benefits; Overcoming plan preparation difficulties (“some of the problems encountered in completing a client’s DRP relate to the difficulty of gathering the essential information to be detailed in the body of the plan”). Key point: 1993 industry study commissioned by then DEC determined that “90% of companies that experience a catastrophic loss of data and equipment and do not have a DRP are out of business within 2 years.” Some of the problems encountered in completing a client’s DRP relate to the difficulty of gathering the essential information to be detailed in the body of the plan.

Jeng, B., Lieng, T. -P. & Hong, M. (1996). Interactive induction of expert knowledge. *Expert Systems with Applications*, 10(3), 393-401.

Article discusses knowledge acquisition. Authors note that KA is a bottleneck in developing KBS. They note that it is a manual approach that elicits domain knowledge by interviewing human experts and that this process has problems because the experts are often unable to articulate their reasoning. They also note that attempted automated approaches that induces knowledge from a set of training cases also suffers from the unavailability of sufficient case studies. The authors present an integrated approach that combines the strengths of both methods to compensate for their weaknesses.

Authors note that there are many different methods for KA, which can be classified as manual, semi-automatic and automatic. Manual methods are highly labor intensive and include structured/unstructured interviews, analysis of protocol, observations, etc.

Induction is a process of general inference from particular instances, whereas, rule induction refers to a concept learning process by which a set of rules is created from training cases or case studies to explain or to predict a problem solving behavior. When the induced structure of knowledge is represented in the form of a decision tree, it is also called (decision) tree induction. A decision tree is considered as a set of rules in a compact form, which can easily be transformed into rules.

Kemp, E. & Purvis, M. (1996). The role of the individual project in teaching knowledge acquisition. *Proceedings of the International Conference in Software Engineering: Education and Practice*, 138-143.

Knowledge acquisition is that stage of knowledge based systems development, which is equivalent to the analysis and design phases of the conventional software life cycle. An account is given of the experience gained when teaching a postgraduate course on this topic. One of the aims of the course, "Topics in Knowledge Acquisition", is to provide students with the opportunity to discover for themselves what this stage of the knowledge based system development life cycle involves. Eliciting, analyzing and modeling domain knowledge, the main activities of the knowledge acquisition process, are fraught with problems. Each student undertakes the task of developing a small expert system in an appropriate domain. The emphasis in this project is not on the implementation but on the processes the students follow. Students are asked to evaluate the elicitation, analysis and modeling techniques they use. The paper describes the information available to the students before they commence the project, gives an account of their experiences during the knowledge acquisition stage, discusses their findings and concludes with lessons learned for the future.

Keyes, J. (1989). Why expert systems fail. *AI Expert*, 4(11), 50-53.

Author states that only 10% of medium to large sized expert systems are successful. Article discusses reasons for failure, including: (1) lack of an available and willing expert, (2) lack of management support, (3) lack of user support, (4) lack of user involvement during early development, (5) failure of KE to capture all or the right knowledge, and (6) inability to fulfill operational requirements. Author also states that the only difference between expert system development methodologies and traditional DP methodologies is the capture of knowledge which she termed "the essence of the system" and that knowledge capture is a "really fuzzy" process that can be performed in many ways and that no single way is best.

LaSalle, A. & Medsker, L. (1991). Computerized conferencing for knowledge

acquisition from multiple experts. *Expert Systems with Applications*, 3, 517-522.

Article discusses the motivations for acquiring the knowledge of multiple experts, the problems related to knowledge acquisition, new issues that arise when multiple experts interact, solutions that can be brought to bear in building multiple domain expert systems (particularly when experts are geographically dispersed), and new tools for knowledge engineers to use when dealing with multiple experts. Authors list some of the problems: (1) ambiguous knowledge, (2) distributed knowledge, (3) disjoint knowledge, (4) critical knowledge, (5) adversarial knowledge, and (6) synergistic knowledge. They noted that in trying to gather knowledge from multiple experts who are geographically dispersed or who cannot meet in real time, new issues arise, including (1) What is the role of the KE (new skills, facilitator, etc.), (2) What consensus mechanisms should be used (voting, judgment of KE), (3) How can conflict among experts be resolved, (4) What new tools will facilitate KA from multiple experts, (5) What feedback mechanisms are appropriate or necessary (rapid prototyping), and (6) How can performance of the integrated knowledge domain expert system be evaluated.

Lichte, K. (1993). Knowledge capture model for expert systems. Proceedings First New Zealand International Two-Stream Conference on Artificial Neural Networks and Expert Systems, 163-164.

The development of expert systems requires the efficient capturing and organizing of diverse knowledge in a form suitable for encoding in an expert system shell. Structuring of the development can be achieved using models such as: (1) models of knowledge analysis to assist system developers in the task of capturing knowledge and designing user interfaces; (2) knowledge capture models to provide developers with a formal method of collecting and organizing data in a form suitable for direct coding in an expert system. Authors present these two models.

KADS is the name of a methodology for the development of knowledge-based expert systems. The process of capturing knowledge is defined as the collection, organization, evaluation, and incorporation of knowledge within a working expert system.

Liebowitz, J. (1993). Educating knowledge engineers on knowledge acquisition. *IEEE International Conference on Developing and Managing Intelligent Systems Projects*, 110 -117.

Article notes that although many universities are offering programs in knowledge engineering (?), it is rare that a course in knowledge acquisition is included. Paper discusses the need for courses in KA that will contribute to improving the state-of-the-art in expert system development as well as a general background in KA.

