

AD-A107 065


A STRATEGIC PLAN FOR SUPPORT OF EXPERT SYSTEMS IN
ORGANIZATIONS(U) AIR FORCE INST OF TECH
WRIGHT-PATTERSON AFB OH SCHOOL OF SYSTEMS AND LOGISTICS
P R BOGGS SEP 87 AFIT/GLN/LSV/875-6

1/2

UNCLASSIFIED

F/G 12/8

NL






MICROCOPY RESOLUTION TEST CHART

1963 A

AD-A187 865



A STRATEGIC PLAN FOR SUPPORT OF
EXPERT SYSTEMS IN ORGANIZATIONS

THESIS

Paul R. Boggs
Major, USAF

AFIT/GLM/LSY/87S-6

DTIC
ELECTE
DEC 04 1987
S H D

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

87 11 24 018

AFIT/GLM/LSY/87S-6

A STRATEGIC PLAN FOR SUPPORT OF
EXPERT SYSTEMS IN ORGANIZATIONS

THESIS

Paul R. Boggs
Major, USAF

AFIT/GLM/LSY/87S-6

DTIC
ELECTE
DEC 04 1987
S D
H

Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

Approved for public release; distribution unlimited

The contents of the document are technically accurate, and no sensitive items, detrimental ideas, or deleterious information is contained therein. Furthermore, the views expressed in the document are those of the author and do not necessarily reflect the views of the School of Systems and Logistics, the Air University, the United States Air Force, or the Department of Defense.

AFIT/GLM/LSY/87S-6

A STRATEGIC PLAN FOR SUPPORT OF
EXPERT SYSTEMS IN ORGANIZATIONS

THESIS

Presented to the Faculty of the School of Systems and Logistics
of the Air Force Institute of Technology
Air University
In Partial Fulfillment of the
Requirements for the Degree of
Master of Science in Logistics Management

Paul R. Boggs, B.A., M.S.
Major, USAF

September 1987

Approved for public release; distribution unlimited

Acknowledgements

I would like to take this opportunity to recognize those individuals who provided inspiration and support in the production of this thesis. Special thanks go to Dr. Ed Mahler, Gary Sauer, Bruce Campbell, Bill Baker, Steve Wagner, Doug Barton, and Leonard Gross for their help and encouragement. I would like to acknowledge my thesis advisor, Major Mary Kay Allen, whose expert and timely advice is reflected in the final product. I would like also to thank my wife, Glenda Boggs, whose constant assertions that I should be working on my Ph.D., was inspiration enough to finish this dissertation. Finally, I would like to thank my sons, Paul, Matthew, and Robert for reminding me that there are other priorities in life.

Table of Contents

	Page
Acknowledgements	ii
List of Tables	v
Abstract	vi
I. Introduction	1
Background	1
Problem Statement	3
Research Questions	4
Scope of the Study	4
Definitions	5
II. Literature Review	8
Chapter Overview	8
Acquisition	8
Performance	13
Documentation	16
Maintenance	17
Organization	18
Personnel	18
Software Modification	19
Maintenance Costs	20
Lessons Learned	21
Summary	22
III. Methodology	28
Chapter Overview	28
Justification	28
Questionnaire	29
The Survey Questions	30
Interviews	38
Population	38
Census	39
Data Collection Plan	40
Data Analysis	41
Summary	42
IV. Findings	43
Chapter Overview	43
Questionnaire Results	43
Descriptive Analysis	45
Interview Results	59
Summary	68

	Page
U. Summary, Conclusions, and Recommendations	69
Chapter Overview	69
Significance of the Research	69
Summary of the Research Effort	70
Conclusions	72
Recommendations for Further Research	81
Appendix A: The Survey Instruments	83
Appendix B: Frequency Tables by Survey Question	97
Appendix C: Transcripts	127
Bibliography	144
Vita	147

List of Tables

	Page
2-1. Reported Expert System Performance Criteria	24
2-2. Acquisition Factors Reported in the Literature	26

Abstract

✓ The purpose of this study was to elicit lessons learned in business and industry in the fielding of expert systems. The study looked at 12 research questions dealing with the acquisition, maintenance, documentation, organizational, performance, personnel, and software modification issues associated with the support of expert systems. The study surveyed the Fortune 500 corporations via questionnaire and telephonic interviews.

The study found that the Fortune 500 corporations viewed expert systems as an instrument of strategic advantage and were very protective of what is considered proprietary information. The study found the corporations to be particularly secretive about costs of acquiring and maintaining expert systems.

Analysis of the questionnaire and interview data found that expert systems in the Fortune 500 are predominately delivered on personal computer hardware and software systems, and primarily developed and maintained in-house. Support issues such as the software modification, performance, and documentation issues are still being addressed.

(continued on page 2)

A STRATEGIC PLAN FOR SUPPORT OF EXPERT SYSTEMS IN ORGANIZATIONS

I. Introduction

Background

An expert system is a computer system that has stored in it the problem solving knowledge of one or more human experts in a particular field. Expert systems have the potential of being able to aid human non-experts in problem solving under conditions of uncertainty and where there are many possible solutions, none of which are clearcut. The past few years have seen a rapid growth in the use of expert systems in business and industry. Just as the private sector views the use of expert systems as a means of enhancing productivity and extending rare in-house expertise, the Air Force has made a conscious decision to apply the technology of expert systems in the management of its resources (2).

The first expert system was fielded in 1979 by John McDermott of Carnegie-Mellon University, for Digital Equipment Corporation (DEC) and was called R1, now XCON (25). R1 was such a success, reportedly saving DEC about \$200,000 a month in staff costs (8:17), that the system's area of "expertise" was expanded from configuration of Vax 11/780 systems only, to ten different DEC systems by the end of 1984 (7).

Since XCON, many prototype expert systems have been reported in the literature as "being developed", "in use", or "being discussed". (8) However, aside from the article "R1 Revisited: Four Years In the Trenches", by Judith Bachant and John Mc Dermott, there is little to be found on evaluations and support of actual, fielded, expert systems. Perhaps the problem is best stated by Wendy Rauch-Hindin:

In the past few years, knowledge-system tools and prototypes have become widespread in business and industry. Moving these tools and prototypes to production environments where they can be used in an organization's everyday operations is this year's challenge (30:69).

The Air Force faces the same challenge as it begins to bring its own expert systems online. One indication that the Air Force has accepted that challenge is evidenced by the establishment of an AI Management Integration Office (AIMIO) at AFLC Headquarters "to develop a coordinated command AI program." (2:3) A further indicator is the four phase introduction of the Inventory Management Assistant (IMA) expert system developed by Dr. Mary Kay Allen as part of her dissertation effort at the Ohio State University (14).

The AIMIO, in implementing the IMA system, will be concerned with its chartered goal to "increase ... productivity in all functional areas with relatively modest investment" (2:1). In following this charter, the AIMIO or similar organizations, need to know what aquisition, maintenance, and documentation costs will be for a particular expert system to properly assess the cost

benefits of implementing a system. The organization needs to know how many personnel will be required to program and maintain the expert system. Acceptable performance must be specified in advance and verification and validation procedures set up to insure the system performs at the desired level. Once fielded, the organization needs to determine proper software modification and configuration control procedures. The manager of the organization implementing expert systems requires information regarding these cost, personnel, performance, procedural, and organizational issues to efficiently and effectively plan support for an expert system. At the present time, the Air Force does not have the direct experience of fielding expert systems that business and industry have.

Problem Statement

Since a review of the literature revealed little useful information about fielded expert systems either in the private or commercial sectors, an attempt was made to obtain this information. Specifically, lessons learned in business and industry needed to be captured and a prescriptive plan developed, which would outline how the Air Force should address the acquisition, maintenance, documentation, organizational, performance, personnel, and software modification issues associated with expert systems. This plan should provide long term support for expert systems to enhance the Air Force's resource management capabilities.

Research Questions

In order to develop a support strategy and plan, several research questions were posed:

1. What should the acquisition costs be for expert systems?
2. How should field performance of the expert system be specified?
3. What are the best ways to verify and validate expert systems?
4. What documentation is required for fielded expert systems?
5. Who will maintain expert systems for the organization?
6. Should the responsibility for expert systems be centralized or decentralized?
7. What are the organic personnel requirements for maintenance of an expert system?
8. How will software deficiency reporting be managed?
9. How often should expert systems be modified?
10. How will configuration control be maintained?
11. How much should expert systems cost the Air Force for maintenance?
12. What specific lessons can be learned from organizations that have successfully fielded expert systems?

Scope of the Study

This study did not examine expert systems within the Department of Defense. The study reviewed only expert systems in industry. Further, the study focussed on the

fielding and implementation of expert systems, not on the development of prototypes.

This study gathered data concerning the maintenance and performance of expert systems in the private sector and applied a lessons learned approach to suggesting a plan for support of expert systems for the Air Force. Much of this plan should also be applicable to other DOD organizations and firms in the private sector.

Definitions

The following terms are used frequently throughout the text. For the purpose of this research effort, they are defined as follows:

Artificial Intelligence (AI). According to Feigenbaum and McCorduck (1981), artificial intelligence is

A subfield of computer science concerned with the concepts and methods of symbolic inference by a computer and the symbolic representation of the knowledge to be used in making inferences. A field aimed at pursuing the possibility that a computer can be made to behave in ways that humans recognize as 'intelligent' behavior in each other" (18:257).

Expert System. An expert system is a computer system that has stored in it the problem solving knowledge of one or more human experts in a particular field (24, 18, 34).

Fielded System. A fielded system is defined as an expert system that has progressed beyond the research prototype phase and is being used in daily operations for the purpose for which it was designed (36).

Frame. A frame is "a knowledge representation scheme that associates an object with a collection of features (e.g., facts, rules, defaults, and active values)" (18:260).

Each feature is stored in a slot and a frame is that set of slots that represent an object (18:260).

Knowledge Engineering. Knowledge engineering is the process of building an expert system (36:392). This process involves interviewing and working with domain experts to capture their knowledge before, during, and after programming the expert system.

Maintenance. Maintenance of an expert system involves the actual software modification of the system (18, 34).

Performance. Performance, as it pertains to an expert system, is evaluated three ways: is the system efficient, effective, and accepted by the user? The expert system is efficient if it saves time for the user and/or the expert. The system is effective if an appropriate answer is arrived at during a consultation. Whether or not a user accepts the expert system is usually a function of the user friendliness of the system and the user interface (36).

Prototupe. An expert system prototype is "...an initial version of the expert system, usually from 25 to 200 rules, that is developed to test overall knowledge representation and inference strategies being employed to solve a particular problem" (18:265).

Rule. In an expert system a rule is a two part conditional statement. "The first part, comprised of one or

more IF clauses, establishes conditions that must apply if a second part, comprised of one or more THEN clauses, is to be acted upon" (18:265).

Support. As used in this research effort, support is defined as the maintenance and configuration control of, and documentation for, an expert system.

Validation. Validation is the process of ensuring the expert system is producing an appropriate answer (26, 36).

Verification. Verification is the process of reviewing the entries in the knowledge base for accuracy, completeness, and consistency (11:2). This process is similar to debugging in conventional programming (26:42).

II. Literature Review

Chapter Overview

The purpose of this chapter is to provide an understanding from the literature of the importance of the support issues for expert systems and to discuss what answers to the research questions may be found in the literature. This chapter is organized into ten major sections. The first section discusses the acquisition issues in fielding expert systems. The second section deals with the performance issues and the third section discusses the documentation requirements for fielded expert systems. The next three sections deal with the maintenance, organizational, and personnel issues. The seventh section reviews the software modification issues. Maintenance costs are discussed in the eighth section. The ninth section examines some of the lessons to be learned from organizations that have fielded expert systems. The final section summarizes the chapter.

Acquisition

There are three issues relevant to the acquisition of a system which must be examined: the size of the expert system; the time required to develop the expert system; and the complexity of the expert system. This section looks at each of these issues in turn.

Size. One of the most important steps in planning an

expert system is the issue of size. Van Horn notes that "the capacity of your computer often limits the capabilities of an expert system" (34:74). The question not only arises of how much capacity is needed but also how capacity and capability are measured in expert system terms. Stephen Leibholtz, president of Analytics and manager of several artificial intelligence projects, suggests that "a measure of the size and complexity of an expert system is provided by the number of rules it contains" (24:37). Bill Belew, manager of information services at Texas Instruments, also notes that "the TI PC or IBM PC with 640 Kbytes of memory can handle systems in the 100 to 500 rule range, depending on the complexity of the rules" (16:80). Although the memory capacity of a computer, expressed in bytes or kilobytes, will actually limit the size of an expert system, the 500 rule/PC measure was useful as a means of comparison in reviewing the literature .

The size of the expert system, measured in rules or frames varied widely in the literature. At one end of the spectrum were expert systems of forty rules reported by Dupont (13:105), 100 rules for an expert system at Ford Motor Company (10:24), and 160 rules for Informart's expert system, called SNAP (16:78). At the intermediate levels, Northrop reports an expert system with 500 rules (10:22), Westinghouse has an expert system with 400+ rules (16:77), and Dupont claims systems up to 500 rules (13:105). Companies with larger expert systems were American Express

with its Authorization Assistant, 1000-1500 rules (17:32); Coopers & Lybrand's ExperTax, 2000 rules(23); Westinghouse's CORA, 3000 frames (10:23); and Digital Equipment's XCON, with 8000 rules (12).

Complexity. Most expert systems are either rule or frame based systems. However, rules alone cannot be used as a measure of complexity. Rule based systems predominate and use IF-THEN statements. For example, IF a given situation exists, THEN take these actions. The more complex the expert system, the more situations or actions can be addressed. For example, one rule may be composed of several IF antecedents, followed by a "consequent" or action. The antecedents are linked by the logical conjunctions AND or OR (22:32). Thus a simple IF-THEN rule becomes a more complex IF A AND B OR C AND D, THEN X rule. Similarly, in frame based systems, the more "slots" per frame, the more complex the expert system. Data in this format was unavailable in the literature for fielded expert systems.

Another problem with using the number of rules as the only indicator of expert system complexity is that the data base is not accounted for. An expert system that accesses a general data base is an incomplete system without the data base (27:101). A systems approach to planning an expert system requires that the data base hardware and software requirements be considered in the planning phase for both current and projected needs.

When considering the hardware and software requirements for the delivered system, the manager should carefully review the environment in which the system will operate.

The greatest technology is useless if few people understand it. The power of knowledge system development tools lies in their ability to be used by people with no AI experience (30:71).

Leibholtz pointed out that "the AI tools best for research and advanced development are not the tools most appropriate for operational use" (24:38). Leibholtz also noted that "...once the expert system has been developed, it should be migrated to a more conventional system" (24:38). While LISP is recognized as an excellent research tool, the slowness and expense of this development tool discourages practical applications (24). LISP development requires a much higher degree of AI expertise, which is both expensive and in short supply, in both the programmer and the end user. However, commercial software tools called shells require much less training for both the programmer and the user and are far less expensive than the AI research tools. An expert system shell is software that provides the programmer with an "inference engine", the generic problem solving framework for the rest of the expert system, and some basic development tools, such as a text editor, debugger, and graphics generator. (16:76, 36:391) The programmer has to complete the expert system by adding the domain knowledge of the expert to the initially empty knowledge base (24:37). Some data were available in the literature as to the

hardware and/or software being used to field expert systems in the private sector.

Of all the references in the literature, only the APEX expert system, PlanPower, was delivered on a Xerox 1186 LISP machine (10:19). Several companies had expert systems running on mainframe computers. American Express (28), Metropolitan Life (31), and R.R. Donnelley & Sons (10) all use an IBM mainframe computer. Digital Equipment Corporation and Boeing run expert systems on VAX minicomputers (12, 31). Boeing also runs an expert system called "Expert Executive" on an Apollo workstation (20:44). The following firms were reported to have expert systems running on personal computers: Beckman Instruments, Chicago First National Bank, Coopers & Lybrand, Dupont, Ford, General Electric, Infomart, Travelers Insurance, and Westinghouse (35, 31, 37, 10, 32, 16). Westinghouse also reported development of an expert system on the Texas Instruments PC (16:78).

Time. In planning for the acquisition of an expert system, a manager needs to know how much time should be allotted for development of the expert system from conception to fielded system. The data available from the literature were presented in the form of development time from concept to working prototype. American Express reportedly developed its Authorization Assistant to the first prototype in six months but the system was not expected to become operational for another three to five

months (28:65). Ford Motor Company developed its 100 rule prototype in six weeks using two programmers (10:24). According to Dr. Ed Mahler at Dupont, development time from conception to a useable system takes about one hour per rule (13:105). Other firms reporting development times were Northrop, two months; IBM, 8 man-months; and Digital Equipment, one year for XCON (10, 15, 12).

