

**DESIGNING EXPERT SYSTEMS IN A BUSINESS ENVIRONMENT**

**James Clifford, Matthias Jarke  
and Henry C. Lucas, Jr.**

May 1985

Center for Research on Information Systems  
Computer Applications and Information Systems Area  
Graduate School of Business Administration  
New York University

**Working Paper Series**

CRIS #99

GBA #85-62

To appear in : L.F. Pau (ed.): **Artificial Intelligence  
in Economics and Management**, North-Holland 1985.

## Table of Contents

|  |    |
|--|----|
| 1. INTRODUCTION  | 2  |
| 2. EXPERT SYSTEMS IN THE BUSINESS ENVIRONMENT              | 3  |
| 3. THE UNDERWRITING EXPERT SYSTEM                          | 7  |
| 3.1. <u>The Information-Gathering Decision</u>             | 7  |
| 3.2. <u>The Underwriting Decision</u>                      | 9  |
| 4. KNOWLEDGE ACQUISITION                                   | 10 |
| 4.1. <u>General Issues in a Business Environment</u>       | 10 |
| 4.2. <u>Underwriting System in Particular</u>              | 10 |
| 5. UNDERWRITING IN CONTEXT                                 | 11 |
| 6. RELATIONSHIP TO TRADITIONAL SYSTEMS ANALYSIS AND DESIGN | 13 |
| 6.1. <u>Traditional Systems Analysis and Design</u>        | 14 |
| 6.2. <u>Prototyping</u>                                    | 15 |
| 6.3. <u>Expert Systems Design</u>                          | 15 |
| 7. SUMMARY AND CONCLUSIONS                                 | 17 |

## Abstract

The integration of an ES into a business environment presents a different set of problems to the designer. First, it can be difficult to isolate and draw boundaries around the domain of the business problem. Second, there is often a need to generate a large number of expert decisions in a short period of time. That is, there can be many transactions requiring some expertise to process, such as applications for life insurance. Finally, there may exist a number of computerized transactions processing systems which interact with very large databases and there may be a need to integrate the ES with these existing systems. This paper discusses these general issues involved in developing expert systems for business applications, with particular examples drawn from the domain of insurance underwriting.

## 1. INTRODUCTION

An expert system (ES) is a computer application that provides decision support similar to that of a human expert in solving a problem. The systems analyst who builds an expert system is called a knowledge engineer; his or her role is to learn the expert's approach to problem solving. Usually the knowledge engineer can formulate the expert's techniques as a series of rules. Then, an Artificial Intelligence language like Prolog can be used to design a computer system that exhibits some of the decision making characteristics of the expert.

Most expert systems to date incorporate knowledge that is primarily scientifically oriented; for example, applications exist in medicine, chemistry and biology [CJV 83]. These systems can be characterized as having a clear task and expertise that is relatively easy to locate. Generally the problems to which these applications are applied require substantial time, such as the diagnosis of a patient's diseases. Frequently, existing ES's are used in environments where there is no current computer support.

The conditions imposed by a business environment differ from those described above. First there appears to be extensive interaction among different tasks; drawing the boundaries for what the expert does is not easy. Second, there may be substantial expertise required to perform a given task, but the expert is often functioning in a high volume, high pressure environment. He or she often has to generate a large number of expert decisions in a relatively short period of time. Finally, successful business utilization of expert systems technology will require that these systems interact heavily with existing business subsystems, such as transactions processing systems and database management systems.

This paper reports on the authors' experiences and reflections on the design of an expert system for life insurance underwriting. The particular application is underwriting applications for life insurance for a large insurance company. The problem, as in many ES, is essentially one of classification: determining into which insurance category an applicant falls for insurance purposes. The problem is an important one, and is a routine part of the business. Transactions volumes are high and the underwriter makes use of a large volume of information in coming to a decision. In addition, the underwriter must interpret facts and frequently apply judgment based more on experience or rules of thumb than on clear-cut policies. While insurance firms are some of the largest users of computer-based processing systems, the fact that there are very few systems (for an exception, see the CAUSE underwriting DSS [A1 80]) to aid in the underwriting decision provides evidence that significant human judgment is required for this task.

It should be noted that this research is part of a larger NYU Expert Systems Project which is also concerned with developing a general interface mechanism between expert systems and database management systems [JV 84].