Authors note that KA bottleneck has been well documented; that one day of the expert's time is usually needed for every four days of the KE's time; because of the difficulty in acquiring and eliciting knowledge, potential KE's need a stronger foundation in the

fundamentals, methodologies, techniques, and tools for acquiring knowledge; that George Washington University is one of the few schools that do offer such a course. They noted that the KA process includes: (1) discovery of prelim knowledge and problem range; (2) identifying sources of information; (3) acquiring detailed knowledge from sources; & (4) analyzing, coding and documenting knowledge. They quote Cullen for KA problems: i.e. quality of expertise, communication, etc. They also quote that a major reason why KA is a difficult process relates to the Paradox of Expertise which states that the more competent domain experts become, the less able they are to describe the knowledge they use to solve problems. Specifically: experts have a tendency to state their conclusion and the reasoning behind them in general terms that are too broad for effective machine analysis; the pieces of basic knowledge are assumed and are combined so quickly that it is difficult for the expert to describe the process; that in explaining his conclusion or hunch, the expert will repeat only the major steps, leaving out most of the smaller ones (note: this is a concern in my research). They also quote for other problems such as: Experts may lack time or may be unwilling to cooperate; Methods for KA many be poorly defined; Use of one source when relevant knowledge is scattered across several sources; Builders may attempt to collect documented knowledge rather than use experts; difficulty in recognizing specific knowledge from irrelevant data; Experts may change their behavior when they are observed or interviewed; and Interpersonal communication factors between KE and expert.

Liou, Y.I. (1999). Expert system technology: knowledge acquisition. In J. Liebowitz Ed.), *The Handbook of Applied Expert Systems*. (pp. 2.1 – 2.11). New York: CRC Press.

Article presents chapter on knowledge acquisition. Author discusses selection of experts based on (1) domain expertise, experience and reputation (2) personal characteristics & attributes, & (3) availability. He references McGraw's book for problems when using one expert; references himself for problems using multiple experts; and references himself for the positive effects of using multiple experts.

Problems when using one expert include: (1) difficulty in allocating time by a key individual in the organization, that may create a bottleneck; (2) personal bias that may affect the performance of the system; and (3) limitation to a single line of reasoning that may affect the usefulness of a system since expert systems that are developed based on a single, perhaps narrow line of reasoning, do not emulate most real-life decision making; and (4) incomplete domain expertise that may affect the performance of a system.

Problems with multiple experts: (1) difficulty of merging each individuals expert's knowledge into one group structure that provides the underlying problem solving expertise of the ES (i.e., conflict resolution); and (2) difficulty of the generation of group knowledge that does not reside in any one individual expert but resolves as a result of group interaction. I will note that a solution to this is a strong KE with strong knowledge in the domain.

Positive effects of multiple effects include: (1) it assures that the KB can be complete; (2) improves the likelihood of obtaining specialized knowledge in subdomains of the problem (important for my research); (3) increases the quality of the acquired knowledge; (4) assures that the facts that are included in the KB are important ones; (5) enhances understanding of the domain knowledge through discussion, debate and exchange of the hypothesis between members of the expert team; and (6) encourages interactions among experts and creates a synergy such that the acquired group knowledge is sum of the individual's knowledge.

McGraw, K. L., & Harbison-Briggs, K. (1989). *Knowledge acquisition: Principles and guidelines*. Englewood Cliffs, NJ: Prentice-Hall.

Textbook written for novice knowledge engineers. Presents a practical approach to the KA process. KA methodologies (Chapter 1-3) and techniques (Chapter 4-11) are detailed, including the benefits of working with multiple domain experts and the most effective means for resolving problems that may arise. Authors note that KA is the most critical element in the development of an expert system and that it is both people and time intensive (p.2). They state that most of the literature related to ES development alludes to the troublesome nature of KA, even when only a single domain expert has contributed to the knowledge base. The already difficult KA process becomes even more complicated when the ES being developed requires that knowledge engineers interact with multiple experts (p. 242). Complex, real life problems are seldom so simple that they can be solved based on consultations with a single expert. More often than not, a complex problem requires access to different types of knowledge and thus, to more than one expert (p. 243).

Motoda, H., Mizoguchi, R., Boose, J. & Gaines, B. (1991). Knowledge acquisition for knowledge-based systems. *Paper presented at the IEEE Expert*.

The authors' work reported at the first Japanese Knowledge Acquisition for Knowledge-based Systems Workshop is discussed, providing both an overview of the field and an introduction to a series of articles on knowledge acquisition. The discussion covers tools, methods, and mediating representations; real-time problem solving, the system-model-operator metaphor; an interview architecture based on dynamic analysis, inductive knowledge acquisition from structured data; research in Japan; how to make application programming easier; justification-based knowledge acquisition; integrating knowledge acquisition and performance systems; tasks, methods, and knowledge; rule induction, hypertext' explanation-based learning and case-based reasoning; and interviewing.

Money, W. & Harrald, J. (1995). The identification of group support systems to knowledge acquisition for disaster recovery planning. *Proceedings of the Twenty-eighth Hawaii International Conference on System Sciences*, 4, 468-474.

This paper describes the results of the use of a Group Support System (GSS) and two

"hindcasting" scenario exercises. The GSS is used to implement the hindcasting, knowledge mapping, and knowledge representation techniques as disaster planning tools. This work improves our understanding of the knowledge engineering process in the areas of knowledge acquisition, validation, and representation when the knowledge must be collected from the multiple experts, and transferred to a knowledge base. The experiences from the exercises illustrate how the collection of the knowledge of experts may be supported by a GSS when the experts perceive their task to be one of collaboratively deducing a rational causal explanation for a complex event. The results indicate that managers may improve their ability to integrate multiple experts' knowledge into the disaster preparedness planning process using GSS processing techniques.