Performance

At issue in the performance category are basically the answers to two questions. First, how should the expert system be specified in the planning stage? Second, how should the expert system be verified and validated?

Waterman suggests that one of the pitfalls in system testing and evaluation is that users may find the performance disappointing in terms of both quality and utility of answers produced (36:198). One of the first issues to be resolved in the planning of an expert system is how it will be specified and whether on-going testing or evaluation of the the system is necessary once the system is fielded. Waterman suggests that the prudent course is to "...be sure to specify the minimum acceptable performance that will allow the system to be considered a success" (36:198).

Specification. Specific lessons learned in the private sector were difficult to determine from the literature. Some companies specified an accuracy rate as their criterion

for a successful system. The most common indicators were cost savings expressed in dollars or as a reduction in the time required to do a particular job without the aid of the expert system. Whether these indicators were specified before or after fielding is unknown.

Three firms expressing accuracy as a standard were Campbell Soup, Digital Equipment, and IBM. Campbell Soup's goal was an expert system "that could diagnose 99 percent of a sterilizer's malfunctions (33:69). Digital Equipment says its expert system, XCON, has achieved a 98 percent accuracy rate (12). Dr. Herbert Schorr, group director of products and technology at IBM, says one of IBM's expert systems is 100 percent accurate (15:65).

A few companies have indicated a cost savings in using an expert system. Digital Equipment, for example, claims a 25 million dollar a year savings using XCON (17:31). IBM claims a five million dollar per year savings (15:65). Ed Mahler, head of AI activities at Dupont, does not specify a dollar amount saved, but claims an 800 percent return on investment for just ten Dupont expert systems (17:37). R.R. Donnelley & Sons is using an expert system that allows a 50 percent reduction in mailing costs per single mailing (10:25).

Some form of productivity increase was the predominant method corporations used to describe their system's performance in the literature. Some expressed the performance as a percentage increase in productivity, but

most used examples of reduced time to do a certain task. Dupont, for example, claims a payoff ratio increase of 10:1 for a successful expert system (37:35). Evensky & Brown uses an expert system called PlanPower to reduce the amount of computer time required to produce a financial plan from 50 hours to 10 or 15 hours (10:19). American Express claims a 25 percent reduction in decision time using the Authorization Assistant expert system (17:32). Beckman Instruments' expert system, SPINPRO, allows a 70 percent reduction in centrifuge run times (35:5). The Delco Product division of General Electric uses an expert system to reduce a four week job to less than an hour (17:31). Lockheed uses the Lockheed Expert System (LES) to shorten expert system development time from 18 months to four months (22:40). Northrop uses an expert system to reduce process planning time usually taking from eight to 12 hours to just five minutes (10:20,22). Steelcase reduced job times from 24 hours to minutes (31:102). Travelers Corporation also used PlanPower to reduce the time required to produce a financial plan from 30 to 12 hours (17:30). Westinghouse's expert system, CORA, reduces from one or two days to 15 seconds the time needed to select a suitable relay device for a customer (10:24).

Verification and Validation. An expert system is never completely finished (33:70). Van Horn notes that "fine tuning, expanding its capabilities, and even major revisions can continue indefinitely." (34:84). In 1984 McDermott and

Barchant predicted that XCON would "...continue to grow and evolve for as long as there is a configuration problem" (7:21). At that time XCON had about 3300 rules (7:23). XCON now has 8,000 rules (12). Because of this continual growth characteristic, the verification and validation of an expert system is the key to maintaining the accuracy and consistency of the results the system is supposed to produce. The greatest benefit of an expert system is consistency and uniformity in its results (29:99). Van Horn also states that "consistency and reliability ... are paramount" (34:197).

Verification is simply checking the knowledge base for accuracy, consistency, and completeness (11:2). Validation is the process of ensuring the expert system is producing an appropriate answer. One method suggested in the literature to validate an expert system is to maintain a library of test cases to test the expert system after modification (5:82). If the expert system arrives at the appropriate answer for the test case, then the expert system is validated. Specific examples of corporate validation procedures were not found in the literature.

Documentation

Documentation is a part of the support for an expert system. An assumption was made that, at one end of the spectrum, documentation might be simply a set of instructions, verbal or written, on how to power up the

system. A set of online or onscreen instructions would then help the user navigate through the system to a conclusion of that session. A more complex set of documentation would be required for a maintainer of the expert system.

Documentation might not be required for the end user if the expert system was embedded in another decision support or management information system being used and thus was transparent to the user. The user would simply be unaware that he or she was using an expert system (30:70). At the opposite end of the spectrum, a user's manual for the expert system and the rule listing provide more sophisticated examples of documentation (3:4). No specific information on existing corporate documentation was found in the literature.

Maintenance

For any expert system, there must be someone to maintain the system and to continue to serve as the source of expertise for the expert system. In some cases, the expert might be the person who programs the expert system. However, the person who maintains the system, once fielded, may not be the individual who originally programmed the system. At issue is whether the expert system development and/or maintenance will be performed in-house or contracted to a vendor. In some cases, the relative risk involved and the amount of resident AI expertise will heavily influence the decision to contract the maintenance of the expert

system. Firdman suggests that in-house expertise is a good idea. "However, ... let people develop in-house tools only if it is absolutely necessary and you have indisputable expertise in the field" (15:69).

Organization

Whether or not the organization will develop the expert system in-house, one issue to be considered is who will have the responsibility for that system. Harmon notes that "it's a good idea to make one person responsible for entering new rules whenever data or procedures change or whenever questions arise that the current system could not answer" (18:263). This approach advocates centralized, rather than decentralized responsibility for maintaining the system. There were no examples in the literature of decentralized control or responsibility. Other data, related to a corporation's AI organization, such as the size of the staff, if any, devoted to expert systems, the experience level of the individuals on the staff, and the duties of the staff were not found in the literature.

Personnel

One of the things a manager needs to know in planning the support for an expert system is the number of in-house personnel that will be required to maintain the system. This section assumes that the decision to maintain the system in-house has already been made by the firm. As mentioned earlier, the person who maintains the system, once

fielded, may not be the individual who originally programmed the system. Specific information on the number of personnel needed to maintain fielded expert systems in the private and commercial sectors was not found.

Software Modification

The software modification issues will be covered in this section. The issues are: how should software deficiency reporting be handled?; how often should the expert system be modified?; and how should configuration control of the system be maintained?

Deficiency Reporting. Deficiency reporting is the process of making the responsible person aware of a suspected or known problem with the expert system (3:4). The process may be formal or informal and may include suggested enhancements (3:5). If only one person will be using the system, a reasonable assumption would be that a formal organization or procedure is not necessarily needed or wanted. If the expert system serves several users, then a more structured approach might be needed to make sure the person or persons responsible for modifying the expert system are aware of the problems or suggested enhancements. The literature contained no information regarding the software deficiency reporting procedures corporations use in their organizations.

Modification. Assuming a responsible person or organization within the company is receiving software

deficiency reports and that the reported problems or suggested enhancements are valid, how often should the system be modified? Should the system be modified on a continuous basis, on an as required basis, or should reported deficiencies and suggested enhancements be made on a fixed periodic basis? The answers to the questions at issue were not found in the literature.

Configuration Control. Once an expert system has been prototyped, validated, and fielded, the question of how to maintain an "official" version of the expert system arises. An official version is the current approved edition of the expert system software with validated changes. The official version also reflects company policies. The problem is of greater magnitude in a larger organization that maintains and distributes the official version to various remote locations. Allen, Lammers, and Jenkins note that, within the Air Force Logistics Command, "a key concern will remain the desire to limit the proliferation of hardware and software systems, without curbing innovation" (6:7). If the Air Force is concerned about configuration control, one could logically assume that similar concern exists in the private sector. However, evidence to prove or disprove that assumption was not found in the literature.

Maintenance Costs

This is an area that is not covered specifically in the literature. An assumption was made that maintenance costs

would include personnel costs, system growth costs, documentation costs, and contractor costs. System growth is defined in this research effort as an increase in the size of the expert system due to an increase in the number of rules or frames, the amount of memory required, or the migration of the system from one hardware or software system to another. Specific lessons learned in this area were not available in the literature.

Lessons Learned

Lessons learned are the real objectives of this effort. Lessons learned were sought in the literature under all the previous categories but were found to be lacking in specific detail. For example, most firms that are willing to talk to the press about expert system activities relate only the advantages perceived by that company. Few mentions were made of disadvantages, problems, or drawbacks to an expert system. Some of the problems that were mentioned were keeping the knowledge base up to date (10:19), networking (28:65), a shortage of knowledge engineers (38:48), the slowness of expensive LISP research tools (24:38), extracting expertise from the expert (24:36), migrating LISP systems to personal computers (30:82), development and production costs (22:36), integration with existing hardware and software (30:69), and the lack of a natural language interface (9:62). The last point appeared to be a common complaint in the literature for development tools that

require technical proficiency beyond that of the average business user (16:77, 24:38). These problems and drawbacks were noted but more specific information was desired for the research effort.

Summary

The purpose of this chapter was to review the literature for answers to the research questions. The results of the review are summarized below, as reflected in tables 2-1 and 2-2, located at the end of this chapter.

As indicated by the gaps in tables 2-1 and 2-2, information sufficient to answer the research questions was not available from the literature. Specific, detailed information was not found for fielded systems but, a considerable information exists on research prototypes and ideas for expert system applications. Many of the references cited various companies as "developing" expert systems or describe an expert system "under development" but little is written about what is being done with operational systems.

Most books and articles keep covering the same old story -- the tired examples from the late 1970's or the heavily funded research projects at universities or corporations. Unfortunately, examples of everyday problems being solved using this technology are not readily available. --
Terry Hengel (19:8)

Cost and productivity benefits of using expert systems are reported but the methodology used to derive the figures is unclear. Further, almost nothing has been published

about the maintenance and support issues regarding specific problems and solutions to those problems. Harvey Newquist III, editor of AI Trends Newsletter, suggested that the proprietary nature of expert systems is responsible for the secrecy associated with expert systems (1:54). Another editor suggested that the problems being solved might not be "flashy enough to warrant publishing" (19:8).

Whatever problems are solved, minimal acceptable performance of the system needs to be specified in the planning stage (36:198). Standards of accuracy and speed need to be predetermined. These standards will help to ensure user acceptance when coupled with a user friendly interface (36:199).

The lack of published material to answer the basic research questions drove the need for data collection via some other avenue. A questionnaire/interview approach was decided upon as the methodology to be followed. The methodology is discussed in Chapter III.

TABLE 2-1

Reported Expert System
Performance Criteria

COMPANY	EXPERT SYSTEM	PRODUCTIVITY INCREASE	DOLLAR SAVINGS
Allied Signal			
American Express	Authorizatr Assistant	reduce decision time 25%	
Arthur D. Little	CORP		
Beckman Instrument	SPINPRO	70% runtime reduction	
Bell Labs	ACE		
Boeing			
Boeing	Expert Exec.		
Campbell Soup		99% accuracy	
Chicago 1st Natl.			
Coopers & Lybrand	ExperTax	"reduces time"	
DEC	XSITE		
DEC	XSEL		
DEC	XCON	98% accuracy	25 mill.
Dupont	300+	10:1 payoff	800% ROI
Evensky & Brown	Planpower	35-40 hours time reduction	
Ford			
GE/Delco Products		4 weeks to 1 hr	
GTE	Compass		
General Instrument			
General Dynamics	ARBY		
General Electric	CATS-1		
General Motors			
Honeywell	Mentor		

TABLE 2-1. continued.

COMPANY	EXPERT SYSTEM	PRODUCTIVITY INCREASE	DOLLAR SAVINGS
IBM		100% accuracy	5 million
IBM	100+		
IBM	DART		
ITT			
Infomart	SNAP		
Litton			
Lockheed	LES	18 mo. to 4 mo.	
Martin Marietta			
McDonnell Douglas			
Metropolitan Life			
Motorola			
Northrop		8-12 hrs. to 5'	
Proctor & Gamble			
R.R. Donnelley	More/2	70-80% response	50% of mail costs
RCA			
Raytheon			
Rockwell			
Sperry			
Steelcase		24 hr. to minutes	
TRW			
Texas Instruments			
Travelers Corp.	Planpower	30 to 12 hrs.	
Travelers Ins.	M.1		
U.S. Treasury			
Westinghouse	CORA	1-2 days to 15 sec.	

Sources: 1, 10, 12, 13, 14, 16, 17, 20, 21, 22, 23, 28, 31, 32, 33, 35, 37

TABLE 2-2

Acquisition Factors Reported
in the Literature

COMPANY	COST [\$]	TIME TO DEVELOP	HARDWARE	SIZE
Allied Signal American Express		6 mo.	IBM MF	1000- 1500 r.
Arthur D. Little Beckman Instrument	2600		IBM PC	
Bell Labs Boeing Boeing Campbell Soup Chicago 1st Natl.			Mini/MF Apollo WS IBM PC	 151 r.
Coopers & Lybrand DEC DEC DEC Dupont		1 yr. 20K each 1 rule/hr.	IBM PC-AT VAX-11 VAX-11 VAX-11 PC	2000 r. 8000 r. 40-500 r.
Evensky & Brown Ford GE/Delco Products GTE General Instrument	45K 5K	12 m/wk.	1186 Xerox IBM PC	 100+ r.
General Dynamics General Electric General Motors Honeywell IBM		4 mo.	PC	1500 r.
IBM IBM ITT Infomart Litton			PC	160 r.

TABLE 2-2. continued.

COMPANY	COST [\$]	TIME TO DEVELOP	HARDWARE	SIZE
Lockheed		18 mo.		
Martin Marietta				
McDonnell Douglas				
Metropolitan Life	300K		IBM MF	
Motorola				
Northrop		2 mo.		500 r.
Proctor & Gamble				
R.R. Donnelley			IBM MF	
RCA				
Raytheon				
Rockwell				
Sperry				
Steelcase				
TRW				
Texas Instruments				
Travelers Corp.				
Travelers Ins.			IBM PC	
U.S. Treasury				
Westinghouse			UAX/PC	3000 frms
Westinghouse			TI Pro.	400+ r.

Sources: 1, 10, 12, 13, 14, 16, 17, 20, 21, 22, 23, 28, 31, 32, 33, 35, 37

III. Methodology

Chapter Overview

The purpose of this chapter is to describe the methodology used to collect and analyze the data which was collected to answer the research questions posed in Chapter I. The first section provides justification for use of the survey approach to answer the research questions. The survey instruments are discussed in the next three sections.

The population of interest in the research effort and the rationale for taking a census rather than a sample of the population are discussed in the fifth and sixth sections, respectively. The next section describes the data collection plan. Data analysis is explained in the eighth section and a summary of the chapter is included as the final section.

Justification.

Since the information necessary to answer the research questions posed in Chapter I was unavailable in the literature, three options were available to collect the data. These options were mail surveys, telephonic interviews or personal interviews. The first option, survey by mail, was selected initially as the best method of obtaining a large sample of data from the target population. This method also made possible a census of the population. Another advantage was the relatively impersonal nature of

the mail questionnaire. This attribute was reinforced by a written promise of anonymity for the respondents. More candid responses were expected if the respondents knew that responses were to be grouped and not attributed to any particular person or company.

Questionnaire

A questionnaire was used in this study to collect data to answer the research questions posed in Chapter One. An existing questionnaire, developed by Major Mary Kay Allen and Lt Col Robert Peschke, AFIT School of Systems and Logistics faculty, was found and extensively modified to collect the data needed. The questionnaire was submitted to two members of the faculty of AFIT for initial comment. Then the questionnaire was subjected to a pre-test by other AI experts at four companies and one Air Force organization known to be working with expert systems. All of the AI experts involved in the pretest possessed extensive experience in working with expert systems. After several minor revisions and changes to the format, the questionnaire was mailed out to the Fortune 500 companies.

The questionnaire was divided into two sections: one for all companies with a planned or fielded expert system(s) and a second section only for those companies with a fielded system. If a company had more than one expert system, the respondent was asked to respond to the questions for that system for which the company had the most experience. The

first section requested the following information:

Company name.
Product line.
The existence or development of an expert system at that company.
Information on the expert system:
 The Functions,
 Objectives, and
 Users of the system.
 Other pertinent information.
 The stage of development.
 The hardware host.
 The shell or programming language used.
 Whether development was in-house, contracted, or both.
 The size of the system.
 The date development began.
 The date the system was or would be ready to field test.
 The advantages or disadvantages of the system.
 Any drawbacks to the application.
 The rationale if no system either existed or was planned.

The second section derived information to answer the research questions based on actual fielded system experience. The questions were designed to provide an indicator of the experience the company had in working with fielded systems and to elicit the lessons learned by the company.

The Survey Questions

Questions on the survey mailed to the Fortune 500 were designed specifically to answer the research questions. The following discussion reiterates the research questions and relates them to the questions on the survey designed to answer them.