## 2. EXPERT SYSTEMS IN THE BUSINESS ENVIRONMENT

One generally accepted macro view of the managerial control processes in an organization defines the structure as consisting of a hierarchy of three distinct but cooperating levels of activity and decision-making [An 65, Si 60, GS 71]. Operational control, the lowest level in the hierarchy, is concerned with all of the details of the individual transactions which constitute the activity of the business. One level higher is the managerial control level, where more summary information about the business is needed in order to support medium-level decisions affecting the management of the business functions. Finally, the top level consists of the strategic planning and policy making function which sets the goals and defines the overall nature of the business organization.

Traditional data processing systems have primarily concerned themselves with the lower two levels of this hierarchy. When one thinks of business processing systems one thinks of such transaction processing systems as order-entry or inventory management, and such managerial reporting systems as accounts receivable and general ledger. More recently, a class of somewhat less-defined systems known as Decision Support Systems have begun to be discussed and researched [SC 82]. These systems are primarily aimed at supporting the kinds of decisions that are required at the higher levels of this hierarchy.

We believe that expert systems and expert systems technology have a role to play at all levels of this organizational hierarchy. The reason for this is that, orthogonal to this hierarchic view of the nature of the business activity is a view of the complexity of the activity being carried on. Accepting this classification of business activities along a complexity scale ranging from "structured" to "semi-structured" to "unstructured", these two parameters, level of activity and complexity of activity, produce a matrix of business activity types. We might view the situation therefore as in Figure 1, and ask ourselves (a) where in this matrix does a particular activity lie, but, more importantly for this discussion, (b) where in this matrix are traditional DP techniques and the newer ES techniques most suited?

|             | STRUCTURED | SEMI-STRUCTURED | UNSTRUCTURED |
|-------------|------------|-----------------|--------------|
| TRANSACTION | DP         | ES              | ES*          |
| MIDDLE-LVL  | DP/DSS     | ES/DSS          | ES*          |
| STRATEGIC   | ES/DSS     | ES*             | ?            |

DP = traditional Data Processing System  
 DSS = traditional Decision Support System  
 ES = Expert System, Current Technology  
 ES\* = Expert System, Future Technology  
 (Integrated with DSS Concepts)

Figure 1: Taxonomy of Computer Applications

We believe that traditional DP techniques have been and will continue to be applied successfully to both transaction and middle-level information systems, but primarily at problems of a routine level of difficulty. What is beginning to happen, however, is that the techniques of ES and DSS technology are being applied at some routine problems at the strategic level, and also at complex problems at all three levels. For the really difficult problems at all three levels, it seems that the jury is still out; we will have to wait both for some more powerful ES/DSS techniques, languages and tools, and also for an evaluation of the success/failure of these techniques at the complex level.

What this analysis, if correct, means, is that business have barely begun to automate or semi-automate processes, decision-making, and planning in their organizations -- they have simply gone through the first stages of a computer revolution, which have helped to solve some of their easier problems. Now that techniques of AI have been developed that enable the automation of much more intelligent processes, businesses will begin to re-examine the whole range of their activities, from transaction processing through strategic planning, searching for new candidates for automation, from among those problems previously thought too difficult.

The underwriting expert system project can be viewed as an attempt to apply ES technology in a complex transaction processing system. Before discussing some of the particular characteristics of this problem, some general remarks about transaction processing and ES technology are in order.

The vast majority of the ES's that have been built have been highly interactive, and intended for an environment in which there was a long interaction with a human and the system in an effort to come to a reasonable solution or decision for a single problem. Whether this problem was the diagnosis of a patient's illness (MYCIN [Sh 76], INTERNIST [P 84], PUFF [Kea 78], etc.), the synthesis of a promising new chemical compound (SYNCHEM [Gea 77]), the likelihood of a mineral deposit in a particular geographical area (PROSPECTOR [DR 84]), or whatever, the striking characteristic is one of prolonged interactive problem-solving.

R1 [M 82] is probably the first ES reported in the literature which moved closer to the transaction processing realm. It also appears to be one of the few expert systems in actual routine use. This system advises systems engineers at Digital Equipment in configuring VAX computers. Although advertised as "the" major success in expert systems to date, it should be noted that its development costs have reportedly far exceeded the current or expected benefits of the system. We can conclude that (a) ways must be found to reduce the costs of building expert systems, and (b) one has to be very careful in selecting an application domain with sufficient potential payback. As noted above, the latter can be achieved by supporting either important single decisions or mass-processing of transaction-like, medium-complexity decisions.