Motoda, Mizoguchi, Boose & Gaines, 1991

Nguyen, T., Perkins, W., Laffey, T. & Pecora, D. (1987). Knowledge base verification. In J. S. Chandler & T.-P. Liang (Eds.), *Developing Expert Systems for Business Applications*. (pp. 69-77). Columbus, OH: Merrill Publishing.

Article describes CHECK, a program whose function is to check a knowledge base for consistency and completeness, i.e., knowledge base verification. Verification is defined as a technique or methodology for testing the consistency and completeness of a rule set. The program detects several potential problems in verification, including redundant rules, conflicting rules, subsumed rules, unnecessary IF conditions, and circular rules. CHECK also attempts to verify completeness in the knowledge base by looking for potential gaps, including unreferenced attributed values, illegal attribute values, missing rules, unreachable conclusions, and deal-end goals. Authors in their research note that many changes and additions to the rule sets occur during the development of a knowledge base.

Niederman, F. (1996). Acquiring knowledge about group facilitation: research propositions. *Proceedings of the 1996 conference on ACM SIGCPR/SIGMIS*. 58-67.

This paper argues for the desirability of investigating the expertise that facilitators use to lead groups. It reviews the importance and difficulties in studying meetings. It reviews the nature of knowledge acquisition techniques and suggests a knowledge acquisition model for developing a map of the expertise in the domain of group facilitation. Sets of propositions regarding (1) the domain of knowledge for expert facilitators and (2) differentiation of facilitators by level of expertise are presented. Finally, potential benefits of further investigation of facilitator expertise are discussed.

Nuseibeh, B., Kramer, J., & Finkelstein, A. (1994). A framework for expressing the relationship between multiple views in requirements specifications. *IEEE Transactions on Software Engineering*, 20(10), 760-773

Article discusses multiple viewpoints or perspectives in requirements specifications.

Authors note that even in a single activity such as requirements elicitation, it is likely to involve multiple development participants who will hold multiple perspectives of a single domain. Multiple perspectives may overlap, complement or contradict each other. In their use of ViewPoints, they note that in attempting to integrate multiple requirements specifications, overlaps must be identified and expressed, complementary participants made to interact, and cooperate, and contradictions resolved.

O'Keefe, R., Balci, O, & Smith, E. (1990). Validating expert system performance. In J. S. Chandler & T.-P. Liang (Eds.), *Developing Expert Systems for Business Applications*. (pp. 91-102). Columbus, OH: Merrill Publishing.

Article seeks to establish validation as an important concern in expert system research and development. It discusses problems in expert system validation and presents the qualitative and quantitative methods for validating expert systems. It was noted that engineers typically have validated expert system performance by running test cases through a system and comparing results (i.e., the classification, final certainty factors, and advice given) against known results or expert opinion. The authors note problems with this process. Validation is often confused with verification. Simply stated, validation refers to building the right system (i.e., substantiating that a system performs with an acceptable level of accuracy) whereas verification refers to building the system "right"; (i.e., substantiating that a system correctly implements its specifications).

O'Keefe, R. & Lee, S. (1990). An integrated model of expert system verification and validation. *Expert Systems with Applications*, 1(3), 231-236.

Paper presents a development model, based on Boehm's spiral model that integrates verification and validation into development stages. The model is loose and adaptable and can easily be adopted in whole or part as a development strategy for small expert systems. Model present v & v techniques throughout the SDLC. Key points noted:

- KA is defined as the process of gathering domain knowledge, typically from experts and transforming it to a certain knowledge representation scheme. Verification in KA can be considered a major task in the iterative development of expert systems.
- Boose and Gaines suggest that KE's prepare diagrams representing a summary of the problem solution process and give it to the experts for verifications.
- Several methods of v & v: (1) mechanical verification, (2) case testing (historic cases are run through the expert system and the expert, and the results are compared), (3) field testing (prototype is implemented for one or more users).

O'Leary, D. (1990). Validation of expert systems with applications to auditing and accounting expert systems. In J. S. Chandler & T.-P. Liang (Eds.), *Developing Expert Systems for Business Applications*. (pp. 78-90). Columbus, OH: Merrill

Publishing.

Paper presents a theory-based framework that is useful for guiding the validation of an expert system and also for eliciting other validation issues. The framework addresses validation in terms of validity, objectivity, cost-benefits and accuracy. Validation is important to the decision-making success of an expert system and to the continued use of the expert system. An expert system that does not make the correct decisions will lead to a loss of confidence in it and non-use. Therefore, expert systems must be validated prior to implementation in the field for real like use. Author also noted that validation requires: (1) ascertaining what the system knows, does not know, or knows incorrectly; (2) ascertaining the level of expertise of the system; (3) determining if the system is based on a theory for decision-making in the particular domain; and (4) determining the reliability of the system.

O'Neil, M. & Morris, A. (1989). Expert systems in the United Kingdom and evaluation of development methodologies. *Expert Systems*, 6, 90-99.

Article discusses survey performed in late 1987, early 1988, of expert system developers which was aimed to establish the nature and scope of expert system projects in Britain, the skill set and methodologies used in their development and the background and experience of the personnel employed to develop them. Major part of survey focused on the human input and organizational impact of the expert systems. It includes details of the organizational background of those working in the area, review of the nature, number and scope of expert system projects, the knowledge engineering task and methodologies used, the skills and attributes of the KE, and issues of updating and extending expert systems.