Research Question 1. How much should expert systems cost the Air Force for acquisition?

Although one survey question, number 27, directly requested information about the acquisition costs for the respondent's system, an assumption was made that this one question might not capture the required data. Certainly, question 27 alone would not provide a good indicator of cost without also considering other factors such as the size of the system, the complexity, the hardware host, the programming language or expert system tool, external links to hardware, software programs, data bases, etc., personnel, and the time required to field the system.

In my opinion, several factors should be considered in the acquisition of an expert system. The hardware factor appears quite important. A system running on a PC will most likely cost less than one running on a LISP machine, mainframe, or minicomputer. A LISP machine, for example, costs between \$50,000 and \$100,000 (22:40). A PC costs about \$5000 (37:36). The software is also likely to be an important factor. Popular commercial expert system development tools are ART, \$65,000; KEE, \$30,000; and Knowledge Craft, \$50,000 (38:47). The choice of the expert system development tool may be dependent on the hardware host. If the expert system will operate on a mainframe, the software cost will be higher than if the hardware host were a PC. For example, Aion's ADS shell costs \$60,000 for the IBM mainframe version and \$7,000 for the IBM PC version

(38:47). An expert system hosted on a PC but accessing a data base located on a mainframe will require external hardware and software links to the data base. Each piece of the system adds incremental cost to the whole system and should be considered. Development time is another important factor. The longer the development time, the more costly the system becomes. The amount of capital and company AI expertise available will affect the decision to develop an expert system in-house. Size and complexity are also factors. A larger, more complex expert system will require more time to develop. Dupont uses an estimate of a rule per hour to develop expert systems using shells (37:36). Development of expert systems is faster using shells, regardless of the size of the computer (16:77). All of the following survey questions were designed to address these acquisition factors:

Survey Question 4. Is the application of an expert system, within your organization, based on a (check all that apply):

☐ Personal Computer (e.g., IBM PC, Apple Macintosh, etc.)?

☐ Large, Dedicated AI (Artificial Intelligence) Workstation?

☐ Minicomputer (e.g., VAX)?

☐ Mainframe (e.g., DEC, IBM, etc.)?

Survey Question 5. For each expert system in development/use, please identify the Shell or Programming Language that is being used:

Survey Question 6. What, if any, links to external (outside of the expert system(s)) hardware, programs, data bases, and/or

communications are in use? Please describe.

Survey question 7. Has your Expert System(s) been developed, or will it be developed:

____ Completely In-House?

____ Partly In-House and Partly Contracted Out?

____/____ Please indicate approximate percentage of each.

____ Completely Contracted Out?

Survey Question 8. If your system(s) was developed in-house, was your system(s) developed:

____ by end users?

____ by a service organization within your company?

____ by a mix of both the above?

____/____ Please indicate approximate % of each.

Survey Question 10. Please indicate the SIZE of your Expert System(s) (approximate number of rules or frames being used and the amount of computer memory required to operate the system in its current/planned configuration, including data base):

What is the average size of each rule or frame (i.e., the number of conjunctions/disjunctions in the rule antecedent or the number of slots/facets in each frame)?

Survey Question 11. Please indicate the approximate date system development began and the approximate date the system was (will be) ready to field test.

Survey Question 27. What were/are the approximate system costs for: Acquisition (from decision to buy until system was ready to field test)

Survey Question 28. In hindsight, how could costs of acquiring and maintaining the system(s) have been reduced in the planning and implementation phases?

Research Question 2. How should field performance of the expert system be specified?

Questions 12a and 26 from the survey were expected to elicit answers to the performance specification issues.

Survey Question 12a. What results/advantages have you experienced from this application to date? (e.g., percent increase in productivity, personnel reduction, cost-savings, etc.)

Survey Question 26. How is the system's performance measured? (e.g., what criteria, such as accuracy of advice, time to complete a consultation with the system(s), etc., are used to evaluate performance/usefulness?)

Research Question 3. What are the best ways to verify and validate expert systems?

Question 25 on the survey was designed to elicit a response resulting from lessons learned in the firm.

Survey Question 25. What verification/validation procedures work best for the company?

Research Question 4. What documentation is required for fielded expert systems?

Survey question 17 was expected to elicit responses to the documentation issue.

Survey Question 17. What documentation (e.g., specifications, test plan, user's manual, etc.) has been established for the system(s)? Do you feel this documentation is sufficient? If not, why not?

Research Question 5. Who will maintain expert systems for the organization?

Questions 18 and 23 were expected to address the maintenance and personnel issues directly.

Survey Question 18. How is the system maintained?

☐ Completely Contracted Out

☐ Partly Contracted Out, Partly In-house

☐ Totally In-house

Why did you choose to maintain the system in this manner? (lack of in-house expertise, cost, etc.)

Survey Question 23. Who performs the modifications?

☐ Contractor and/or ☐ In-house knowledge engineer(s)

What lessons have you learned regarding system updates?

Research Question 6. Should the responsibility for expert systems be centralized or decentralized?

Survey questions 19 and 20 were expected to help answer the organizational issue.

Survey Question 19. Is there a central organization within your company with primary responsibility to oversee the acquisition, fielding, and maintenance of your expert system(s)?

☐ Yes ☐ No

If so, how many people staff this function?

If not, has your company ever had, or planned to have, such an internal organization?

☐ Yes ☐ No

Survey Question 20. What lessons have been learned by your company as to the necessity of centralized versus decentralized control of expert systems within your organization?

Research Question 7. What are the organic personnel requirements for maintenance of an expert system?

Question 9 addresses the personnel requirements for organic maintenance.

Survey Question 9. What is the size of your staff devoted only to expert systems? _____

Research Question 8. How will software deficiency reporting be managed?

Survey question 21 seeks a response based on lessons learned in the company.

Survey Question 21. How is software deficiency reporting for the expert system(s) handled in your company? (special forms required, who reports, who receives reports, who acts on reports, etc):

Research Question 9. How often should expert systems be modified?

Question 22 addresses another of the software modification issues.

Survey Question 22. How often is the system modified? (e.g., quarterly, whenever deficiency detected, policy changes, etc.?)

Research Question 10. How will configuration control be maintained?

Question 24 on the survey addresses the third issue in the software modification area.

Survey Question 24. How is configuration control of the system being maintained? Is configuration control a problem? What recommendations do you have regarding configuration management of expert systems?

Research Question 11. How much should expert systems cost the Air Force for maintenance?

Questions 14, 15, 16, 27, and 28 were expected to elicit responses regarding the maintenance of an expert system in terms of both dollars and lessons learned:

Survey Question 14. How long has this system(s) been in operational use within your company? How many end users actually use the system(s)?

Survey Question 15. Has this expert system(s) grown in size? (rules/frames/memory requirements) If so, what was the original size? Has the inference engine been changed?

Survey Question 16. Has it been moved from one hardware/software system to another (e.g., hardware: IBM PC to mainframe, etc.; software: LISP to "C", etc.). If so, what were the lessons learned?

Survey Question 27. What were/are the approximate system costs for:

Annual Maintenance (documentation, 24-hour call-in service, etc.):

Survey Question 28. In hindsight, how could costs of acquiring and maintaining the system(s) have been reduced in the planning and implementation phases?

Research Question 12. What specific lessons can be learned from organizations that have successfully fielded expert systems?

Several questions in the questionnaire solicited responses to lessons learned under the various areas. Questions 12b and 13 were designed to stimulate responses of a more negative nature. These questions were open-ended so that the respondent would feel constrained to restrict his or her answer to a particular area.

Survey Question 12b. What difficulties, if any, were encountered in scaling the expert system up from the prototype model to the fully operational system?

Survey Question 13. Please list any drawbacks to this application:

This section has reiterated the research questions and related them to the survey questions designed to answer them. The following sections discuss the interview instrument, the population, rationale for a census, the data collection plan, and data analysis.

Interviews

Telephonic interviews were conducted with management and AI experts at five of the companies surveyed based on the experience of the companies involved with fielded expert systems. The amount of experience with fielded expert systems was gleaned from a review of the literature and from the returned questionnaires.

An interview guide was prepared using the results from the questionnaire data analysis and the literature review. The intention of the telephonic interview was to amplify and validate questionnaire findings based on those companies having the most experience with fielded systems. Questions not answered or that had ten or fewer responses were initially identified for the interview guide. Other questions were added to amplify or enrich a response to a previous question. Specific responses that were unclear to the researcher were identified prior to a company's interview and included as part of the interview for clarification during the interview. Questions that were answered on the questionnaire were not duplicated in the interview process unless clarification was necessary. Interviews were tape recorded. A copy of the interview guide is included as part of Appendix A, Survey Instruments.

Population

The population of interest in this research effort was defined as the Fortune 500 companies listed in the 28 April

1986 issue of Fortune magazine. This population was chosen as the desired population of interest for two reasons. The literature review revealed that two-thirds of the Fortune 500 companies have AI projects staffed and underway. (35:6) The literature also revealed that two-thirds of AI applications in manufacturing firms are in the subfield of expert systems. (8) Second, the Fortune 500 companies are large, successful firms. An assumption was made that successful firms actively work at being innovative and productive and usually are the first to use new technology. Since expert systems are still quite expensive and capital intensive, these are the type of companies that might have both the will and capital to implement expert systems in their organizations.

Census

A census of the population was taken rather than a sample for the following reasons. Only a few of the Fortune 500 with ongoing AI projects were known. If at least two thirds of these companies have AI projects underway (35:6) and two thirds of all AI projects tend to be expert systems (8), approximately 200 of the Fortune 500 should be actively working with expert systems. In order to maximize the return of questionnaires from companies with expert system experience, a census was both desirable and necessary. In addition, the size of the population, 500, was small enough to take a census. Finally, the anticipated return rate was 10 to 15 percent. With such a low anticipated return rate,

a census was considered to be necessary to lend validity to the data. Second, if the return rate were to prove to be significantly higher, the increased validity of the results would be an added bonus for the research effort.

Data Collection Plan

When the company was known to have an AI department, the envelope was addressed to that department or the head of that department, if known. When this information was not readily available, the questionnaire was addressed to the chief executive officer of the company. A pre-addressed, stamped envelope was enclosed with the questionnaire to facilitate return. A cover letter was attached to the questionnaire explaining the purpose of the questionnaire, requesting participation, and offering an executive summary to those companies responding to the survey. A sample cover letter and the questionnaire are included in Appendix A. The questionnaires were mailed during the second week of May, 1987. Questionnaire data collection was suspended 4 August 1987.

Due to the time constraint imposed by the Air Force Institute of Technology thesis due date, a followup survey was not possible.

Interviews were conducted between 7 and 14 August 1987. An unplanned opportunity to visit with AI experts at Dupont, UNISYS, and DEC on 4 and 5 June 1987 was afforded the researcher prior to the return of the questionnaire and completion of the interview guide. The meetings at Dupont

and UNISYS were tape recorded and relevant portions of those meetings are presented in Appendix C. The meeting at DEC was not tape recorded nor agreement reached with DEC on publishing the notes from that meeting. For those reasons the notes are not included as part of this report.

Data Analysis

The data were collected from the completed questionnaires. Since most of the data were in the form of qualitative answers, each questionnaire was hand scored and the applicable data were transferred to manual charts. Responses were tabulated and summary statistics calculated using a handheld calculator. Summary statistics used were frequency of response for both numerical and character data, mean and median response for numerical data, and modal response for character data. The standard deviation was also computed for numerical data. Due to the open-ended nature of many of the survey questions, some responses contained a numerical range of values. Whenever a numerical range of values was given as a response, the average value in the range was used in order to calculate the mean, median, and standard deviation of that set of data. A detailed discussion of how the results of each specific survey question were analyzed will be provided in Chapter IV.

Summary

This chapter has discussed the methodology used to conduct research to answer the research questions posed in Chapter I. First, the absence of the necessary information in the literature required a data collection effort. The survey approach to collecting the necessary data was justified as the most efficient method of collecting the data. A census, rather than a random sample, was determined to be desirable because of the relatively small size of the target population and to maximize coverage of those companies having expert system experience. The survey questions were related to the research questions and a data collection plan outlined. Finally, analysis of the data was discussed. The results are summarized in Chapter IV.

IV. Findings

Chapter Overview

The purpose of this chapter is to present the findings. Findings are discussed in relation to the research questions originally posed in Chapter I.

Questionnaire Results

Return Rate. Although the questionnaire return rate was less than anticipated, the low rate of response lent credence to Harvey Newquist's statement that companies are very secretive about expert systems because they use the technology for strategic advantage (1:54). Of those companies returning the questionnaire, 26 percent requested anonymity for the answers that were submitted.

Of the 500 questionnaires mailed out, 37 were returned from 35 different companies. Two companies returned two questionnaires, one each for two separate expert systems. One of the 35 companies returning the questionnaire chose not to participate in the research effort and merely returned the blank questionnaire. Four additional firms responded that they preferred not to participate, and did not return the questionnaire. The overall response rate was 8.2 percent for 41 responses, with a participation rate of 6.8 percent, or 34 companies.

As indicated from the literature review, if two-thirds of the Fortune 500 companies had AI projects staffed and under way in late 1986, then a reasonable expectation would be a similar proportion in the returned questionnaires. Of the 34 companies that responded, 23 had expert systems, a relative frequency of 67.6 percent.

Demographics. Of the companies that participated in the survey, 14 ranked in the top 100 of the Fortune 500. Twenty-two of the 34 participants, 64.7 percent, ranked in the top half of the Fortune 500. The mean ranking, using the 1986 Fortune 500 listing, was 195.68th out of 500. The average change in ranking from 1985 to 1986 was -10.52. In other words, those corporations participating in the survey rose an average of 10.52 places in the 1986 Fortune 500 rankings.

Although there was a bias toward respondent corporations being in the top half of the Fortune 500, 64.7 percent, based on the 1986 rankings, the 1985 rankings indicated a more evenly distributed cross section. Although 37.5 percent of the respondents fell in the top 100, 21.9 percent fell in the 101 to 200 range, 12.5 percent fell in the 201 to 300 range, 12.5 percent in the 301 to 400 range, and 15.6 percent were in the bottom 100 rankings. The questionnaire was completed by individuals with a wide range of job titles. Respondents ranged from technicians to managers, directors, to the chief executive officer of the firm. Five of the respondents held jobs in information

systems, five had job titles directly related to expert systems or AI, and four were scientists or engineers. Thirteen respondents were managers and six were directors. One respondent was an assistant to the corporate vice president and one respondent was the chief executive officer.

Descriptive Analysis

A narrative summary of the findings based on the original research questions is presented.

Research Question 1. How much should expert systems cost the Air Force for acquisition?

Only two responses to the survey question asking for costs of acquisition were received in dollar terms. One respondent reported an acquisition cost of 1000 dollars per system and the other reported a cost of 15,000 dollars. When costs are compared to rule size, the cost of these systems are approximately 25 and 80 dollars per rule, respectively. Both systems ran on personal computers using an expert system shell and averaged between 80 and 100 rules for the first system and between 170 and 350 rules for the second system. Three responses were received in terms of man-months or man-hours, 24 and 36 man-months and 80 man-hours. The average number of man-months was 30 for the two responses, 24 and 36. The 80 man-hour response translates into 4 man-hours per rule and the other two systems approximate an acquisition cost of .048 and 7.68

man-hours per rule, respectively. The first two systems ran on a minicomputer and the other ran on a personal computer and included learning time for the expert system shell. The first two systems used KEE and Fortran or COBOL, respectively.

Forty-two percent of the respondents indicated that their expert system was hosted on a personal computer (PC). AI workstations comprised 22 percent of the responses, closely followed by minicomputer based expert systems, at 20 percent. Mainframe based expert systems accounted for only 13 percent of the total responses. One respondent indicated an expert system based on a special architecture.

LISP was the most common language in which the expert system was written, but accounted for only 12.2 percent of the total responses. Expert system shells, such as M.1, S.1, Insight, and others, accounted for 59.6 percent of the responses. Higher order programming languages, such as Fortran and COBOL, accounted for only 7.0 percent of the total. LISP and OPSS together accounted for 22.8 percent of the responses as general AI programming languages.

The most common response to the question about external links to the expert system reflected links to a data base or data bases. Almost 29 percent of the responses involved a link to one or more data bases. Links to existing hardware and software and no links were the next most common responses, at 18.4 percent and 15.8 percent, respectively. Communications links were next with 7.9 percent.

Thirty-two respondents answered the question concerning whether the company's expert system was developed in-house or contracted. Of the total responses, exactly 75 percent were that the company had developed the expert system totally in-house. Only 6.25 percent said the development was completely contracted out. The remaining respondents said their expert system was a mixture of in-house and contracted development. Eighty percent of this latter group of respondents gave approximate percentages for the amount of in-house effort versus contracted effort. The responses ranged from 80 percent down to 10 percent for organic development, with an average of 45 percent of the systems having been developed in-house.