What are some of the characteristics of transaction processing systems? The first is that they must be able to handle a high volume of transactions per unit of time, i.e., they should not spend a long time agonizing over the decision for any particular case. A second is that they typically have a high level of interaction with other systems in the business, e.g. a DBMS, other transaction systems and systems at the middle and strategic levels of the organization. Finally, a third characteristic is that they must be able to track transactions through various stages of processing, for example from initial processing through an "awaiting action" state through a "final disposition" state, since in many applications not all transactions can be completely processed in one pass. Each of these factors should influence the shape of business ES's that might be built for complex problems at the transaction level.

The first issue, that of high volume, dictates that the ES needs to be organized to be extremely fast, and probably would best operate in a batch rather than an interactive mode. However, many people feel that ES's will not be trusted to make decisions unilaterally in complex tasks, but rather will only be effective when built to operate in conjunction with human experts in the domain. The inherent tension between these two issues involved in complex transaction level systems, namely high volume with its need for speed, and decision importance with its need for confidence, will require creative systems design if ES's are to succeed in this business area. As we shall see, the underwriting system addressed this issue by partitioning the transactions into disjoint classes of transactions with different characteristics.

The second issue is a crucial one in determining the scope of the ES and its relationship to its environment. Most ES's to date have been built by systems designers who were free to construct their knowledge base in isolation from any existing organizational databases, policies, or constraints. They have typically chosen a knowledge representation scheme appropriate to the problem to be solved and to the language in which the system was being built. By contrast, a business ES will undoubtedly have to rely upon an existing DBMS for the major part of its factual knowledge about the enterprise in question, and probably upon existing business subsystems -- transaction systems, mathematical modelling systems, decision support systems, etc. -- for some of the data processing that will be relevant to its proper functioning [JV 84]. No business could tolerate the indiscriminate replication of data, of modelling algorithms or processing elements such as summary reports and statistical data on its operations. Not only would such replication be costly, but it would open the door to data and processing inconsistency and lack of security -- all of the things that modern DP techniques and policies, and DBMS's have worked hard to prevent. These components of the business -- data and processing information -- have attained the status of an important corporate resource, and ES technology, if it is to succeed in this environment, will have to accommodate itself to this view and the policies and structures that it has generated.

In practice this means that the ES will have to rely heavily -- as more and more all business systems do -- upon the DBMS as the source of most of its data. One of the major thrusts of the research that involved us in the insurance underwriting ES was the exploration of a general mechanism for just such ES-DBMS interaction, in particular in the context of a Prolog-based ES and a relational DBMS [JCV 84, VCJ 85]. However it is clear that the DBMS is not the only existing system with which the ES will have to contend. For instance, the underwriting function in insurance is but one step in the process of processing an insurance application. This process begins with the salesperson-client interaction, which generates an application. Some of the information on the application is then coded into an order-entry subsystem. This generates a file of pending applications, which together with the entire application is sent to the underwriters. After the underwriting is completed -- which may involve a pending file awaiting the arrival of additional information, as we shall see -- the application is then either sent on to policy issue, where the actuarial subsystem computers the premium and creates an entry in the corporate database, or it is sent to the reject subsystem which must generate a file and a letter explaining the reason for the rejection. If an ES is to be built to help perform some or all of the tasks of the underwriting function, clearly it must fit into an elaborate, pre-existing structure of systems, subsystems, policies, procedures, and legal strictures.

Before discussing some of the more general issues involved in business ES deployment, the next section provides a context by discussing the particular application of insurance underwriting.



### 3. THE UNDERWRITING EXPERT SYSTEM

An insurance policy transaction begins when the sales representative and the customer complete an application form and forward it to a regional office for underwriting. Here the application is assigned to an underwriter based on the size of the requested insurance policy and the type of medical information provided on the form. At this point the underwriting process begins.

An insurance underwriter is basically charged with determining whether or not the company should accept a customer. This overall task involves examining the potential customer's application and all other accompanying material to determine either (a) accept the customer in such and such a rating category, or (b) reject the customer for such and such reasons, or (c) withhold a final decision pending the acquisition of additional supporting information. In the company cooperating in this research, the type of product involved is life insurance of any kind, term or whole life.