Noted in survey were:

Only 22% of expert systems were designed using multiple experts. Overriding reason for the use of a single expert seemed to be that it was easier both to elicit the knowledge and to avoid contentious issues and conflicting opinions. Several developers had tried to use multiple experts but had faced problems of conflicting information, team disagreements and individual experts trying to assert their own authority over the group. However, they noted that in cases where it was impossible to obtain all the information needed from one expert, it was necessary to use multiple experts and that the KE needed to take steps to avoid conflict or dissent by: (1) asking experts to provide documentary backup evidence; (2) the use of probabilities and uncertainty factors to express degrees of agreement or disagreement amongst experts; and (3) creating systems modularly so that different experts could be specifically called on for particular sections and consultation. I will take these recommendations into my KA/expert system design approach by requiring all experts to state the reasons they need the information they are asking for (which I will also use as a validation technique and for the explanation facility of VP-Expert). In addition I will design the ES modularly and have each modular done by a specific domain expert, although all will eventually review the others to ensure all matters are covered.

Olson, J. & Rueter, H. (1987). Extracting expertise from experts: methods for knowledge acquisition. *Expert Systems*, 4(3), 152-168.

Authors note that knowledge acquisition is the biggest bottleneck in the development of expert systems and that the process of translating expert knowledge to a form suitable for expert system development can benefit from methods developed by cognitive science to reveal human knowledge structures. There are two classes of these investigative methods: direct and indirect. Authors provide reviews, criteria for use, and literature sources for all principal methods. Direct methods discussed include interviews, questionnaires, observation of past performance, protocol analysis, interruption analysis, closed his all curves, and inferential flow analysis. Indirect methods include multidimensional scaling, hierarchical clustering, general weighted networks, ordered trees, and repertory grid analysis.

Rees, P. L. (1996). User participation in expert systems. *Industrial Management & Data Systems*, 93(6), 3-7

Article discusses the need for user involvement in the design and implementation of expert systems and emphasizes that the lack of such is a reason for failure of expert systems. The article first outlines some of the work that has been carried out in the area of participative design of new technology. It then goes on to describe three expert system projects and uses them to illustrate the difficulty of participative design and the dangers of not undertaking it. The article concludes that, just because participation is difficult, this is insufficient justification for rejecting it as an important goal.

Author notes DEC's XSEL as one of the most successful use of participative design and discusses Mumford's methodology for participative design, ETHICS (Effective Technical and Human Implementation of Computer-based Systems), which provides practical guidance for user participation.

Rook, F. & Croghan, J. (1989). The knowledge acquisition activity matrix: a systems engineering conceptual framework. *IEEE Transaction on Systems, Man & Cybernetics*, 19(3), 586-597.

Paper presents a systems engineering conceptual framework that provides a structure for the specification of KA activities throughout the SDLC. Objectives include: (1) review general KA requirements for the development of applied knowledge-based systems; (2) present a SE conceptual framework in the form of the KA acquisition activity matrix (KAAM), for structuring KA activities; and (3) discuss the specific goals, constraints, & output of KA steps in the context of the SDLC. Authors also note the KA has often been described as the bottleneck in the KBS development and that most responses to countering this bottleneck phenomenon focus on specific knowledge elicitation methods

and techniques. Authors suggest a SE conceptual framework (KAAM) to address the bottleneck.

Rudolph, C. Business continuation planning / disaster recovery: a marketing perspective. *IEEE Communications Magazine*, 28(6), 25-28.

Article discusses requirements for restoring telecommunication services to businesses following a disaster such as the central office fire in the Chicago suburb of Hinsdale in 1988 and several California earthquakes. Also discussed was US West's disaster recovery strategies and disaster recovery marketing tool (called Preparedness Audit). An illustrative case study is presented. Author noted that evidence strongly supports the idea that companies can survive disaster if they: (1) plan for the possibility of a disaster; (2) formulate strategies for recovering critical business functions; (3) implement technologies to aid the recovery of automated functions and systems; and (4) train employees to implement those strategies.

Survey conducted at the University of Minnesota revealed that many businesses cannot survive with data processing capabilities beyond 5 days. Maximum Time To Belly Up (MTTBU)=term is being used in the industry to denote the time frames that contingency planning coordinators have to restore not only the DP applications, but the associated telecommunications networks as well. Reality dictates that proactive, preplanned solutions for restoring communications capability be a part of the contingency planning process for companies which rely on telecommunications for critical business functions.

Sangster, A. (1996). Expert system diffusion among management accountants: a U.K. perspective. *Journal of Management Accounting Research*, 8, 171-182.

Author reports findings of mail survey of expert system diffusion among over 4000 UK based professionally designated management accounts. That revealed in level of awareness of expert systems, a very low level of the diffusion of the technology, and that management accounts will tend to resist change. Author notes: Expert systems more than any other IT tool, have the potential to help management accountants develop and expand their roles. While experts systems have been very successfully adopted throughout the industry and commerce permeate area of development has been in accounting and finance related tasks. Two key benefits of expert systems include: (1) to system and permanently reproducible performance, and (2) in capability provide expert level support, even in situations where no human expert is present. By 1995 more than 70 percent of the top 500 U.S. companies were using some form of AI in their operations (Business Week, 1995). Accounting related applications were the most frequent users of expert systems (Hayes-Roth/Jacobstein, 1994)

Smith, P. (1996). *An introduction to knowledge engineering*. London: International Thompson Computer Press.

Textbook is about knowledge engineering and development of knowledge base systems. Chapter 1 provides an introduction to the subject. Chapter 2 focuses upon the subject while Chapter 3 explains the life cycle. Chapter 4 provides detailed coverage of knowledge acquisition (bottleneck). Chapter 5 covers knowledge representation, while Chapter 6 presents knowledge base systems implementation in a programming language and two expert system shells. The last chapter (8) provides coverage of management issues in knowledge acquisition.

Smith, P., Ross, P., Awad, E., Green, C. & MacIntyre, J. (1994). A survey of the skills and personality attributes of the knowledge engineer in the United Kingdom.