Of the 25 respondents who said that their expert system was developed completely in-house, only 8 percent were developed by the actual end users. A service organization within the company developed the expert system, according to 40 percent of the respondents, and 52 percent said that the development was completed by a combination of both end user and service organizations. Of the seven respondents that reported a mixture, three reported that end users developed 30 percent of the system, two said that end users did 50 percent, and two respondents said end users do 80 percent of expert system development in the company. One respondent said that 95 percent of expert systems development was done by end users in his firm.

There were 23 responses to the question on the size of

the expert system, including the size of the associated data base. Eighty-seven percent of the respondents said that their expert system contained 500 rules or frames or less, and required 640 kilobytes of memory or less. One respondent gave a response of 2 to 16 megabytes, one respondent claimed a 20,000 rule system, and two respondents did not know the size of the system. There were five separate responses for the size of the data base, three in the megabyte range and two in the gigabyte range.

In response to the related question about the average size of a rule or frame, 61.9 percent of the respondents gave an answer between 3 and 10 conjunctions/disjunctions per rule or slots per frame. Thirty-three percent of the respondents gave a response of unknown, NA, or "?". One respondent said the size "varies".

Twenty-two respondents answered the question concerning development time of the expert system. Of these, 27.2 percent said that development time was six months or less and 22.7 percent said that development time was between seven and 12 months. Twenty-seven percent said development time was between 13 and 24 months and 13.6 percent claimed between 25 and 36 months. Thirteen of the respondents also answered the question about system size in rules. When system size was divided by time to develop, the mean time to develop an expert system was 14.68 rules per month, with a standard deviation of 11.17 rules.

There were six responses to the question on the survey

regarding hindsight and reduction of the cost of acquiring expert systems in the planning phase. Those responses were:

1. Better selection of problem.
2. Through different knowledge representation techniques, better user interfaces, etc.
3. In-house knowledge engineering capability.
4. Had we waited until commercial tools had become more mature, upgrade costs could have been reduced and training might have been better.
5. If people trained in expert systems had been available, training and design implementation could have been improved.
6. Hard to say. Not sure costs could be reduced on the first application.

Research Question 2. How should field performance of the expert system be specified?

There were 26 responses to the survey question on the results or advantages of the company's expert systems. Respondents were given three examples, percent increase in productivity, personnel reduction, and cost savings. Fifteen and one half percent of the responses were related as percentage increases in productivity, such as a 200 percent productivity increase and an 80 percent personnel reduction. Forty-six percent of the responses were couched in generalities such as "better problem solution", "enhanced product performance", and "increased quality". The remaining 38.5 percent were responses such as "too early to predict", "unknown", "difficult to quantify", or "still learning".

There were only ten responses from nine companies to the survey question asking how the expert system's performance was measured. Two of the responses involved accuracy: accuracy of advice and accuracy of output. Three responses

were negative: "not measured", "a non-issue...", and "not applicable". The remaining responses were "improved system utilization", "ease of use", "monitoring work samples", "compare to other statistical (fielded) systems", and "the only criteria is does it make money?".

Research Question 3. What are the best ways to verify and validate expert systems?

Ten respondents answered the verification and validation question. Two of the responses stated that test cases were used. One respondent said that actual use in problem diagnosis was her company's approach while another respondent said that his company's procedure was comparing the expert system's results against the results of other methods. Four respondents said their procedures involved user inputs or testing. One respondent reported a procedure involving signatures and the remaining respondent did not know what verification and validation procedures were being used by her company.

Research Question 4 What documentation is required for fielded expert systems?

There were 16 responses to the survey question on established documentation for the respondent's expert system. User manuals, guides, operating procedures, or instructions constituted 37.5 percent of the responses. Specifications, program listings, and a decision tree diagram comprised 25 percent of the responses. No documentation was used by 12.5 percent of the respondents.

Only seven responses were received for the second part

of the documentation question, which asked for a response as to whether the company's documentation was sufficient. Five out of seven said that the documentation was sufficient. However, one company reporting on-screen documentation as sufficient, also reported a problem with the user interface. Another company reporting insufficient documentation also reported a problem with the user interface. One of the remaining responses stated that "paper documentation (vs. online) is almost impossible" and that no documentation had been approved yet in that firm.

Research Question 5. Who will maintain expert systems for the organization?

Eleven responses were received for this question. Nine out of eleven, almost 82 percent, of the responses were that maintenance was done totally in-house. One respondent said that maintenance was completely contracted out and the remaining respondent said that his company contracted 90 percent of the maintenance.

In answer to the query as to why maintenance was handled in the manner reported above, there were eight responses. The two respondents that reported contract maintenance gave lack of in-house expertise and knowledge engineering capability as the rationale for contract maintenance. The responses for the totally in-house maintainers were cost, the retention of a competitive edge, the ability to have total control over the system, the development of in-house expertise, the desires of the user, and the fact that the

system was developed in-house.

Eleven respondents answered the question concerning who modifies the expert system. A distinction is made here between maintenance and modification. Maintenance refers to making changes to the expert system to insure a level of performance. Modification refers to updates made to the system to reflect policy changes or enhanced capability. There was a one-to-one correlation of contract and in-house maintenance and modification of the system. The nine in-house maintainers stated that in-house knowledge engineers or domain experts modify their systems. The contractor modified 100 and 90 percent of the other two systems, respectively. Regarding lessons learned in the maintenance area, the respondents made five recommendations or observations:

1. Keep an audit trail of up-dates and validate the system after changes.
2. A more straightforward design and simpler coding equates to 100 percent in-house maintenance
3. You cannot "just add a rule" to add more knowledge.
4. The domain experts are never satisfied.
5. Maintenance of the expert system can be accommodated more readily than conventional software systems.

Research Question 6. Should the responsibility for expert systems be centralized or decentralized?

There were eleven responses to the survey question inquiring as to the existence of a central organization within the company with primary responsibility for acquiring, fielding, and maintaining expert systems. Seventy-one percent of the companies responding said there

was no such organization within their company. Only one of these companies said that such an organization was planned. Of the three respondents with a central organization, two reported the size of their organization as 80 and 1, respectively.

The other survey question in the centralization of responsibility area asked for lessons learned as to the necessity of centralized versus decentralized control of expert systems within the respondent's organization. There were eight responses to this question. Of the eight, three respondents were positive towards user or decentralized control. One positively favored centralized control to preserve the integrity of the system. One respondent said that decentralized control meant loss of economies of scale and duplication of training effort. One of the respondents believed both approaches had merit, one said his company had not addressed the issue, and one respondent observed that her AI group had limited success in selling itself in some areas.

Research Question 7. What are the organic personnel requirements for maintenance of an expert system?

Twenty-three companies answered the survey question that asked for the number of people on the company staff devoted only to expert systems. The mean size of the staff was 6.85, with a standard deviation of 11.0, and a range of zero to 50. Fifty-seven percent of the responses were between .5 and 3 for the staff size. Only 13 percent reported a staff

size of zero. The median staff size was three people.

Research Question 8. How will software deficiency reporting be managed?

Eight respondents answered the survey question on how software deficiency reporting is handled within the respondent's company. Fifty percent of the respondents reported some type of user responsible procedure, such as user reports to the design engineer, or user reports to the responsible party at corporate headquarters. One respondent indicated that no formal system was planned for deficiency reporting. Another respondent said that deficiency reporting was not done except locally. The remaining responses were: "verbal" and "request forms to EDP".

Research Question 9. How often should expert systems be modified?

Nine companies responded to the survey question on how often their expert system was modified. Three of the respondents, 33.3 percent, reported continuous or constant modifications, another 44.5 percent reported as required type modifications, and the remaining 22.2 percent reported modifications for policy changes or to meet expanded functions.

Two basic types of modifications were determined from the surveys: planned and unplanned. Planned changes are defined here as those charges not immediately critical to system performance. Unplanned modifications are those necessary to maintain system performance. Unplanned modifications were characterized by words such as constant

and continual, or phases such as whenever necessary, as required, or whenever a deficiency is reported. Based on the reported results of the surveys, almost 80 percent of modifications are unplanned.

Research Question 10. How will configuration control be maintained?

This question was on the survey as a three part question. The first part asked how configuration control was handled in the respondent's company. The second part asked if configuration control was a problem. The third part requested recommendations for configuration management of expert systems. Eight companies responded to individual parts of the question; however, one company's response was "no comment" to all three parts of the question.

Three responses were given for part one. One respondent reported that the user had configuration control. Another respondent said that her company needed to address the issue and the other respondent said that he did not understand configuration control.

Four respondents reported no problems with configuration control. The four companies together had a mean of 13.5 months operational experience with expert systems. One of the four indicated one month of operational experience and the other three had a mean of 17.6 months of experience with operational expert systems. One company said yes, there was a problem, and that the company was studying conventional configuration control products.

There were no responses to part three of the question.

Research Question 11. How much should expert systems cost the Air Force for maintenance?

There were seven responses to the question as to the annual maintenance costs for the respondent's expert system. One respondent said 1000 dollars; one said 100 to 200 dollars per year, including upgrades; two respondents gave 18 man-months as their response; two respondents said "NA"; and one respondent said that the "end users are maintaining it". For one of the companies, 18 man-months meant an average of .036 man-months per rule spent in maintaining the system. The other company reporting 18 man-months did not reveal the size of the expert system.

Eleven companies responded to the question concerning how long the system had been operational. Just over 27 percent had been operational for six months or less. Slightly less than 46 percent had been operational for between 13 and 24 months, and another 18 percent reported operational systems between 25 and 48 months. One respondent, the remaining 9 percent, reported an operational system for 10 years. The average length of time for an operational system for the responding companies, excluding the 10 year response, was 18.3 months with a standard deviation of 16.3 months. The median operational time was 15 months.

Ten companies reported a total of 3329 users, for a mean of 332.9 users per company and a standard deviation of

939.2. The median response was 9.5 users per company.

Responses ranged between one and 3000.

The next question in this area asked the respondent to report whether the expert system had grown in size, the original size if it had, and whether the inference engine had changed. Ten respondents answered the first part of this question, 11 respondents answered the last part of the question. Eighty percent said yes, the system had grown. Responses ranged from a 50 percent growth to 300 percent growth in the number of rules or rule-equivalents. Seventy-three percent of the respondents said the inference engine had not changed.

Respondents were then asked if the expert system had been moved from one hardware or software system to another and if any lessons had been learned. Those that said yes, the system had been moved or a move was in progress, comprised 50 percent of the ten responses to this question. Of the five respondents that said yes, three reported hardware moves and two reported software moves, from LISP to KEE and from Prolog to "C". The lessons learned were:

1. Double the time given by the vendor to convert from LISP to KEE.
2. The programming language "C" is much faster than Prolog.
3. Concerning a particular hardware change, "the system should have been redesigned".

The last question in this area was survey question 28, which was also discussed under research question one in this chapter. The question asked the respondent to use hindsight

and recommend ways of reducing maintenance costs in the planning and implementation phases. The pertinent recommendations were that in-house knowledge engineering capability, better user interfaces, and more mature communications products would have lowered upgrade costs.

Research Question 12. What specific lessons can be learned from organizations that have successfully fielded expert systems?

The first question in this area concerned any difficulties in scaling the expert system up from the prototype to a fully operational system. Twenty-five different responses were received from 19 companies on this question. Seven respondents said there were either no difficulties or no major difficulties. Five responses dealt with performance, i.e., too slow, "performance", and "real time processing". Four responses related to knowledge engineering or software problems, two of which mentioned LISP as part of the problem. Two responses mentioned user interfaces as a difficulty and one response dealt with a hardware problem. Three responses were "not applicable", in process, or "not scaled up yet". Two responses dealt with difficulties connecting the expert system to the data base.

The last question asked the respondent to relate any drawbacks to the expert system's application. Nineteen responses were received to this question. Six respondents reported no drawbacks to the expert system. Six responses were related to performance of the system, four related to knowledge engineering problems, one response dealt with the

difficulty in maintaining the system, and one response was a complaint that the system was not exercised enough. One respondent stated that the development of the expert system took too long.

Interview Results

Four questionnaire respondents and one person from a non-responding company identified from the literature review were interviewed by telephone between 7 and 14 August 1987. A narrative summary of the interview results based on the original research questions is presented.

Research Question 1. What should the acquisition costs be for expert systems?

One respondent stated flatly that this information was proprietary and could not be disclosed. Another respondent was evasive and did not answer the question. The remainder of responses to a direct question on costs were 1000 dollars, 60,000 dollars, and 12 man-years. Based on other information supplied by the literature or questionnaire on the size of their systems, the cost per rule or rule-equivalent was computed as \$25, \$70, and 0.24 man-years per rule, respectively. The respondents also indicated that the latter two systems were first systems for the company and that the development time was much longer for the first system.

Persons interviewed were also asked how the costs of acquiring and maintaining the systems were justified and how the cost benefits were identified. The general response was

that return on investment or payback was the key. One respondent stated that no immediate payback was expected or computed but that the capture of expertise was expected to pay dividends in the future. Answers were phrased in general terms such as "reduced number of problems", the amount of expert system useage, the number of orders handled by the system, and so forth.

When asked to speculate on how costs could have been reduced in planning and implementing an expert system, most of the respondents indicated that this question was hard to answer. Responses ranged from having in-house knowlege engineering capabilities to doing a better job of transferring technology, to better education of the user, user maintenance, and using a different knowledge representation technique.

Research Question 2. How should field performance of the expert system be specified?

Respondents were asked to elaborate on results or advantages of the expert system. Responses were consistency of results, a 1500 percent return on total cash effort, a 75 percent reduction in time required to accomplish a task, and reduced downtime of equipment.

When asked how these results were measured, one respondent replied that his company used to measure the results but that now the system was an integral part of the overall system and that the performance of the expert system is accepted. Another respondent said that the real cost is

compared to the estimated cost by not using the expert system and the difference is the cost-savings. The other respondents indicated that measurement of results was difficult to quantify.

Respondents were also asked how the system's performance was measured and how often formal evaluations were repeated. One respondent said that the only criterion was if the system made money. Other responses indicated that if the system does the job it is supposed to do, then the performance is acceptable and this is a measure of success. One respondent went further and said that the number of users on the system was a measure of performance. Only one person interviewed said that formal evaluations were done and that these evaluations were performed with each new release of the expert system software.

Research Question 3. What are the best ways to verify and validate expert systems?

Respondents were asked to explain what verification and validation procedures work best for their company and how does the company insure that verification and validation is done correctly. Three of the five people interviewed indicated that the verification and validation procedure was a cooperative effort between users and developers and that test cases were used to validate changes. One respondent said that after the initial testing, the system was subjected to "trial by fire" and that there was no conscious effort made at creating exhaustive test cases.

Research Question 4. What documentation is required for fielded expert systems?

The respondents were asked what documentation had been established for their expert system and if the documentation was sufficient. One respondent said that there was no documentation because the expert system development tools used were self documenting. On-screen documentation was the choice for two more of the respondents while the other two said that user's manuals, guides, specifications, and/or test plans were used. All those interviewed thought the documentation was sufficient but one respondent said that perhaps too much documentation existed at her company.

Research Question 5. Who will maintain expert systems for the organization?

Of the five companies interviewed, four maintain their expert systems totally in-house, while the other company has contracted maintenance only. The company that contracts maintenance indicated that time constraints and lack of in-house expertise lead both to the initial contract development of the expert system and the subsequent maintenance requirement. The other companies were asked a series of questions on in-house maintenance.

All those interviewed indicated that in-house development and maintenance was the desired goal of the company. One respondent explained that the expertise was in the company, so the development and maintenance should be, too.

Research Question 6. Should the responsibility for

expert systems be centralized or decentralized?

The respondents were asked if a central organization existed within their company with responsibility for expert systems, the size of the staff, their duties and backgrounds, and lessons learned about the necessity of centralized control.

Four of the five companies have a centralized organization with primary responsibility for the development of expert systems although one company sees its central organization as a training base and catalyst for expert system development. The other firm follows an overall policy which espouses decentralized control throughout the company. The size of this organization ranged from one to 35. Duties of the staff were researching and developing, coding, interfacing, and problem solving. Backgrounds ranged from computer scientists to engineers, with one company reporting no formal AI experience on the staff.

The responses on the necessity of centralized control were varied. One company's philosophy is "extremely decentralized". Another company preferred the centralized approach. The other three companies provided an "it depends" type answer to the question. One said that centralization depends on the application. A large system affecting the entire corporation should be under centralized control but a smaller application with local impact could be under decentralized control. The respondent also indicated that centralized maintenance was not desirable. Another

response was that control should be where the knowledge is, and centralized at that location. The company with decentralized control because of company policy indicated that there could be some advantages to centralized control due to duplication of effort and availability of scarce resources. Further, centralized control could be more closely attuned to the user.