Underwriting is not a single decision at one point in time. Rather, there are two distinct decision styles. The underwriter first makes a decision on what information to request, for example, a medical examination, report from the applicant's personal physician or several other pieces of information. The application is placed in a file pending the arrival of the information needed for underwriting. When all data are complete, the underwriter examines the information and makes a decision that classifies the applicant into one of a number of categories, each having a different price for a given level of insurance coverage.

#### 3.1. The Information-Gathering Decision

Figure 2 shows the different sources of information available to the underwriter. The applicant signs a release giving consent for all of this information to be made available to the company, and FOR the results of the underwriting process to be reported to a computerized database. Much of this information costs the insurance company money to obtain; in addition, the more information requested, the longer the underwriting process is likely to require, resulting in the potential loss of the customer to the competition.

|   |
|---|
| <p>A check of computerized databases MIB/AIF of applicants for insurance for most large carriers in the U.S.</p> <p>A medical examination by varying levels of medical personnel from a paramedic to a specialist.</p> <p>A report from the applicant's attending physician.</p> <p>A mercantile report or credit and community investigation.</p> <p>A driving report from the Department of Motor Vehicles.</p> <p>Tax returns and/or a financial statement from the applicant.</p> |
|---|

Figure 2: Sources of Underwriting Information

One of the most important sources of information for the underwriting decision is the applicant's previous insurance history. This information can come from one of two sources. There is an industry-wide database maintained by the major U.S. insurance companies. Each time a company rejects an applicant for insurance, it reports that fact, and the reasons for the rejection, to this shared database.. When an individual applies for insurance at the company in this study, there is an automatic search request issued to this database in certain instances, for example is the amount of coverage applied for is over a certain threshold, or the applicant is over a certain age. A second source of similar, though more detailed, information is the company's own internal database. Again, certain policy characteristics trigger a query to this database.

The next information that might be requested is a medical examination. Such factors as the age of the applicant and the amount of insurance requested will determine the scope of a medical examination. In addition, there are questions on the application about medical history. If the applicant reports one or more of a number of conditions, the underwriter will generally ask for a medical examination. This decision is complicated by the fact that there are several types of medical exams; the applicant may need a paramedical exam, a doctor, or a company-appointed specialist.

Certain answers on the application cause the underwriter to ask for a report from the applicant's attending physician. Usually the doctor photocopies the most recent part of the patient's record and sends it to the company. Naturally the complexity of reading a handwritten report, difficult enough for the human "expert," were sufficient to place the examination of these reports outside the scope of the proposed system.

A mercantile report is like a credit investigation. However, it focuses on the behavior of the applicant and may be requested for a number of reasons.

Again, a major reason for authorizing this type of report is the face amount of the requested policy. Other triggers include such factors as the occupation of the applicant, a history of motor vehicle problems, or perhaps a medical disability.

A mercantile report always includes a query to the department of motor vehicles to determine the applicant's driving record. Of particular interest are any arrests for driving while intoxicated. If the applicant does not require a mercantile, but is considered a high-risk driver for some other reason, especially age, and if additionally the policy amount warrants it, the company will query the department of motor vehicles.

Finally, for certain very large policies, the applicant will be requested to submit a financial statement, usually in the form of a copy of a tax return.

### 3.2. The Underwriting Decision

Given the requested information, the underwriter can evaluate the application. The underwriter assigns scores for various conditions and diseases; the more points assigned to an application, the higher the cost of the insurance. There are a number of possible categories into the which the application will be placed, each one representing a different premium schedule.

There can be a substantial dollar difference in premiums between the various categories, so the underwriting decision is an important one for the applicant. In addition to the classification, certain medical conditions will warrant a surcharge of so many dollars for every thousand dollars of insurance for a period of years.

What is the source of the underwriters' scoring system? The medical and actuarial departments furnish underwriters with a medical guide. The guide's point totals assigned for different conditions reflect the results of actuarial studies and represent the risk the applicant poses to the insurance company.

However, not all risks are additive, as assumed in the scoring system. In such cases, the underwriter applies judgmental rules. For example, an application is rejected if the applicant has a history of mental problems and is unemployed, even though none of these properties have a high score by themselves. Because the agent in the field has done some screening and since the company must insure customers to remain in business, there are relatively few absolute rejects, on the order of 5% of applications.