Paper describes a 1994 survey in the UK in order to closely scrutinize the skills and personality of the knowledge engineer. The survey was based on a similar one done in the US. The UK results confirmed a number of findings of the US survey and established that certain skills and personality attributes are commonly held to be important for knowledge engineers in both countries.

Song, I. -Y., Strum, S. & Medsker, C. (1991). Design and implementation of a database design aid using VP-Expert. *Proceedings of the IEEE/ACM International Conference on Developing and Managing Expert System Programs*, 1991, 15-23.

Article discusses the design and implementation of a database modeling aid called Database Designer using VP-Expert. Authors discuss their approach to system design and implementation issues specific to the expert system design shell, VP-Expert, such as combining backward chaining and forward chaining (this combination will be used in my research project), procedural techniques, and transparency and visibility of VP-Expert statements. Authors note that VP-Expert allows one dimensional arrays and forward chaining; that an ES for design applications should utilize forward chaining, since its desired goals are not known in advance or are too numerous to list; that forward chaining is most useful in creating data driven processes and ES's with unknown goals.

Starr, M. (1997). The state of disaster recovery planning in Texas small businesses. Unpublished doctoral dissertation. Nova Southeastern University, Florida.

Purpose of the dissertation was to determine the DRP practices of small businesses in Texas. Study surveyed existence of DR plans, systems included in plans, DR prevention and recovery tools used, incidence of testing, and geographic and economic factors affecting DRPs. Results noted that only 34% of the surveyed companies have a DRP and that only 44% of those have been tested

Stein, E. W. (1993). A method to identify candidates for knowledge acquisition. *Journal of Management Information Systems*, 9(2), 161-178

Article talks about network analysis as a method for selecting candidates for knowledge acquisition. Network analysis is based on the measurement of interactions or relations between individuals. Author noted that network analysis can help the knowledge engineer and managers identify experts who can be identified based on their centrality as information providers within the organization.

Network analysis seems to be time consuming and looks broadly at organization's members to determine who has knowledge. However, validation seems to be based on what other members of the organization think about an individual. Question=do they really have the expertise? For example, an accountant may know the most about accounts receivable in his company, but is he an expert in accounts receivable? These are really two different questions. Note: Network analysis may work if the knowledge acquisition is specific to the organization itself.

Stylianou, A., Smith, G., & Madey, R. (1993). Selection criteria for expert system shells. A sociotechnical framework. *Communications of the ACM*, 35(10), 30-48.

The evaluation and selection of a specific shell are important parts of expert system development projects. Use of the wrong shell could lead to major compromised or project failure. It is difficult to sort through the hype and evaluate many shells and vendor claims; therefore many evaluate and select arbitrarily and unsystematically. The objective of the article was to provide a systematic process for the selection and evaluation.

Authors note benefits of expert systems including: cost reduction, increased output, improved quality, consistency of employee output, reduced overtime, captured scarce expertise, flexibility in providing service, easier operation of equipment, increased reliability, faster response, ability to work with incomplete and uncertain information, improved training and increased ability to solve complex problems.

The shell evaluation and selection process include: (1) identification of required capabilities and features; (2) identification of potential shells; (3) evaluation of shells based on required capabilities and features; and (4) selection of the appropriate shell. Based on their survey, they determined the shell capabilities considered to be the most important are: embeddability (refers to the ability of ES shells to be built into conventional applications thereby providing these applications with the advantages of a KB system), rapid prototyping, backward chaining, explanation facility, ability to customize explanations, linkage to databases, and documentation comprehensiveness/readability.

Stylianou, A., Smith, G., & Madey, R. (1995). An empirical model for the evaluation and selection of expert system shells. *Expert systems with applications*, 8(1), 143-155.

Authors develop an empirical model for selecting and evaluating expert system shells. The model was based on a survey of 271 knowledge engineers and end-user and can be used to identify critical ES shell attributes and capabilities, which can then be used as evaluation criteria. On page 153, authors show a diagram of the expert system shell evaluation process (model) they recommend. Authors define ES shells as software packages containing a generic inference engine, a user interface and a collection of other tools that enable users to develop and use expert systems. Using the shell's tools, a KE can develop new KBs and can structure, add, delete, and modify the knowledge contained in them.

There are more than 100 commercial ES shells available in the market (note: this has decreased tremendously in 1999 since interest in expert systems is fading). The selection of a specific ES shell for a particular application is an important decision. If the wrong shell is selected it could result in an inefficient or ineffective system or even in project failure. The difficulties encountered during the process of selection and evaluation are not unlike those for other software packages. The selection depends on factors such as: application characteristics (e.g. problem domain); project characteristics (developer's sophistication with AI techniques); capabilities of shell, sophistication of users; ease of integration with existing software and hardware; and vendor support.

Tsai, N, Necco, C. & Wei, G. (1994). An assessment of current expert systems: Are your expectations realistic? *Journal of Systems Management*, November 1994, 28-32.

Article presents a comprehensive view of current expert systems in order to assess the technology more objectively. Authors conducted a survey of 500 companies with DP budgets greater than \$250,000. Major conclusions reached:

- (1) ...expense & risk continue to be the major obstacles to the use of ES technology.
- (2) Lack of connectivity and integration with existing data processing systems (47%), scarcity of experienced knowledge engineers (40%), human resistance to change (43%), and lack of top management support are major problems being experienced currently in implementing ES technology.
- (3) The most difficult task were determined to be KA since it involves communication between people with completely different backgrounds, human experts and KES who must formulate the concepts, relations and control mechanisms needed for the ES, and verification & validation.
- (4) Cited weaknesses of user's current approaches: Picking appropriate domain experts (10%) and choosing experienced KES (6%).
- (5) Major reasons for not using ESs: Hard to find proper KEs (20%) and hard to find adequate domain experts (16%).
- (6) Major reasons for using ESs: #1=to enforce consistency; #2=enhance product/service quality; #6=to train less experienced employees; #8=reduce costs; and #9=monitor vast amounts of information.