Research Question 7. What are the organic personnel requirements for maintenance of an expert system?

The respondents were asked who maintained the system, was it a full time job, did the maintenance person have another job, and did the same maintenance person have responsibility for more than one system. All four of those interviewed indicated that the systems were maintained by users. Three of the four indicated definitely that the maintenance was only part of the person's job. Half said that the maintainer had responsibility for only one system and the other half said that the responsible party might maintain more than one system.

Research Question 8. How will software deficiency reporting be managed?

One company has established a problem hotline and requires a regular monthly user report to evaluate problems and to verify that the system is being used. The others interviewed do not have formal reporting procedures. The user reports the deficiency to a responsible party such as a design engineer or the expert system "owner" in these organizations.

Research Question 9. How often should expert systems be modified?

The consensus of opinion on this question was that the system should be modified whenever a deficiency is detected or "as required".

Research Question 10. How will configuration control be maintained?

The consensus of opinion in the area of configuration control was that configuration control was not a problem, at least not yet. One company said that the potential was there for a major problem while another said that configuration control was "not a big deal". Three of the five companies do not have a configuration management system. The other two corporations indicated that configuration control is maintained through limited releases of new versions of the expert system software. Lessons learned in configuration management were:

- a. users must be heavily involved.
- b. keep management informed of changes.
- c. the more knowledge gained about the problem, the more problems encountered in maintaining the system.

The last lesson learned refers to a peculiarity of expert systems noted by both one of those interviewed and one questionnaire respondent. The peculiarity is that as more of the problem is analyzed, more knowledge is gained. As more is known about the problem, more rules are needed in the expert system. New rules affect old rules so that, as mentioned in the questionnaire results, "you can't just add

a rule" to add new knowledge.

Research Question 11. How much should expert systems cost the Air Force for maintenance?

The respondents were asked for an opinion as to the annual maintenance costs for an expert system. The responses were 100-200 dollars per year; 1.5 million dollars per year; depending on the size and application, 1-2 man-months or 1-2 people full time; and 1 part-time person. The 100 to 200 dollar per year costs were for PC hosted systems developed using a shell, averaging 80 to 100 rules in size, and maintained by the domain expert. The system that costs 1.5 million dollars annually to maintain is a 10,000 rule system written in OPS5, hosted on a Vax minicomputer, and maintained by a team of knowledge engineers. When the dollar costs are divided by the number of rules in the system, the dollar cost per rule could range from \$1.25 to \$2.00 per rule for a personal computer hosted system, to \$150 per rule for a very large, mainframe based system.

Research Question 12. What specific lessons can be learned from organizations that have successfully fielded expert systems?

Lessons learned were the main reasons for conducting interviews in addition to conducting a questionnaire survey. Additional questions, not on the original questionnaire, were placed on the interview guide to ask during the interview process. Those questions are covered in this section.

Two additional lessons learned in scaling up the expert

system from the prototype to the operational system were suggested by one of the persons interviewed. The lessons learned were to pay attention to the performance issue and to insure more user involvement during the development phase.

Respondents were asked how their companies decided on applications to pursue. The following criteria were suggested by those interviewed.

- a. Survey for needs.
- b. Look for an application where the technology can be applied.
- c. Look for long range problems to solve, not problems that must be solved next quarter.
- d. Look for retiring expertise and try to capture that expertise.
- e. Pick projects that will be conspicuous.
- f. Look for applications that will produce a finite payback, increased productivity, and/or enhanced product line.
- g. Make sure the scope of the project is bounded.
- h. Let the domain expert pick his or her own problem, where the domain expert is also the developer of the expert system.

The respondents were also asked if any cost benefit analysis was done to decide on which applications to pursue and, if so, what factors were considered in doing the analysis. One corporation said that about three minutes of "back of the envelope" analysis was done. Two firms said that they had tried it but that gross assumptions had to be made and that quantification of benefits was difficult. Another firm looked at previous experiences with the software, the existing hardware, and the training of the personnel involved. The other corporation reported that no

cost benefit analysis was done; the company simply looked for the best match between need and technology for feasible applications.

The corporations were also asked how long it took to get from a prototype to a fielded system. Two firms stated that the time depends on the size of the system, the problem, the domain expert, and the development tool. The actual time estimates ranged between 1-2 months on a personal computer using a shell, to six months contracted development on a personal computer, to between 12 and 18 months on an AI workstation.

Summary

This chapter has reported the findings of the survey instruments. The findings were discussed in relation to the research questions originally posed in Chapter I. Summary tables of the results by survey question may be found in Appendix B. Chapter V will discuss the conclusions of the study and recommendations for further research.

V. Summary, Conclusions, and Recommendations

Chapter Overview

The purpose of this chapter is to summarize the research effort, present the conclusions, and recommend areas for further research. Conclusions are discussed in response to the research questions originally posed in Chapter I.

Significance of the Research.

Although the questionnaire return rate was less than anticipated, the low rate of response lent credence to Harvey Newquist's statement that companies are very secretive about expert systems because they use the technology for strategic advantage (1:54). The effects of a low response rate were minimized by conducting a census of the target population, the Fortune 500 corporations. Conclusions are assumed to reflect the population in general.

The research effort clearly indicates that corporations are still learning the answers to many of the support issues for expert systems. Business and industry, although more experienced than the DOD in fielding expert systems, still have very few answers.

However, the questionnaire and interview guide, developed to probe the lessons learned in fielding expert systems, provide the foundation tools for further research in this area. These tools are believed to represent a

significant contribution to the field.

Summary of the Research Effort

The last few years have seen a rapid growth in the use of expert systems in business and industry. The private sector views expert systems as a means of enhancing productivity, extending rare in-house expertise, and insuring strategic advantage. The Air Force also views the technology of expert systems as a means to more effectively manage resources and to enhance productivity. However, most of the experience gained in fielding expert systems exists in business and industry.

The purpose of this research was to capture lessons learned in fielding expert systems and to apply those lessons learned to plan the support of expert systems. Specifically, the support issues associated with expert systems such as the acquisition, maintenance, documentation, organizational, performance, personnel, and software modification issues needed to be addressed. Research questions were developed to address the support issues and a review of the literature was conducted to obtain the data to answer those questions.

Twelve research questions were investigated:

1. What should the acquisition costs be for expert systems?
2. How should field performance of the expert system be specified?
3. What are the best ways to verify and validate expert systems?

4. What documentation is required for fielded expert systems?

5. Who will maintain expert systems for the organization?

6. Should the responsibility for expert systems be centralized or decentralized?

7. What are the organic personnel requirements for maintenance of an expert system?

8. How will software deficiency reporting be managed?

9. How often should expert systems be modified?

10. How will configuration control be maintained?

11. How much should expert systems cost the Air Force for maintenance?

12. What specific lessons can be learned from organizations that have successfully fielded expert systems?

Since the information necessary to answer the research questions was not found in the literature, a survey methodology was decided upon to gather the required data. Initial data was gathered through the use of a questionnaire. Further data was collected by conducting telephonic interviews. The target population was the Fortune 500 corporations because of the perceived successfulness of those companies. A census was undertaken because of the relatively small size of the target population and to maximize the return of the questionnaire. Questionnaire data were collected from mid-May 1987 until 4 August 1987. Telephonic interviews were conducted between 7

and 14 August 1987.

The data were analyzed and summary statistics such as frequency of response, mean, median, and standard deviation were computed for each question involving numerical data. Frequency of response and the modal response were statistics computed for non-numerical responses. The findings were presented in Chapter IV.

Conclusions

The conclusions are presented with emphasis on the original research questions.

Research Question 1. How much should expert systems cost the Air Force for acquisition?

Conclusion 1. The cost of an expert system hosted on a personal computer (PC) and using a commercial shell could range from 25 to 80 dollars per rule or rule equivalent.

The results of the survey for other hardware hosts, such as the minicomputer varied too widely to permit a conclusion to be drawn for those systems. No cost data were received for mainframe hosted expert systems.

Conclusion 2. Personal computers appear to be the most common hardware host for expert systems. Personal computers were twice as common as either workstation or minicomputer hosted systems. The survey indicated that PCs were in the majority, at 42 percent. Further, indications from the literature and the Dupont and UNISYS visits are that PC hosted expert systems are growing in popularity.

Conclusion 3. Most expert systems are developed using

expert systems shells. This conclusion complements the previous conclusion. Development time is much less using a commercial shell because the time to develop a custom made development tool is eliminated. The size of expert systems that can be developed on a PC is somewhat limited. However, this restriction is offset by the fact that many applications do not require more memory than that afforded by a PC.

Conclusion 4. The most important external links to a expert system appear to be links to data bases, existing hardware and software, and communications equipment. Links to data bases, particularly mainframe data bases, and to existing hardware and software were mentioned most often in both the questionnaire and the interview. Some expert systems are desired solely because of the problems associated with searching large data bases, typically hosted on a large mainframe computer. The ability of an expert system to interface with existing hardware and software is very important to companies with large investments in existing hardware and software.

Conclusion 5. Almost all expert systems are developed in-house by a service organization or a combination of a service organization and end users. Ninety-two percent of the respondents on the survey said their expert systems were developed this way.

Conclusion 6. The majority of expert systems are 500 rules or less and use 640 kilobytes of memory or less.

Eighty-seven percent of the survey respondents reported that their systems fell within those parameters.

Conclusion 7. The average size of a rule appears to be between 3 and 10 conjunctions/disjunctions per rule.

Although a wide range of rule sizes resulted from the survey, almost two-thirds, 62 percent, of the rule sizes fell within this range. The wide range of response was expected, as each expert system has different requirements, programmers, and so on.

Conclusion 8. The time required to develop most expert systems is 12 months or less. Half of the companies reported development times of 12 months or less. Dupont, the most prolific of the companies developing expert systems, reports an average development time for a PC hosted expert system, developed using a shell, and composing about 80 rules, as one to two months. Obviously, a much larger system will take much longer to develop and some companies are faster than others at developing expert systems.

Conclusion 9. A delivery system in LISP almost guarantees that in-house development will be very expensive and prohibitive to end user maintenance. LISP requires more expensive hardware to develop the expert system and personnel with more technical expertise to maintain the system, once developed. This level of expertise was noted to be above that which would be expected in a normal user in business.

Research Question 2. How should field performance of

the expert system be specified?

Conclusion 10. Companies are still learning the answers to the performance issues. Only 15 percent of the companies were able to describe the results or advantages of their expert systems in quantifiable terms, such as a 200 percent productivity increase. The great majority, almost 85 percent, could not quantify the results.

Conclusion 11. Minimal acceptable performance should be specified in advance. It appears from the results of the survey instruments that minimal acceptable performance is not specified in advance. It would follow then, that the corporations do not have a clear idea of the results or advantages of their systems because of a failure to specify what a successful system will be able to do for the company.

Research Question 3. What are the best ways to verify and validate expert systems?

Conclusion 12. The best verification and validation procedures involve user inputs and testing. Based on the results of the survey instruments, it appears that the preferred method of verification and validation involves some form of testing and user inputs.

Research Question 4. What documentation is required for fielded expert systems?

Conclusion 13. Corporations are still learning the answers to the documentation issues. User manuals, guides, operating procedures, instructions, specifications, program listings, and a decision tree diagram comprised the majority of responses to the survey instruments but there was no

clear preference for any particular form of documentation. The wide range of responses could indicate that companies are still experimenting to find the best way to document the expert system.

Research Question 5. Who will maintain expert systems for the organization?

Conclusion 14. In-house maintenance appears to be the preferred source for expert system maintenance. Almost 82 percent, of the responses were that maintenance was done totally in-house. The rationale for in-house maintainers was based on cost, the retention of a competitive edge, the ability to have total control over the system, the development of in-house expertise, the desires of the user, and the fact that the system was developed in-house. Lack of in-house expertise and knowledge engineering capability was the only rationale for contract maintenance.

Research Question 6. Should the responsibility for expert systems be centralized or decentralized?

Conclusion 15. Control of expert systems should be centralized at the same location where the knowledge is located. If the knowledge exists at corporate headquarters and the expert system affects the entire company, then control should be at corporate headquarters. Locally developed systems with only local effects, should be under decentralized control.

Conclusion 16. Responsibility for training personnel to develop and maintain expert systems should be centralized. Centralized training avoids duplication of effort, affords

economies of scale in both funding and training effort, and provides standardization of training, including company policies concerning expert systems.

Conclusion 17. Maintenance of expert systems should be decentralized as much as possible. Expert systems should be maintained at the level where the knowledge resides. If the knowledge required to maintain the system is located at corporate headquarters, then the maintenance should be located at corporate headquarters. Some corporations indicated that the goal of their company was end user maintenance. This goal appears to be in harmony with user involvement and acceptance policies or goals.

Research Question 7. What are the organic personnel requirements for maintenance of an expert system?

Conclusion 18. A personnel requirement of between one and three people per system would be expected. More than half of the respondents had staffs between .5 and 3 people and the median staff size was three people.

Research Question 8. How will software deficiency reporting be managed?

Conclusion 19. Corporations are still learning the answers to the software deficiency issues. Although fifty percent of the respondents reported that in their company the user reports to a responsible party, almost none of the companies had established any type of formal reporting procedure. Informal procedures, such as user reports to design engineer, or user reports to responsible party at corporate headquarters were in use but the majority of

companies appear to have not yet reached a stage where they feel that a more formal system for deficiency reporting is needed.

Research Question 9. How often should expert systems be modified?

Conclusion 20. Expert systems should be modified whenever a deficiency is reported and on an as required basis. The majority of modifications, 78 percent, were found to be unplanned modifications that have to be done to maintain system performance. The remaining modifications were for policy changes or to meet expanded functions, which could be classified as planned changes. Planned changes may be accomplished at periodic intervals but can also be done on an as needed or an as required basis.

Research Question 10. How will configuration control be maintained?

Conclusion 21. Companies are still learning the answers to this question. From the small number of responses and the types of responses to the survey question, companies apparently have not yet had time to address this issue. The companies responding to the question only had an average of 13.5 months operational experience with expert systems, which may not be enough time for problems with configuration control to manifest themselves.

Research Question 11. How much should expert systems cost the Air Force for maintenance?

Conclusion 22. Annual maintenance costs could cost from one to two dollars per rule for a PC hosted system to 150

dollars per rule for a mainframe hosted system. The disparity in costs is explained partially by the size of the systems involved. The PC system averages 80 to 100 rules and is frequently maintained by the domain expert. The mainframe system is 10,000 rules and is maintained by a team of knowledge engineers. There are other factors involved but the maintenance costs will generally depend on the application. The rule size is a convenient means of comparison.

Conclusion 23. Most expert systems are based on the original inference engine. Based on the results of the survey, the great majority, 73 percent, of expert systems are still based on the original inference engine.

Conclusion 24. The delivery system should not be in LISP if speed of program execution is important. LISP was singled out by respondents as a problem area for applications where speed of program execution was important.

Research Question 12. What specific lessons can be learned from organizations that have successfully fielded expert systems?

Conclusion 25. There were several lessons learned in fielding expert systems. Some of the lessons learned are presented.

1. Development.
 - a. Heavy user involvement is a must from the beginning.
 - b. Top level management support is a must from the beginning.
 - c. Survey for needs.
 - d. Look for an application where the technology can be applied.
 - e. Look for long range problems to solve,

not problems that must be solved next quarter.

f. Look for retiring expertise and try to capture that expertise.

g. Pick projects that will be conspicuous.

h. Look for applications that will produce a finite payback, increased productivity, and/or enhanced product line.

i. Make sure the scope of the project is bounded.

j. Let the domain expert pick his or her own problem, where the domain expert is also the developer of the expert system.

k. An expert system without an owner, the person charged with the "duty of care" of the system, is a recipe for failure. -- Dr Ed Mahler, Dupont, 4 June 1987.

l. An expert system is never finished.

m. The user interface is very important to user acceptance of the expert system.

2. Performance.

a. Minimum acceptable performance must be specified in advance (36:198).

b. "LISP is not the language of the angels" (24:38). The expert system tools best for research are not the best tools for operational systems (24:38).

3. Maintenance.

a. In-house maintenance is desirable for cost, security, and control.

b. Maintenance should be as decentralized as possible.

c. The ultimate maintenance goal for an expert system is to be user maintained.

d. Keep an audit trail of up-dates and validate after changes.

e. You cannot just "add a rule" to add more knowledge.

f. Maintenance of expert systems is easier than for conventional software systems.

4. Responsibility and Control.

a. Control should be centralized where the knowledge is located.

b. Centralized or decentralized control should depend on the scope of the application.