## 4. KNOWLEDGE ACQUISITION

### 4.1. General Issues in a Business Environment

Business expertise is seldom localizable. In most businesses it is distributed among various people at different levels, and among various forms of documentation and processing elements. Most expert systems have been built in a different environment, where the expertise of one person was evoked during knowledge acquisition, systematized, and incorporated into the knowledge base.

In a business system, the knowledge acquisition process is complicated by the distributed nature of the expertise. Among the possible sources for relevant expertise are existing applications programs and other computerized systems; manuals, rule books, policy statements or other written materials representing compiled expertise of the organization; human experts from various departments or functional areas. Among the problems that this diversity of knowledge sources poses are the following: inconsistency across the sources; incompatibility of the view of the problem space; potential compromising of a human expert who may disagree with a stated policy or rule. The issue of resolving conflicts or disparities among multiple experts is one that will increasingly be an issue in the development of expert systems [H 85].

The security of the organization's rules is an added issue in business ES's. The power and flexibility of the knowledge representation tools used to build ES's, i.e., the ability both to represent rules and policies declaratively and to access and modify these rules during program execution comes with a price. Just as data in a DBMS is seen as a corporate resource which must be protected from mis-use, so too (perhaps even more so?) business policies, strategies and rules encoded into the knowledge base of the expert system must be protected. One offshoot of this necessity to preserve the confidentiality of business rules is the difficulty of publishing results in the area of business expert systems; most of the real progress involves the development of prototype systems in confidential application areas.

### 4.2. Underwriting System in Particular

For the underwriting expert, the design team interviewed various underwriters. At first there was interest in an expert for the high value policies, those with face amounts in the millions of dollars. However, it turns out that these policies are underwritten by a committee rather than a single individual. Construction of an ES for the committee decision did not look promising because of the number of individuals involved and the complexity of the decision process.

Below the level of policies requiring a decision by committee, there is a large volume of underwriting at the district offices. The human expert used in the knowledge acquisition and prototyping for this system is a senior lay medical underwriter. He can underwrite policies up to about a quarter of a million

dollars, and other underwriters send to him any applications that need some medical judgement.

The design team met with the expert for a period of six months prior to writing any code. The meetings began with the expert describing his job and decisions. Soon, the designers asked the expert to bring various applications and talk through his decisions on each one. A protocol analysis of these applications reviews provided enough understanding of the process to involve the expert in design. The designers and human expert worked together to complete Figure 2 and to describe the rules that apply for requesting each type of information.

The system was then developed piece by piece, first designing the data structures needed to represent all of the information relevant to the decision, and then building the rules that determine when a request for individual pieces of additional information are necessary. Moreover, there were rules for developing a composite "score" of the application, which could potentially contribute both to a request for additional information and ultimately to the insurance classification of the application. The development proceeded in the usual ES development mode, namely as a sequence of successive refinements to the quality of the system's decision making. Each meeting we would review the performance of the system on an application with our human expert, and discover where and why it had errors or been incomplete. In the next meeting we would have developed new rules to handle the previously discovered problems, and then proceed to discover new ones.

At this point, an expert system is working for the first underwriting decision, on what information to acquire. In a limited "test" of the system against eight applications chosen by the underwriter to reflect a broad spectrum of application types, the system's performance was comparable to our advisor. In a production environment the system would need substantially more thorough testing, in a carefully planned and controlled test against many more cases and also involving other underwriters. Current plans call for developing a full expert system to underwrite either a class of applicants or applicants with a particular disease to learn how an expert can assist in the actual classification of an application.

##### 5. UNDERWRITING IN CONTEXT

The insurance firm is organized into regions and there is a steady flow of applications. However, business does tend to be seasonal with a definite peak in December as agents try to achieve their sales goals for the year. In contrast to systems that might advise on unsticking an oil well drill bit or on where a valuable body of one might exist, there are literally hundreds of underwriting decisions made each day.

In contrast with other ES, this system involves a number of relatively low value decisions, the cumulative effect of which is of high importance to the firm. It should also be noted that there is a lack of a fine grained performance measure for the underwriter. These individuals are not given feedback on how the cases they underwrite perform, though such information is available in the company.