Turban, E. & Aronson, J. (1998). *Decision support systems and intelligent systems*. Upper Saddle River, NJ: Prentice-Hall.

Textbook has 5 chapters on expert systems and includes fundamentals of expert systems, knowledge acquisition and validation, knowledge representation, inferences, explanations and uncertainty, and building expert system. Book is a good source of definitions. Chapter 16 has a good model of the ESDLC.

Vanthienen, J. & Dries, E. (1995). Restructuring and simplifying rule bases. *Proceedings of Seventh International Conference on Tools with Artificial Intelligence*, 484-485.

The development of high quality knowledge base systems requires methods, which support the three major steps of the SDLC: knowledge acquisition and structuring, validation and verification, and implementation. The major issue when generating this full knowledge based application is how to implement the decision to logic. Two transformation categories can be distinguished: transforming the decision tables into a tree and transforming them into the single set of rules. In this paper, authors concentrate on one of these generation options, namely the minimal rule translation. This translation generates for each relevant action in the decision table one rule that describes the entire application field of that action in an optimal way, i.e., with as few conditions States as possible in the premise

Vanthienen, J. & Robben, F. (1993). Developing legal knowledge based systems using decision tables. *Proceedings of the Fourth International Conference on Artificial Intelligence*, 282-291.

Paper describes the methodology, tools used and experience in developing an expert system used to determine financial benefits for the disabled in Belgium. Authors discuss use of decision tables for KA, representation and for testing consistency (verification and validation) of legal knowledge. Includes definition of decision tables and steps for constructing decision tables.

Warren, J., Warren, D. & Freedman, R. (1993). A knowledge-based patient data acquisition system for primary care medicine. *Proceedings of the Second International Conference on Information and Knowledge Management*, 547-553.

Article describes an expert system (PDAS) for use in emergency rooms, designed to extract patient data for primary-care medicine. PDAS produces the printed report contains the patient perception of the symptoms and problems and has three components: (1) an opening interview with the patient, (2) an inferencing system that analyzes natural language statements of the problems and symptoms to determine which

physical systems should be reviewed, and (3) detailed reviews of the physical systems, which are investigated been detailed to provide the physician with information to confirm or rule out possible diagnosis. Process uses forward chaining.

Authors note that the literature reports several computer-based interview systems developed to gather medical data from patients

Waterman, D.A. (1985). *A guide to expert systems*. Reading, Mass: Addison-Wesley Publishing.

Book is one of the early textbooks on the subject of expert systems and is well known as an excellent introduction to expert systems. It is a good source of definitions in the subject as well as a reference for expert system tools, knowledge representation and methodology. Also includes: knowledge acquisition (chap 14), difficulties in development (chap 16), and common pitfalls in planning (chap 17), in dealing with domain experts (chap 18) and in development (chap 19).

He defines a domain expert as an articulate, knowledgeable person with a reputation for producing good solutions to problems in a particular field. He defines knowledge acquisition as the process of extracting, structuring and organizing knowledge from several sources, usually domain experts, so it can be used in a program.

Watkins, A., Dimopoulos, N., Neville, S. & Li, K. (1993). Flowtool: a procedural-knowledge acquisition tool. *IEEE Pacific Rim Conference on Communications, Computers & Signal Processing*, 1, 31-34.

Authors present FLOWTOOL, a procedural-knowledge acquisition tool developed specifically to acquire procedural knowledge associated with the diagnosis and calibration procedures of engineering systems. It uses a graphical user interface to acquire procedural knowledge and delivers a knowledge base complete with links to related hypermedia document(s). FLOWTOOL has been implemented and operates in a Unix/X-Windows environment and creates rules for NEXPERT OBJECT and PROLOG. FLOWTOOL has been used to capture diagnostic and calibration knowledge pertaining to C-COR main trunk amplifiers.

Wolfgram, D., Dear, T., & Galbraith, C. (1987). *Expert systems for the technical professional*. New York: John Wiley & Sons.

Book is a general text on expert systems. Authors discuss introduction to expert systems, structures, components, and construction of expert systems. Key to text is chapter on knowledge engineering and knowledge acquisition including definitions, techniques and tools. Authors also discuss types of applications suitable for expert systems and stages of development. Expert systems (ES) are defined as “computer systems, comprising both

hardware and software that mimic an expert's thought processes to solve complex problems in a given field" (domain). ES address and solve knowledge-intensive problems (large information sets) that can have multiple correct or acceptable answers. Suitable applications for ES include: interpreting and identifying, predicting, diagnosing, designing, planning, monitoring, debugging and testing, instructing and training, and controlling. ES are appropriate where there are not established theories, where human expertise is scarce or in high demand, and where the information is cloudy or fuzzy. ES analyze and present the best possible answer with advice and recommendations. According to Barr & Feigenbaum (1981), production rules are best used in domains of diffused knowledge. These are domains consisting of many facts, such as clinical medicine.

Authors note that one of the most critical responsibilities of the knowledge engineer is KA. KA is defined as the process of identifying, extracting, documenting and analyzing the information processing behavior of domain experts in order to define an expert system's knowledge base and inference engine. ES will be only as good as the expertise it is built on---quality of knowledge is key. For some domains, it is sufficient that only a minimal level of expertise be achieved; in other domains, such as medical diagnostics, it is absolutely necessary that the best expertise is incorporated. KA techniques include: interviews (unstructured, open-ended, etc.), protocol analysis, walkthroughs, questionnaires, expert reports. Protocol analysis (AKA, think-aloud method), particularly a set of techniques known as verbal protocol analysis, is by far the most common method by which the KE acquires detailed knowledge from the expert.