5. Software Deficiency Reporting.

a. Software deficiency reporting procedures should be tailored to the organizational structure, desires, and policies.

- b. Deficiency reporting depends upon the application and the level of responsibility for maintenance.
- 6. Documentation. Some form of documentation is necessary for the expert system user.
- 7. Personnel.
 - a. Some staff personnel are needed for training of the users.
 - b. If user maintenance of the expert system is a goal then no additional personnel are needed for maintenance of the system itself.
- 8. Configuration Control.
 - a. Users must be heavily involved.
 - b. Keep management informed of all changes.
 - c. The more knowledge gained about the problem, the more problems are encountered in maintaining the system.

Recommendations for Further Research

There are several areas of expert systems management suggested by this study for further research. The areas recommended are costs of acquisition and maintenance, configuration control, documentation, centralized responsibility, and software deficiency reporting for expert systems.

Due to the observed reluctance of the Fortune 500 corporations to discuss costs, further investigation into the reasons why this information is so sensitive could be very instructive. A hypothesis of the researcher is that the cost is so high, for those companies still using research tools rather than commercial shells to develop non-user friendly applications, that these companies are reluctant to admit a low rate of return on investment.

Another hypothesis suggested by the study is whether the

company with centralized responsibility for expert systems is more successful than the company with decentralized responsibility.

Lack of experience in other areas of this research effort, such as configuration control, documentation, centralized responsibility, and software deficiency reporting most likely contributed to the low response rate to questions concerning those issues. A repeat of the survey is recommended after an interval of one or two years to capture the additional corporate experience.

As the Air Force proceeds into the 1990s, artificial intelligence will play an increasing role in the management of scarce resources. As a subfield of artificial intelligence, expert systems promise the most immediate return on investment. As in most endeavors, lessons learned and applied are mistakes avoided. This research has demonstrated that there are valuable lessons to be learned in fielding expert systems.

Appendix A: The Survey Instruments



Name
Company Name
Street Address
City, ST Zip Code

Dear _____:

In the past few years, knowledge systems have evolved from laboratory experiments in artificial intelligence to fairly widespread use in industry and business. The transition from research prototypes to fielded systems is viewed by many as a major challenge facing managers today. In an effort to meet this challenge, SRA is assisting in a research effort designed to help the U.S. Air Force develop a strategic plan for the introduction and support of expert systems in its organizations.

As the essential first step in this research, we are gathering data from Fortune 500 Companies on their planned and currently fielded expert systems. I am offering you the opportunity to participate in this mutually beneficial research effort by completing and returning the enclosed, Air Force-developed questionnaire. Your responses will be combined with others; they will not be attributed either to you personally or to your company.

Your participation is vital. The results will be provided to the Defense Department to be used to help shape its strategy for introducing, exploiting, and supporting expert systems in the future. We will also provide an Executive Summary of the analyzed responses to any respondent to the survey that asks for it.

If, when you return the completed questionnaire, you can include documentation on your expert systems(s) that would help explain or clarify your answers, it would be greatly appreciated. We are not soliciting proprietary information, but rather we seek to develop a comprehensive picture of what the Fortune 500 Companies are learning in fielding expert systems. It would be a great help to us in meeting our schedule if we could get your response, via the enclosed pre-addressed envelope, by June 12, 1987.

For further information, contact Dr. Mary Dee Harris, at (703) 558-7849. Thank you in advance for your invaluable contribution to this research.

Sincerely,

Sherman Greenstein, Director
Artificial Intelligence Division

LESSONS LEARNED IN FIELDING EXPERT SYSTEMS
A Questionnaire

Please check here if specific details of your work/answers are NOT to be publicly released: _____

SECTION I

Company Information: _____

Company Name: _____

Nature of Product Line: _____

Individual Completing Questionnaire: _____

Position/Title: _____

For the purpose of this questionnaire, EXPERT SYSTEMS are defined as:

A computer system that has stored in it the problem solving knowledge of one or more human experts in a particular field.

1. Given the definition of Expert Systems presented above, does your company have one or more expert systems in use, or under development, to support any in-house activities?

_____Yes _____No

If Yes, how many systems? _____

If No, why not?

2. If the answer to question 1 was yes, please identify the activity and briefly describe [for the system that has been in use/development the longest]:

(a) its functions, (b) its objectives, (c) who uses the system, and (d) any other pertinent information regarding the system.

(a) _____

(b) _____

(c) _____

(d) _____

3. Please identify the Stage of Development, from the four stages defined below, that most closely describes the Expert System(s) you are currently using/developing. Can your system be defined as a:

_____ QUALIFIED SYSTEM?

Definition: A problem has been determined to be a good expert system opportunity but as of yet, no work has been done.

_____ DEMONSTRATION PROTOTYPE?

Definition: The system solves a portion of the problem undertaken, suggesting that the approach is viable and system development is achievable.

_____ FIELDDED PROTOTYPE?

Definition: The system displays good performance with adequate reliability and is being revised based on extensive testing in the user environment.

_____ OPERATIONAL SYSTEM?

Definition: The system exhibits high quality, reliability, fast and efficient performance and is being used on a regular commercial basis.

4. Is the application of an expert system, within your organization, based on a (check all that apply):

Personal Computer (e.g., IBM PC, Apple Macintosh, etc.)?

Large Dedicated AI (Artificial Intelligence) Workstation?

Minicomputer (e.g., VAX)?

Mainframe (e.g., CDC, IBM, etc.)?

5. For each expert system in development use, please identify the shell or programming language that is being used.

6. What software, besides the external, outside of the expert system's database, programs, databases, and/or communications are in use? Please describe.

7. Has your expert system(s) been developed, or will it be developed:

completely in-house

by a service organization or by contract?

by a mix of both the above? (Please indicate percentage of each.)

Completed: _____

8. If your system(s) was developed in-house, was your system(s) developed:

by end users?

by a service organization within your company?

by a mix of both the above?

Please indicated approximate % of each.

9. What is the size of your staff devoted only to expert systems? _____

AD-A187 865


A STRATEGIC PLAN FOR SUPPORT OF EXPERT SYSTEMS IN
ORGANIZATIONS(U) AIR FORCE INST OF TECH
WRIGHT-PATTERSON AFB OH SCHOOL OF SYSTEMS AND LOGISTICS
P R BOGGS SEP 87 AFIT/GLM/LSV/875-6

2/2

UNCLASSIFIED

F/G 12/8

NL





10. Please indicate the SIZE of your Expert Systems(s)
(approximate number of rules or frames being used and the amount
of computer memory required to operate the system in its cur-
rent/planned configuration, including data base):

What is the average size of each rule or frame (i.e., the
number of conjunctions/disjunctions in the rule antecedent or the
number of slots/facets in each frame)?

11. Please indicate the approximate date system development began
and the approximate date the system was (will be) ready to field
test.

Development Began: _____
System Ready to _____
Field Test _____

12. What results/advantages have you experienced from this
application to date? (e.g., percent increase in productivity,
personnel reduction, cost-savings, etc.)

What difficulties, if any, were encountered in scaling the
expert system up from the prototype model to the fully operational
system?

13. Please list any drawbacks to this application:

SECTION II

The remaining sections of the questionnaire are solicited in an effort to see what lessons have been learned by those companies that have actually put expert systems into operation and to elicit "this is the way expert systems should be developed/fielded/supported" comments where appropriate.

*** ANSWER THE FOLLOWING QUESTIONS ONLY IF ***
*** YOU HAVE AN OPERATIONAL EXPERT SYSTEM ***

If section II does not apply to you, thank you for your time and interest.

14. How long has this system(s) been in operational use within your company?

How many end users actually use the system(s)? _____

15. Has this expert system(s) grown in size: (rules/frames/memory requirements) If so, what was the original size?

Has the inference engine been changed? _____

16. Has it been moved from one hardware/software system to another (e.g., hardware: IBM PC to mainframe, etc.; software: LISP to "C", etc.). If so, what were those changes and what lessons were learned?

17. What documentation (e.g., specifications, test plan, user's manual, etc.) has been established for the system(s)? Do you feel this documentation is sufficient? If not, why not?

18. How is the system maintained?

☐ Completely Contracted Out

☐ Partly Contracted Out, Partly In-house (specify approximate
split ____%/____%)

☐ Totally In-house

Why did you choose to maintain the system in this manner? (lack
of in-house expertise, cost, etc.)

19. Is there a central organization within your company with
primary responsibility to oversee the acquisition, fielding, and
maintenance of your expert system(s)?

☐ Yes

☐ No

If so, how many people staff this function? _____

If not, has your company ever had, or planned to have, such
an internal organization?

☐ Yes

☐ No

20. What lessons have been learned by your company as to the
necessity of centralized versus decentralized control of expert
systems within your organization?

21. How is software deficiency reporting for the expert system(s)
handled in your company? (special forms required, who reports,
who receives reports, who acts on reports, etc):

22. How often is the system modified? (e.g., quarterly, whenever
deficiency detected, policy changes, etc.?)

23. Who performs the modifications?

___ Contractor and/or ___ In-house knowledge engineer(s)

What lessons have you learned regarding system updates?

24. How is configuration control of the system being maintained? Is configuration control a problem? What recommendations do you have regarding configuration management of expert systems?

25. What verification/validation procedures of the expert system(s) work best for the company?

26. How is the system's performance measured? (e.g., what criteria, such as accuracy of advice, time to complete a consultation with the system(s), etc., are used to evaluate performance/usefulness?)

27. What were/are the approximate system costs (in dollars or staff-months or both) for:

Acquisition (from decision to buy until system was ready to field test): _____

Annual Maintenance (documentation, 24-hour call-in service, etc.): _____

28. In hindsight, how could costs of acquiring and maintaining the system(s) have been reduced in the planning and implementation phases?

Thank you for your time and interest.

Interview Guide

1. I will be recording this interview as a backup for my notes.
2. Go over questions from questionnaire where responses were left blank or responses were unclear.
3. What were/are your approximate system costs for:

Acquisition (from decision to buy until system was ready to field test):

Annual Maintenance (documentation, 24-hour call-in service, etc.):
4. How do you justify the costs?
5. How do you identify the cost benefits?
6. In retrospect, how do you feel the costs of acquiring and maintaining the system(s) could have been reduced in the planning and implementation phases?
7. What results/advantages have you experienced from this application to date? (e.g., percent increase in productivity, personnel reduction, cost-savings, etc.)
8. How do you measure those results?

9. How is the system's performance measured? (e.g., what criteria, such as accuracy of advice, time to complete a consultation with the system(s), etc., are used to evaluate performance/usefulness?)

10. How often are formal evaluations of systems repeated?

11. How is your system maintained?

____ Completely Contracted Out

____ Partly Contracted Out, Partly In-house

____ Totally In-house

12. Why did you choose to maintain the system in this manner? (lack of in-house expertise, cost, etc.)

13. In-house: Who maintains it, is it a full time job, or does the maintenance person have another job?

14. In-house: Do the same maintenance personnel have responsibility for several different systems?

15. How is software deficiency reporting for the expert system(s) handled in your company? (special forms required, who reports, who receives reports, who acts on reports, etc):

16. How often is the system modified? (e.g., quarterly, whenever deficiency detected, policy changes, etc.?)

17. Who performs the modifications?

____ Contractor and/or ____ In-house knowledge engineer(s)

What lessons have you learned regarding system updates?

18. How is configuration control of the system being managed?

19. Is configuration control (standardization) a problem?

20. What recommendations do you have regarding configuration management of expert systems?

21. What verification/validation procedures work best for the company?

22. How do you make sure the verification and validation are done correctly? (test cases?)

23. Is there a central organization within your company with primary responsibility to oversee the acquisition, fielding, and maintenance of your expert system(s)?

____ Yes ____ No

24. If so, how many people staff this function?

25. What are their duties?

26. What are their backgrounds?

27. If not, has your company ever had, or planned to have, such an internal organization?

____ Yes

____ No

28. If not, why not?

29. What lessons have been learned by your company as to the necessity of centralized versus decentralized control of expert systems within your organization?

30. What documentation (e.g., specifications, test plan, user's manual, etc.) has been established for the system(s)?

31. Do you feel this documentation is sufficient? If not, why not?

32. How long has/have your system(s) been in operational use within your company?

33. How long does it take to get from a prototype to a fielded system?

34. How many end users actually use the system(s)?

35. What difficulties, did you encounter in scaling the expert system up from the prototype model to the fully operational system?

36. Were there any drawbacks to the application?

37. How do you decide on the applications to pursue?

38. Do you do any cost-benefit analysis to decide which applications to pursue?

39. What factors did you consider in doing the cost benefit analysis?

40. Has this expert system(s) grown in size since the the system was first fielded? (rules/frames/memory requirements) If so, what was the original size? What is the size now?

41. Has the inference engine been changed? How? Why?

42. Has your system been moved from one hardware/software system to another (e.g., hardware: mainframe to PC, etc.; software: LISP to "C", etc.)?

43. If so, what were the lessons learned?

Appendix B: Frequency Tables
by Survey Question

SURVEY QUESTION #1.
Number of Expert Systems per Company.

Response	# of responses
0	11
1	6
2	3
3	3
4	5
5	2
6	2
12.5	1
20	1
35+	1
600	1
726.5+	36

no. of companies participating: 34 [Note 1]

no. of questionnaires returned: 37 [Note 2]

mean number of expert systems
per company: 22.9

median number of expert
systems per company: 2

standard deviation: 105.6

Note 1: Two companies returned two questionnaires each.

Note 2: Of the companies not participating,
only one returned the questionnaire.

SURVEY QUESTION #2a.
Expert System Function.

Response	# of responses
diagnosis/troubleshooting	8
simulation	2
cost estimation	1
configuration	2
scheduling	2
inventory control	2
quality control	1
selection of alternatives	1
expert assistant	1
embedded system	1
auto data processing	1
sequencing job orders	1
	24

no. of companies
answering this question: 34

most common use of
expert systems: diagnosis

SURVEY QUESTION #2c.
Users.

Response	# of responses
engineers (various types)	6
analysts (various types)	3
technicians	3
managers	2
mechanics	2
pilots	2
sales personnel	2
scheduling personnel	1
inventory control personnel	1
quality control personnel	1
production control personnel	1
supervisors	1
specialists	1
scientists	1
operating plant personnel	1
chemists	1
petrologists	1
operators	1
marketing	1
hourly workers	1
foremen	1
	34

no. of companies
answering this question: 34

most common user of expert systems: engineers

SURVEY QUESTION #3.
Development Stage.

Response	Company Experience	Number of Systems
Qualified System	3	2009
Demonstration Prototype	8	1014
Fielded System	7	508
Operational System	5	105
	<u>23</u>	<u>3636</u>

no. of companies
answering this question: 23

median stage of expert system
development by a company: Fielded System

Note: The company experience column reflects the most advanced expert system reported by each company, an operational system being the most advanced.

SURVEY QUESTION #4.
Expert System Hardware.

Response	# of responses
Personal Computer (PC)	19
Large, Dedicated AI Workstation	10
Minicomputer	9
Mainframe	6
Special Architecture	1
	45

no. of companies
answering this question: 34

most common hardware base for
expert systems: personal computer

Note: Some questionnaires contained more than
one response to this question.

SURVEY QUESTION #5.
Expert System Programming Language.

Response	# of responses
LISP	7
KEE	6
OPSS	6
ART	5
Goldworks	3
Insight, Insight II+	4
M.1	3
PC+	5
PC Easy	1
TI Personal Consultant	3
COBOL	2
Knowledge Craft	2
Other Proprietary	2
ADA	1
AON	1
ESE	1
Fortran	1
GURU	1
MOBS	1
RIMS	1
S.1	1
	57

no. of companies
answering this question: 23

Note: Some questionnaires contained more than
one response to this question.

SURVEY QUESTION #6.
External Links to Expert System.

Response	# of responses
Data Base(s)	11
Existing Hardware/Software	7
None	6
Communications	3
Mainframe Computer	3
Graphics	2
External Sensors, Devices	2
Library Routines	1
SNA linked to many users	1
Speech Synthesis	1
Text Entry	1
	38

no. of companies
answering this question: 20

Note: Some questionnaires contained more than
one response to this question.

SURVEY QUESTION #7.
Expert System Development.

Response	# of responses
Completely In-house	24
Completely Contracted Out	2
Partly In-house/Partly Contracted Out	6
	32

Percentage of In-house Development
Reported for Partial Response:

80	1
75	1
40	1
20	1
10	1
	5

no. of companies
answering this question: 23

Note: Some questionnaires contained more than
response to this question.

SURVEY QUESTION #8.
In-house Expert System Developed by:

Response	# of responses
End Users	2
Service Organization	10
A Mix of Both the Above	13
	25

Percentage of End User Development
Reported for "Mix" Response:

30	3
50	2
80	2
95	1
	7

no. of companies
answering this question: 23

Note: Some questionnaires contained more than
one response to this question.