Zlatareva, N. (1993). An integrated approach to quality assurance of expert system knowledge bases. *Proceedings of the Second International Conference on Information and Knowledge Management*, 85-94.

Paper presents a quality assurance framework for rule-based expert systems, based on a sound logical theory that: (1) guarantees the completeness of verification and validation analysis; (2) allows rigorous reasoning about the causes of detected structural and functional errors; (3) helps the user to find new relevant sets of semantic constraints which the KB-theory must satisfy; and (4) generates the complete set of synthetic test cases to insure the systematic validation of the expert systems performance if no or a few real test cases are available.

References

- Aasgaard, D. (1978). An evaluation of data processing machine room low and selected recovery strategies. *Working paper MISRC-79-04, Management Information Systems Research Center, University of Minnesota*, 82 pages.
- American Institute of Certified Public Accountants (AICPA). 1988. The auditor's consideration of an entity's ability to continue as a going concern. *Statement on Auditing Standards No. 59*. New York, NY: AICPA
- Andrews, R. A. (1994). An ounce of prevention: Guidelines for preparing a disaster recovery plan. *Proceedings of the IEEE 1994 National Aerospace and Electronics Conference*, 2, 802-806.
- Awad, E. & Lindgren Jr, J. (1992). Skills and personality attributes of the knowledge engineer: an empirical study. *International Association of Knowledge Engineers '92 Proceedings*.
- Baldwin-Morgan, A. & Stone, M. (1995). A matrix model of expert system impact. *Expert Systems with Applications*, 9(4), 599-608.
- Braden, B., Kanter, J. & Kopcsó, D. (1989). Developing an expert systems strategy. *MIS Quarterly*, December, 1989, 459-467.
- Byrd, T. (1992). Implementation and use of expert systems in organizations: Perceptions of knowledge engineers, *Journal of Management Information Systems*, 8(4), 97-.
- Byrd, T., Cossick, K. & Zmud, R. (1992). A synthesis of research on requirements analysis and knowledge acquisition techniques. *MIS Quarterly*, 16(1), 117-138.
- Chen, S. (1993). A knowledge acquisition scheme for ruled-based systems. *Proceedings IEEE Region 10 Conference on Computer, Communication, Control and Power Engineering*, 2, 621-625.
- Cirullo, M. J., & McDuffie, R. S. (1994). Planning for disaster. *CPA Journal*, 64(6), 34-37.
- Christensen, S. & Schkade, L. (1987). Financial and functional impacts of computer outages on business. *Working paper CRIS-87-01, Center for Research on Information Systems, University of Texas at Arlington*, 16 pages.
- Cullen, J. & Bryman, A. (1988). The knowledge acquisition bottleneck: time for reassessment? *Expert Systems*, 5(3), 216-225.

- Durkin, J. (1996). Expert systems: a view of the field. *IEEE Expert*, 56-63.
- Fujihara, H., Simmons, D., Ellis, N. & Shannon, R. (1997). Knowledge conceptualization tool. *IEEE Transactions on Knowledge and Data Engineering*, 9(2), 209-220.
- Garcia-Molina, H. & Polyzois, C. (1990). Issues in disaster recovery. *COMPCON Spring '90. Intellectual Leverage. Digest of Papers. Thirty-fifth IEEE Computer Society International Conference*, 573-77.
- Gorney, D. & Coleman, K. (1991). Expert systems development standards. *Expert Systems with Applications*, 2, 239-243.
- Griswold, J. S., Lightle, T. L., & Loverlady, J. G. (1990). Hurricane Hugo: Effect on state government communications. *IEEE Communications*, 28(6), 12-17.
- Harmon, P. & King, D. (1985). *Expert systems, artificial intelligence in business*. New York: Wiley.
- Hayes-Roth, F. & Jacobstein, N. (1994). The state of knowledge-based systems. *Communications of the ACM*, 37(3), 26-39.
- Hiles, A. (1992). Surviving a computer disaster. *Engineering Management Journal*, 2(6), 271-74.
- Hoffman, R. (1987). The problem of extracting knowledge of experts from the prospective of experimental psychology. *AI Magazine*, 8, 53-64.
- Hwang, G. (1994). Knowledge elicitation and integration from multiple experts. *Journal of Information Science and Engineering*, 10, 99-109.
- Jacobs, J., & Weiner, S. (1997). The CPA's role in disaster recovery planning. *CPA Journal*, 67(11), 20-25.
- Kemp, E. & Purvis, M. (1996). The role of the individual project in teaching knowledge acquisition. *Proceedings of the International Conference in Software Engineering: Education and Practice*, 138-143.
- Keyes, J. (1989). Why expert systems fail. *AI Expert*, 4(11), 50-53.
- LaSalle, A. & Medsker, L. (1991). Computerized conferencing for knowledge acquisition from multiple experts. *Expert Systems with Applications*, 3, 517-522.
- Lichte, K. (1993). Knowledge capture model for expert systems. *Proceedings First New Zealand International Two-Stream Conference on Artificial Neural Networks and Expert Systems*, 163-164.