SURVEY QUESTION #9.
Size of Staff Devoted to Expert Systems.

Response	# of responses
0	3
0.5	1
1	2
1.5	1
2	3
3	3
5.5	1
6	3
7	1
8	1
10	1
15	1
25	1
50	1
	23

no. of companies answering
this question: 23

mean size of expert systems
staff per company: 6.85

median size of expert
systems staff per company: 3

standard deviation: 11.01

SURVEY QUESTION #10a.
Expert System/Data Base Size.

Response	# of responses
Expert System	
Unknown/NA	2
20r./640K	1
30r.	1
50r.	1
60r./400K	1
75r./640K	2
90r.	1
100r./640K	1
125r.+35 Fortran SBR.	1
150r.	1
200r.	1
250r.	1
260r./640K	1
300r./640K	1
450r.	2
500r.	1
500f.	1
650r.	1
20,000r.	1
2-16 Meg	1
	23
Data Base	
8 Mb RAM	1
5Kb-4Mb	1
400 Mb Disk	1
1.2 Gigabytes	1
"Gigabyte range"	1

no. of companies answering
this question: 20

Note: Some questionnaires contained more than
one response to this question.

SURVEY QUESTION #10b.
Expert System Rule/Frame Size.

Response	# of responses
Unknown/NA	3
4	3
?	3
Varies	1
7-10	1
6 C/D per rule	1
3.6 pat/join	1
10 slots/frame	1
10 conj./rule	1
3.7 ant./2.2 actions	1
7 slots/frame	1
3 conj./disj.	1
5-8	1
5 or 6	1
10-20	1
	<u>21</u>

no. of companies answering
this question: 20

SURVEY QUESTION # 11
Expert System Development Time.

Response	# of responses
1-6 mo.	6
7-12 mo.	5
13-18 mo.	3
19-24 mo.	3
25-36 mo.	3
no completion given	2
	22

no. of companies answering
this question: 22

mean: 15.76
standard deviation: 11.03

SURVEY QUESTION #12a.
Expert System Results/Advantages.

Response	# of responses
not applicable	3
difficult to quantify	1
not quantified	1
still learning	1
unknown	1
too early to predict	1
no evaluation available today	1
no objective data available currently	1
200% productivity increase	1
ability to build simulation	
proving concepts	1
more effective use of expert's time	1
enhanced product performance	1
better problem solutions	1
should lead to lower lost	
time outage costs	1
increased quality	1
identification of data anomalies	
in the data base	1
better understanding of analysis process	1
capture scattered knowledge	1
goal: 80% personnel reduction	1
10 to 20% improvement	1
training	1
proprietary information -	
cost savings encountered	1
1500% return on total cash effort	1
significant increase in diagnosis	
success rate	1
	26

no. of companies participating: 21

Note: Some questionnaires contained more than
one response to this question.

SURVEY QUESTION #12b.
Difficulties in Scaling Up System.

Response	# of responses
none	7
system became slower	1
performance	1
proper user interface	1
meeting the requirements	
of a robust system	1
real time processing	1
some degradation in speed	1
biggest problem--common bandwidth	
between mainframe and workstation	
is too slow	1
hooking expert system to data base	1
moving from custom LISP based	
language to a commercial	
shell took 14 months	1
developing appropriate user interface	1
delivery system will have to	
run on a workstation, not	
a LISP processor	1
software engineering in LISP	
is not so well understood	1
in process	1
greater need for knowledge	
acquisitoion tools to trace	
through knowledge elements	1
more work than anticipated	1
not yet scaled up	1
not applicable	1
no major difficulties	1
	25

no. of companies
 answering this question: 19

Note: Some questionnaires contained more than
 one response to this question.

SURVEY QUESTION #13.
Drawbacks.

Response	# of responses
none	6
hardest part is maintaining	1
none of field people know how	
to estimate manually	1
system did not adequately	
reflect user attitudes	1
first application was much more	
difficult than hoped	1
knowledge acquisition	1
getting data from the main	
frame is too slow	1
performance (soon to be moved	
to "C"	1
expert system not used daily	1
appropriateness of LISP to	
meet requirements of a	
robust system	1
lacks probabilistic features	1
heavy reliance on pattern matching	1
expert system results in a complex	
computer environment	
took too long to develop	1
	19

no. of companies
 answering this question: 16

Note: some questionnaires contained more
 than one response to this question

SURVEY QUESTION #14a.
Length of Time Operational.

Response	# of responses
0.5-6 mo.	3
7-12 mo.	0
13-18 mo.	3
19-24 mo.	2
25-36 mo.	0
37-48 mo.	2
10 yrs.	1
	11

no. of companies answering
this question: 11

mean: 18.3

median: 15.0

standard deviation: 16.3

SURVEY QUESTION #14b.
Number of End Users.

Response	# of responses
1	1
2	1
3	1
4	1
5	1
14	1
20+	1
80	1
200	1
3000	1
	10

no. of companies answering
this question: 10

mean number of end users: 332.9

median no. of end users: 9.5

standard deviation: 939.2

SURVEY QUESTION #15a.
Expert System Grown in Size?

Response	# of responses
Yes	8
No	2
	10

no. of companies answering
this question: 10

SURVEY QUESTION #15b.
Inference Engine Changed?

Response	# of responses
Yes	3
No	8
	11

no. of companies answering
this question: 11

SURVEY QUESTION #16.
Has Expert System Been Moved?

Response	# of responses
Yes	4
In progress	1
No	5
	10

Lessons learned:

Double time given by vendor
to convert from LISP to KEE

Converted from Prolog to "C"
because "C" much faster

Re: hardware change, system
should have been redesigned

no. of companies
answering this question: 10

SURVEY QUESTION #17a.
Documentation.

Response	# of responses
user manual	2
specifications	2
none approved	1
user instructions	1
on screen user guide	1
user guide	2
operating procedure	1
user data	1
program listing	1
complete documentation	
of code	1
decision tree diagram	1
security system	1
none	1
	16

no. of companies
 answering this question: 10

Note: some questionnaires contained more
 than one response to this question

SURVEY QUESTION #17b.
Documentation Sufficient?

Response	# of responses
Yes	3
adequate	1
effective	1
a little insufficient	1
no, paper doc. (vs. online)	
is almost impossible	1
	7

no. of companies
 answering this question: 7

SURVEY QUESTION #18a.
Expert System Maintained:

Response	# of responses
Totally In-house	9
Completely Contracted Out	1
Partly In-house/Partly Contracted Out	1
	11
Percentage of In-house Maintenance Reported for Partial Response:	
10	1
	1

no. of companies
answering this question: 11

SURVEY QUESTION #19a.
Central Organization to Oversee Expert Systems?

Response	# of responses
Yes	3
No	8
	11

no. of companies answering
this question: 11

SURVEY QUESTION #19b.
Expert Systems Central Organization Planned?

Response	# of responses
Yes	2
No	5
	7

no. of companies answering
this question: 7

SURVEY QUESTION #20.
Centralized Control Lessons Learned.

Response	# of responses
user control planned	1
AI group had limited success in selling itself in some areas	1
duplication of training muddies employment waters	1
central maintenance preserves the integrity of the system	1
believe both approaches to be beneficial	1
user needs are paramount	1
issue not addressed at this time	1
economies of scale lost with decentralized control	1
duplication of training effort with decentralized control	1
	9

no. of companies
answering this question: 8

SURVEY QUESTION #21.
Software Deficiency Reporting Procedures.

Response	# of responses
request forms to EDP	1
verbal	1
reported by user, field management, industrial engineers	1
reported by users to responsible person at corporate HQ	1
user to design engineer	1
users internally	1
no formal system planned	1
not done except locally	1
—	8

no. of companies
 answering this question: 8

SURVEY QUESTION #22.
How Often Is Expert System Modified?

Response	# of responses
constantly	2
continually updated	1
whenever necessary	1
whenever deficiency reported	2
as required	1
during policy changes	1
occasionally, to meet expanded functions	1
	9

no. of companies
answering this question: 9

SURVEY QUESTION #23.
Who Modifies Expert System & Lessons Learned?

Response	# of responses
in-house knowledge engineers	8
in-house domain experts	1
contractor	1
mix - 10% in-house, 90% contracted	1
	11

Lessons Learned:

domain experts are never satisfied
design and simpler coding equate to 100% in-house
maintenance
can be accomodated more readily than conventional system
keep an audit trail of updates and validate after
changes
you cannot just "add a rule" to add more knowledge

no. of companies participating: 11

Appendix C: Transcripts

The Dupont Visit

EM: All I can do is tell you how Dupont did things and that it's not a prescription for how to do things but an example of how a large, diversified manufacturing company set about a specific problem of AI, but more generally, we're running a management experiment on how you do technology transfer. How do you get it from an idea, whether it be from academia or where it is, and turn that idea into [garbled] production operation day-to-day. And in our case it's almost invariably a 24 hour a day operation, seven days a week. We could stack up papers this high [about three feet] on how not to do it. We've got lots of experience on how not to do it and take ten years longer than you need to....

Let me talk about why we have a correographed corporate effort and contrast AI work to other things, discretionary effort. The return we get on discretionary technical effort, in general, is about 10 to one. That is, we will make, if a technical person costs us \$50,000, we can expect, if they are doing improvement work, that they will create about half a million dollars worth of earnings every year they're doing it...but the annual return, we think, is about 10 to one. Our experience with AI systems has been about 15 to one.

EM: Our approach to it has been to build and deliver on existing hardware. We don't have any deliveries in LISP. We now have about sixty-odd routine commercial systems and another 500 that are in the wings.

MK: Let me stop you there for a minute. When you say you don't have any deliveries in LISP, what about any of the higher level shells like KEE or [garbled]. Everthing you're doing is on PCs?

EM: No, about half the systems we've written are on Vaxs, OK?

MK: What are you using? You said you're not using LISP; so are you using S.1 or something on the Vax?

EM: No, we're using a toolkit we wrote internally, because there wasn't anything written for the Vax, and it turns out that it's written in RS.1.

MK: Did you do any developments, like you started on Symbolics [garbled] and it didn't work or you just never moved in that direction because you didn't need to?

EM: I continue to have about 10 LISP machines in the company, all in research. Currently, I'm sponsoring a four departmental evaluation of KEE. We think if it has a role, it's in the area of conceptual design. If you look at expert systems, we think there are four: diagnosis and therapy; structured selection, how do I sell the guy the right thing from among 20 choices; the configuration problem, really a checklist, how do I get this together

right; and the fourth one, the planning and scheduling problem, which is really the structured selection problem with an infinite number of choices. The big problem is how do you narrow infinity down to a manageable number.

EM: What we've found is that ... if you're going to write one of these four general types of expert systems, and I haven't seen any others, ... the criteria you use is from the day I decide to write it to the day it's commercial, the day it's in routine use by the end user, as the time frame, what you'll find is that if you write it in one of these, let's call them Lotus type packages, because no programming is required, you can write about a rule an hour. If you write it in a language, and it doesn't really matter much which language it is, we've done some in Fortran, and we've done some in LISP, you'll write about a rule a day. There's a factor of eight in productivity of eight, maybe it's ten, but it's on an order of magnitude. You do one of two things, either you are more productive by a factor of ten or you increase your opportunity base by a factor of ten. We found that if you can put it in the hands of the end user, then you've opened up the mainstream of the business.

EB: Your article said the functional guys are doing this, not ADP guys...

EM: Right...

MX: When your functional guys do it, do you give that to

them as their full time job? Or are they still doing other things....

EM: They've still got their other job. It's just part time. In fact, what we find is the calendar time to do one is about four "X" [times] the effort time. People tend to work on it about a fourth of the time, or less. They just have their other job and the reason they'll work on it, and frequently they'll work on it at nights and weekends is because it has such value to them, because it's going to save them the three A.M. call to them.

The reason we think it's so important to force it on those guys is the "duty of care" issue, who's going to maintain it.

EB: How do you standardize it? ... We've found that a lot of the people have got their own set of information in their desk drawer ... not only do they have information in the desk drawer but they have written some unique software at that location that makes the operating system at that location different from another location, when in fact they are supposed to be the same. How do you prevent that from happening?

EM: First off, it's a matter of degree. If we're in the titanium dioxide business we don't care what the electronic imaging business is doing.... The reason people put their cheat sheet in the drawer is usually that the data is bad.... The real problem is that on a consultation, you gotta deliver the facts. Step one in just about everything

we do is: you have to gather up the data, and gather it up electronically....

In Dupont, our knowledge, particularly in the diagnosis and therapy world is so disperse that centralization and standardization is a mistake. You just say, give people aids to fix their job.

Now when you get to things like selection in a marketing environment or configuration, which now you're saying "I want my 50 salesmen to all use the same one." What we try to do there is tag an owner. In an organization, the owner has the "duty of care" of maintenance. That includes the whole ball of wax, that includes building it usually, maintaining it, but making sure that the data, the fact data, gets updated....so what we do on things that do need standardization is make somebody have, owns it. In fact, one of the things we do as a part of our go/no-go decision is, if in the first week we cannot identify an owner, in the sense that we put "owner", and he'll sign his name, we quit, we just don't work on it.

MK: How do you do that go/no-go decision? How do you identify applications to work on?

EM: It depends on the people to identify them.... The facts are that the people down on the floor that are getting beat up every day know what the problems are....

MK: Do you do any cost-benefit analysis at all, to decide whether or not you're actually going to let them spend their time letting them develop this thing?

EM: Well, very subjectively, and very quickly, back-of-the-envelope stuff. And the reasons are, if you're making 1500 percent return, it really doesn't matter if it's 1500 or a thousand, or two thousand, you don't want to spend too much time figuring out how much it is, you just want to do it. If your bailout position is a week, by the time you spend a week's analysis, you could have already decided whether or not you want to do it or not.

MK: How do you do that back of the envelope stuff? What are some of the factors....

EM: Some of the factors are: most people, when they have a problem, in our world, they know what the problem's worth. Or they can say, it's worth about this much. If they don't know, we usually let them go ahead, anyway. And the reasons are: if a problem is creating organizational stress, the mere fact that it creates that stress creates a lot of time spent solving the problem, that you just do it on faith.... If the organization views it as a problem, solving it will be economical. So, those that you really can't nail the quantification on, we don't worry much about, we go ahead and do them. What you find, after you do it, is you can look back, and say, hey, that made a lot of money.

EM: The one thing we won't do centrally, I will not build a system.... I believe that is a prescription for failure, because after I build it, and it doesn't matter what I build it on; I could build it on LISP, whatever's most productive,

and I get it built and I hand it to him and it's like buying a case of eggs. Who's going to take care of it? Who's going to be concerned that it doesn't work properly?

EM: Since we've gone to say, we're going to be in existing hardware, build on existing hardware, we really haven't had to fight for computer resources.

EM: Let me talk a little bit about the central activity we have here because basically, what we've said is that we're going to push everything out as well as we can....

Sometimes the owner is not necessarily the person who actually builds it. The owner is the person who actually has the "duty of care" of getting it built, of getting it maintained, and they may call upon a systems person to do that. One thing you've got to be sure of: if everybody owns it, nobody owns it, and it will die.

MK: Tell us about your AI network and how do those people communicate. Do you have any formal or informal organizational structure?

EM: My group is about a dozen folks and we run the helpline and we have knowledge; we've tried to put together a data base of applications....

EB: It's like a mailbox then, to avoid duplicatins?

EM: Right. The second thing we did, of course, we said 12 people across 30,000 was awful dilute, so we need some

organizational structure. There already existed an information exchange group...it meets every other month and now has about 400 electronic distribution people and the meeting usually draws about a hundred. Last October I formed a group of site coordinators who are really my second arm and that's my local representative, who doesn't report to me direct or indirect. He's simply a volunteer, site coordinator, who has the duty of care of making sure the people at his plant know what's going on at other places. He's kind of the helper, I give him a little more training, whatever he wants, so he is my extension at each little site. And by site, a marketing group might be a site. Some plants are big enough there tend to be two site coordinators.

MK: How many applications have you actually fielded today? How long have you had them fielded and the systems that your people are working, what kind of success/failure rate are you experiencing and why do you think the failures [inaudible]?

EM: Let me talk about the success/failure ratio. Ninety percent, or nine out of ten...

MK: So, 90 percent success rate?

EM: Ninety percent of the ones we start on we finish and are successful. Ten percent fail, and they fail because we can't find the owner.

MK: When you say you can't find the owner, I thought you

said your philosophy was the owner identified the problem?

EM: Sometimes the guy will identify, and the group, it kind of, in any work group, they kind of fiddle around. So sometimes we would come in and help to get some ideas together, we'd participate. If we participate, the point is, that the organization can't look someone in the eye and identify the owner.

EB: How do you know if the things are successful?

EM: We sample, and they could lie to us, but what we do is run a survey. We just got through running a survey...how many systems are you working on, how many were successful, how many failed....We survey every six months. I send out a little one page questionnaire, guaranteed five minutes, it's boxes, check the boxes, and you only have to write one word stuff, with an envelope to send it back to me.

EB: How much training are we talking about...?

EM: The training we give is two and a half days. In that, what we teach is, we assume they know nothing about AI, so there's an introduction to AI, elementary knowledge engineering, not PhD stuff, but enough to get your job done; and then we teach them the mechanics of three knowledge engineering tools.

MK: What tools do you use?

EM: The tools that we teach are toolkit internal, for the Vax, same thing as RS.1. We teach Insight II+, and we teach First Class....85 percent leave the class with a working prototype, usually tends to be about 20 rules.

MK: Your training is two and a half days. How do you select people to go to training? Do you pick anybody? Do they have to have any prior background in programming or?

EM: No, they just volunteer.

MK: Nothing, they don't even have to know how to write a text file?

EM: Most of them don't even know how to write a text file. What we've found is the best user is the one who uses E-mail and a spreadsheet, that level of capability. Anything above that is fine.

MK: You said you had 12 people that work as a central group. What are the kinds of things that they do?

EM: We train. They all do all of it. I rotate the helpline once a week, because you can do anything for a week and the helpline is the most miserable job in the world.... Everybody trains. We teach both here and remotely. We have a set of 10 compacts and rather than have somebody come to central class we go out to the plant and teach....

MK: How many people do you train at a time?

EM: Up to 20.

MK: With one instructor?

EM: Two instructors. It is about, more than 50 percent hands on; very little lecture...what we also do is joint prototyping....

MK: Once the applications get done, how do you make sure the verification and validation's done correctly? Or do you worry about that?

EM: Ah, the business boys do that.

MK: So, that's their job, to make sure it's tested properly

EM: I wouldn't know whether it

MK: and it meets policy and all that?

EM: And all that stuff. Because we don't have a prayer of doing it; everybody's got different policies, and how do you do validation anyway? You do case studies until it gets the wrong answer, then you put in the stuff to fix it. That ought to be done locally.

MK: Do you teach them how to do validation and testing in this two and a half days?

EM: The way we test; we tend to build the prototype and we believe you want to get the ultimate end user involved when it's dirty, as soon as possible.... The prototype turns into the real thing so the validation tends to be, you know, I like to say that expert systems are never done.

PB: You mentioned right at the beginning [garbled] return on investment [noise] how do you figure that?

EM: You only write two checks: to your employees, to your suppliers [garbled] The way you do it is the same way you do it in all other projects. You take the real cost to you and say, "this is the only one I got, how do I create the alternative cost sheet. Well, somebody does an engineering estimate of what it would have been had we not done this.

That's no different than anything else and I don't know any way other than that.

PB: How do you justify the costs, how do you identify the cost benefits?

EM: [inaudible] you can work a week and try to figure out whether it's any good or not [garbled] it's a riskless start [inaudible]

PB: Do you have any planned benchmark for increased productivity or cost reduction on these things?

EM: Globally, I want to make a 150 million dollar impact

PB: For each system?

EM: No, that's total. Now each system, a standard system, is going to make a 100,000 bucks and it's going to take a month to build. That's what we've seen. That's average, some are more, some are less. Some take a week and some take two months.

PB: Do each one of your end users go to this class?

EM: Just the owners, just the builders. To use them, it's just question and answer stuff.

MK: How long does it really take you to deliver on an application? You send somebody to school for two and a half days and you go out into the field for two days and help them prototype. How long before they're actually delivering something back to their management?

EM: ...They deliver in about four months. ...From the time

they get started till it's a commercial and routine use.

MK: When you say you put a system in the field in four months, what's the average size of that system?

EM: Our average size has turned out to be 80 rules.

MK: Everything rule based?

EM: Yes, essentially everything is rule based.

MK: What are the backgrounds of the people in your AI group?

EM: Marketing research with a masters in Comp. Sci. and a PhD in marketing research

Information systems, an MBA

Process control, I believe he is an electrical engineer

Chemical engineer

Chemist, 10 years experience

I don't have anybody that has any formal AI training.

MK: What's your most successful system?

EM: Probably the scheduling system we built. We did it after three different LISP people had failed.... It has economic impact across the company.... The LISP people who were going to write it in code said they could give us a prototype of one piece of the problem in nine months and I forget, 7 or 800,000 dollars. We tried it with Knowledge Craft and spent 100,000 dollars and stopped because it just wasn't going to work. We wrote it completely in a 1000

man-hours.

The UNISYS Visit

DB: I assume this is why you really came. What I am presenting is not the opinion of UNISYS corporation. It's the opinion of us and, in many cases, the opinion of me.... I'd like to go through the lessons I think we've learned, some important lessons I think we haven't learned....

Lesson learned number one is that I think expert systems are helpful and useful in some problems that involve logistics analysis. Things that it is particularly useful for are problems that feature rare but important exception cases...things we can't effectively solve by conventional techniques, and things that are significant cost drivers.

One thing that I'll throw out is that expert systems are not only a tool to build things, but also in effect, an analytical device. The process of building an expert system helps you analyze a problem

One of our most important findings is that expert systems promote a more uniform and objective analysis.... People have diverse backgrounds and we want them all doing the same things. It helps to establish a more objective evaluation procedure.... Ensuring the continuity of knowledge is if people quit you have some grasp of what they were doing....

Another very important one from our perspective is an expert system that's really delivered as part of a more conventional system. By that I mean the communications, the data base...all that is not AI, it's just software engineering.... You should be focussing not on I want to build an expert system to solve this problem but I want to solve this problem. So, it's a systems engineering function, not a knowledge engineering function. That's a very important distinction. A systems engineer should also be a knowledge engineer but he should first be a systems engineer. When he looks at the problem he should first partition the problem and use the appropriate software tools to take care of the various parts of that problem....

Expert systems are effective extensions to existing data base systems.... Building and maintaining data bases is very expensive.... Expert systems are particularly effective if the problem involves large amounts of complex data which a human expert would find difficult to manage.... Expert systems are useful when only a small portion of a large volume of data is necessary.... Allowing the user to accumulate, view and manipulate the data in a workstation environment is a very useful thing to do from a productivity viewpoint....

Lesson learned number five: expert system development requires good front-end analysis, just like a conventional system. By that I mean you focus on the problem and take a

systems engineering approach to the problem.... Build a prototype during concept definition.... Maintain a systems perspective; don't just concentrate on the rule base. Sometimes expert system strategies are applied to problems which are simply ill-defined.

EB: How many rules do you have in this system?

DB: It changes daily, currently there's about 50. I expect to see a 100 to 150.

MK: What is the average size of the rule?

DB: Probably four to five premises and probably three or four actions.

DB: The key to the issue [control] is you have to maintain control over the contents of that rule base. How you choose to do that is entirely up to you but you can't allow two users of different cycles to just go off and start changing things, because if the content of the system starts to diverge, you've lost a lot of the benefit of the system, which is a consistent approach, and God knows, I hope you'll want to verify those changes in some fashion. It's possible to put in rules that directly contradict each other and most shells won't tell you that you did that....

WR: There are users out there that ... if there is a problem or an enhancement they want to make, we'll pop a copy of the official controlled version of the program off

to them, let them make the changes, we'll review it, we'll make sure it's tested right, and when it works, it's in the system in the production mode.... The trick is that the folks have to remember that there are configuration control issues and you can't let that stuff slide....

DB: I tend to feel that most expert systems in logistics will be an expert assistant.... Logistics matters are money matters and decisions that you make affect large flows of money all over the place, therefore, it's highly unlikely that any inventory manager would allow a system to initiate a procurement without signature, without any intervention. If that's true it requires additional emphasis on the interface....

Clearly distinguish between research, feasibility studies, and production development. They're all different, they all have a different intent, and they all require, in my opinion, a different approach....

Reasons why I don't feel that if a conventional system will do it you should build an expert system.... The quality assurance issues: how do you know when you're done? How do you validate and verify? ... The maintenance and support issues, the hardware and software you put it on...

Bibliography

1. "AI Conversations, Interview with Harvey Newquist III," PC AI: 50-54 (Spring 87).
2. Air Force Logistics Command. "AFLC Artificial Intelligence Objectives," September 1986.
3. Air Force Logistics Command. "IMA Implementation Plan". AFLC Artificial Intelligence Management Integration Office, September 1986.
4. Alexander, Tom. "Teaching Computers the Art of Reason," Fortune: 82-92 (17 May 82).
5. Allen, Major Mary Kay. Development of an Expert System. Ph.D. dissertation. Ohio State University, Columbus OH, August 1986.
6. Allen, Mary Kay and others. "Silicon Management: Applying Expert Systems Technology In AFLC," unpublished work, September 1986.
7. Bachant, Judith and John Mcdermott. "R1 Revisited: Four Years in the Trenches," The AI Magazine: 21-32 (Fall 1984).
8. Cook, James H. Artificial Intelligence Applications in Manufacturing. MITIAC Report No. IA-85-01, October 1985.
9. Cotellessa, Robert F. Identifying Research Areas in the Computer Industry to 1995. Park Ridge NJ: Noyes Publications, 1984.
10. Davis, Dwight B. "Artificial Intelligence Goes to Work," High Technology: 16-27 (April 87).
11. Dawson, Paul. AFLC Artificial Intelligence (AI) Program Detailed Instruction. AFLC working paper. 29 July 1987.
12. "Digital and AI". Perspective. Digital Equipment Corporation, Stow MA, June 1986.
13. Fersko-Weiss, Henry. "AI Update 1986 Hits and Misses," Personal Computing: 99-107 (December 1986).
14. "Financial Planner With Nerves of Silicon," Business Week: 108 (October 7, 1985).

15. Firdman, Henry Eric. "Second Thoughts about AAAI'86," AI Expert: 65-69 (December 1986).
16. Gold, Jordan. "Do-It-Yourself Expert Systems," Computer Decisions: 76-81 (January 14, 1986).
17. Guterl, Fred V. "Computers Think For Business," Dun's Business Monthly: 30-37 (October 1986).
18. Harmon, Paul and David King, Expert Systems Artificial Intelligence in Business. New York: John Wiley & Sons, 1985.
19. Hengl, Terry. "Publisher's Note," PC AI: 8 (Spring 1987).
20. Hinden, Harvey J. "General-purpose Machines and PCs Emerge as AI Platforms," Computer Design: 43-48 (November 1, 1986).
21. Juris, Robbin. "Look for Smart Systems," Computer Decisions: 36-39 (January 12, 1987).
22. Kinnucan, Paul. "Computers that Think Like Experts," High Technology: 30-42 (January 1984).
23. Kneale, Dennis. "How Coopers & Lybrand put Expertise Into Its Computers," The Wall Street Journal: 29 (November 14 1986).
24. Leibholtz, Stephen W. "Expert Systems; Performance, Potentials, Promises, Problems," Government Executive: 35-38 (January 1987).
25. McDermott, John. "R1, the Formative Years," AI Magazine: 21-29 (Summer 1981).
26. Marcot, Bruce. "Testing Your Knowledge Base," AI Expert: 42-47 (July 1987).
27. Myers, Ware. "Introduction to Expert Systems," IEEE Expert: 100-109 (Spring 1986).
28. Newquist, Harvey P. III. "American Express and AI: Don't Leave Home Without Them," AI Expert: 63-65 (April 87).
29. "Peter Hart Talks about Expert Systems," IEEE Expert: 96-99 (Spring 1986).

30. Raunch-Hindin, Wendy. "Software Integrates AI, Standard Systems," Mini-Micro Systems: 69-83 (October 1986).
31. Seaman, John. "AI Moves Into the Mainstream," Computer Decisions: 100-104 (December 3, 1984).
32. Turban, Efraim and Theodore J. Mock. "Expert Systems, What They Mean to the Executive," New Management: 45-51 (Summer 1985).
33. Tyson, Reuben and Richard Herrod. "Capturing 32 Years on Computer," Food Engineering: 69-71 (December 1985).
34. Van Horn, Mike. Understanding Expert Systems. New York: Bantam Books, 1986.
35. Wang, Eugene. "Artificial Intelligence for Competitive Advantage," AI Expert: 5-6 (January 1987).
36. Waterman, Donald A. A Guide to Expert Systems. Reading MA: Addison-Wesley, 1986.
37. Williamson, Mickey. "At Du Pont, Expert Systems are Key to AI Implementation," PC Week: 35-36 (January 13, 1987).
38. Wood, Lamont. "Out of the Ivory Tower". Computer Decisions: 47-48 (Jan 26, 1987).

UITA

Major Paul R. Boggs was born on 5 January 1951 in Tyler, Texas. He attended Texas A & M University, receiving a Bachelor of Arts degree in Political Science in May 1973. A member of the Corps of Cadets, he received his commission upon graduation. Reporting to navigator flight training in September 1973, he received his wings in June 1974. After completing Combat Crew Training School at Castle AFB, California, he was assigned to the 7th Air Refueling Squadron at Carswell AFB, Texas as a squadron instructor navigator and flight examiner. In November 1979, he was assigned to the 909th Air Refueling Squadron, Kadena AB, Okinawa, Japan as an instructor navigator, flight examiner, flight scheduler, plans officer, and chief of the Emergency War Order Operations Branch, Operations and Plans Division, 376th Strategic Wing. He received a Masters of Science degree in Systems Management from the University of Southern California in 1982. In May, 1983, he was assigned to the 11th Strategic Group, RAF Fairford, England as an Emergency Actions Officer and Group Executive Officer. Major Boggs was assigned to the School of Systems and Logistics, Air Force Institute of Technology, in June, 1986.

Permanent Address: 1009 W. 10th St.

Tyler, Texas 75701

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED			1b. RESTRICTIVE MARKINGS		
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; distribution unlimited.		
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE					
4. PERFORMING ORGANIZATION REPORT NUMBER(S) AFIT/GLM/LSY/87S-6			5. MONITORING ORGANIZATION REPORT NUMBER(S)		
6a. NAME OF PERFORMING ORGANIZATION School of Systems and Logistics		6b. OFFICE SYMBOL (If applicable) AFIT/LSM		7a. NAME OF MONITORING ORGANIZATION	
6c. ADDRESS (City, State, and ZIP Code) Air Force Institute of Technology Wright Patterson AFB, OH 45433-6583			7b. ADDRESS (City, State, and ZIP Code)		
8a. NAME OF FUNDING / SPONSORING ORGANIZATION HQ AFLC		8b. OFFICE SYMBOL (If applicable) MMT		9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER	
8c. ADDRESS (City, State, and ZIP Code)			10. SOURCE OF FUNDING NUMBERS		
			PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.
11. TITLE (Include Security Classification) A STRATEGIC PLAN FOR THE SUPPORT OF EXPERT SYSTEMS IN ORGANIZATIONS					
12. PERSONAL AUTHOR(S) Paul R. Boggs, B.A., M.S., Major, USAF					
13a. TYPE OF REPORT MS Thesis		13b. TIME COVERED FROM _____ TO _____		14. DATE OF REPORT (Year, Month, Day) 1987 September	
15. PAGE COUNT 157					
16. SUPPLEMENTARY NOTATION					
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number) Expert Systems, Artificial Intelligence, Management, Computer Systems, Support		
FIELD	GROUP	SUB-GROUP			
12	09				
19. ABSTRACT (Continue on reverse if necessary and identify by block number) See reverse. Thesis Advisor: Mary Kay Allen, Major, USAF Assistant Professor of Systems Management					
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT ☐ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT ☐ DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED		
22a. NAME OF RESPONSIBLE INDIVIDUAL Mary K. Allen, Major, USAF			22b. TELEPHONE (Include Area Code) (513) 255-2229		22c. OFFICE SYMBOL AFLC/MMT

UNCLASSIFIED

Block 19

ABSTRACT

The purpose of this study was to elicit lessons learned in business and in industry in the fielding of expert systems. The study looked at 12 research questions dealing with the acquisition, maintenance, documentation, organizational, performance, personnel, and software modification issues associated with the support of expert systems. The study surveyed the Fortune 500 corporations via questionnaire and telephonic interviews.

The study found that the Fortune 500 corporations viewed expert systems as an instrument of strategic advantage and were very protective of what is considered proprietary information. The study found the corporations to be particularly secretive about costs of acquiring and maintaining expert systems.

Analysis of the questionnaire and interview data found that expert systems in the Fortune 500 are predominately delivered on personal computer hardware and software systems, and primarily developed in-house. Support issues such as the software modification, performance, and documentation issues are still being addressed.

UNCLASSIFIED

END
FILMED
FEB. 1988
DTIC