- Liebowitz, J. (1993). Educating knowledge engineers on knowledge acquisition. *IEEE International Conference on Developing and Managing Intelligent Systems Projects*, 110-117.
- Liou, Y.I. (1999). Expert system technology: knowledge acquisition. In J. Liebowitz Ed.), *The Handbook of Applied Expert Systems*. (pp. 2.1 – 2.11). New York: CRC Press.
- McGraw, K. L., & Harbison-Briggs, K. (1989). *Knowledge acquisition: Principles and guidelines*. Englewood Cliffs, NJ: Prentice-Hall.
- Money, W. & Harrald, J. (1995). The identification of group support systems to knowledge acquisition for disaster recovery planning. *Proceedings of the Twenty-eighth Hawaii International Conference on System Sciences*, 4, 468-474.
- Motoda, H., Mizoguchi, R., Boose, J. & Gaines, B. (1991). Knowledge acquisition for knowledge-based systems. *Paper presented at the IEEE Expert*.
- Nguyen, T., Perkins, W., Laffey, T. & Pecora, D. (1990). Knowledge base verification. In J. S. Chandler & T.-P. Liang (Eds.), *Developing Expert Systems for Business Applications*. (pp. 69-77). Columbus, OH: Merrill Publishing.
- Niederman, F. (1996). Acquiring knowledge about group facilitation: research propositions. *Proceedings of the 1996 conference on ACM SIGCPR/SIGMIS*. 58-67.
- Nuseibeh, B., Kramer, J., & Finkelstein, A. (1994). A framework for expressing the relationship between multiple views in requirements specifications. *IEEE Transactions on Software Engineering*, 20(10), 760-773.
- O'Keefe, R., Balci, O, & Smith, E. (1990). Validating expert system performance. In J. S. Chandler & T.-P. Liang (Eds.), *Developing Expert Systems for Business Applications*. (pp. 91-102). Columbus, OH: Merrill Publishing.
- O'Keefe, R. & Lee, S. (1990). An integrated model of expert system verification and validation. *Expert Systems with Applications*, 1(3), 231-236.
- O'Leary, D. (1990). Validation of expert systems with applications to auditing and accounting expert systems. In J. S. Chandler & T.-P. Liang (Eds.), *Developing Expert Systems for Business Applications*. (pp. 78-90). Columbus, OH: Merrill Publishing.
- O'Neil, M. & Morris, A. (1989). Expert systems in the United Kingdom and evaluation of development methodologies. *Expert Systems*, 6, 90-99.

- Olson, J. & Rueter, H. (1987). Extracting expertise from experts: methods for knowledge acquisition. *Expert Systems*, 4(3), 152-168.
- Rees, P. L. (1996). User participation in expert systems. *Industrial Management & Data Systems*, 93(6), 3-7.
- Rook, F. & Croghan, J. (1989). The knowledge acquisition activity matrix: a systems engineering conceptual framework. *IEEE Transaction on Systems, Man & Cybernetics*, 19(3), 586-597.
- Rudolph, C. Business continuation planning / disaster recovery: a marketing perspective. *IEEE Communications Magazine*, 28(6), 25-28.
- Sangster, A. (1996). Expert system diffusion among management accountants: a U.K. perspective. *Journal of Management Accounting Research*, 8, 171-182.
- Song, I. -Y., Strum, S. & Medsker, C. (1991). Design and implementation of a database design aid using VP-Expert. *Proceedings of the IEEE/ACM International Conference on Developing and Managing Expert System Programs*, 1991, 15-23.
- Smith, P. (1996). *An introduction to knowledge engineering*. London: International Thompson Computer Press.
- Smith, P., Ross, P., Awad, E., Green, C. & MacIntyre, J. (1994). A survey of the skills and personality attributes of the knowledge engineer in the United Kingdom.
- Starr, M. (1997). The state of disaster recovery planning in Texas small businesses. Unpublished doctoral dissertation, Nova Southeastern University, Florida.
- Stein, E. W. (1993). A method to identify candidates for knowledge acquisition. *Journal of Management Information Systems*, 9(2), 161-178.
- Stylianou, A., Smith, G., & Madey, R. (1993). Selection criteria for expert system shells. A sociotechnical framework. *Communications of the ACM*, 35(10), 30-48.
- Stylianou, A., Smith, G., & Madey, R. (1995). An empirical model for the evaluation and selection of expert system shells. *Expert systems with applications*, 8(1), 143-155.
- Tsai, N, Necco, C. & Wei, G. (1994). An assessment of current expert systems: Are your expectations realistic? *Journal of Systems Management*, November 1994, 28-32.
- Turban, E. & Aronson, J. (1998). *Decision support systems and intelligent systems*. Upper Saddle River, NJ: Prentice-Hall.

- Vanthienen, J. & Dries, E. (1995). Restructuring and simplifying rule bases. *Proceedings of Seventh International Conference on Tools with Artificial Intelligence*, 484-485.
- Vanthienen, J. & Robben, F. (1993). Developing legal knowledge based systems using decision tables. *Proceedings of the Fourth International Conference on Artificial Intelligence*, 282-291.
- Warren, J., Warren, D. & Freedman, R. (1993). A knowledge-based patient data acquisition system for primary care medicine. *Proceedings of the Second International Conference on Information and Knowledge Management*, 547-553.
- Waterman, D.A. (1985). *A guide to expert systems*. Reading, Mass: Addison-Wesley Publishing.
- Watkins, A., Dimopoulos, N., Neville, S. & Li, K. (1993). Flowtool: a procedural-knowledge acquisition tool. *IEEE Pacific Rim Conference on Communications, Computers & Signal Processing*, 1, 31-34.
- Wolfgram, D., Dear, T., & Galbraith, C. (1987). *Expert systems for the technical professional*. New York: John Wiley & Sons.
- Zlatareva, N. (1993). An integrated approach to quality assurance of expert system knowledge bases. *Proceedings of the Second International Conference on Information and Knowledge Management*, 85-94.