An ES must be designed to work with the information provided on the insurance application and from the sources described above. There can be no lengthy discussion with the underwriter or it will take too long to process transactions. In fact, one of the attractions of an expert system for underwriting is the potential to move part of the decision into the field. An underwriting expert accessible by the sales agent could speed response to a customer. With government deregulation, insurance companies are beginning to offer more financial services. Customers, however, are used to relatively fast response from current service firms. An ES that provided a fast underwriting decision for a large percentage of applicants would be a competitive advantage for an insurance company. Exactly how such a system would be implemented, and its potential impact upon employment, are questions that need further examination.

There is an existing system at the company which accepts limited input from the application and requests reports from both the internal and industry-wide databases when the applicant falls under the rules requiring such information. Further, some of the firm's existing systems are used by the actuarial and medical staff to prepare the medical guide, a manual which covers all of the underwriting policies regarding the medical condition of the applicant. One could eventually envision a series of expert systems, each making extensive use of part of the firm's database. An actuarial expert (cf. [SJ 85]) could be queried by the underwriting expert and could produce a customized risk analysis for each applicant. An overall picture of some of the system interconnections desirable in the insurance company environment is shown in Figure 3.

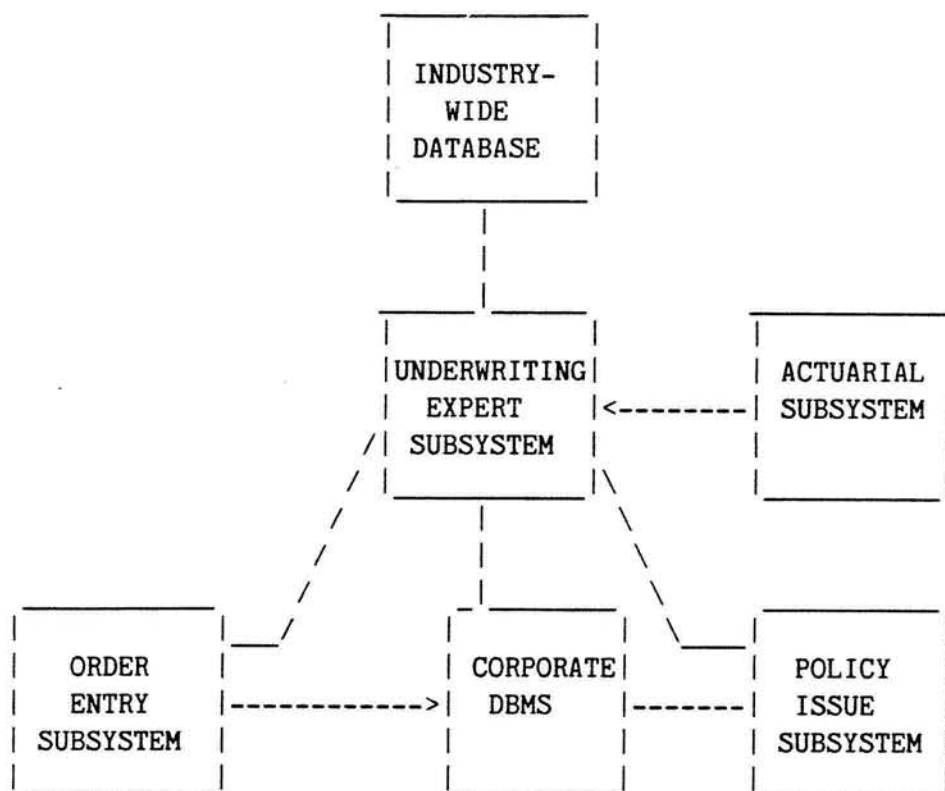


Figure 3: Underwriting System in Context

More immediately, it is not clear yet what kind of interaction the underwriting expert will require with databases. The information gathering part of the expert system accesses very large databases but does not appear to need extensive secondary storage capabilities of its own. However, systems that make actual underwriting decisions must interact with the large set of rules and tables in the medical guide. Design is not sufficiently completed as yet to determine whether the coupling between the database and expert will need to be loose or tight [JV 84].

#### 6. RELATIONSHIP TO TRADITIONAL SYSTEMS ANALYSIS AND DESIGN

To what extent is the design of an expert system different from that of a conventional, computer-based system in a business firm? Table 1 shows the stages in the conventional life cycle for a business information system [L 85].

| LIFE CYCLE        | TRANSACTION PROCESSING SYSTEM    | PROTOTYPE  | EXPERT SYSTEM   |
|-------------------|----------------------------------|--|---|
| Inception         | Idea for system                  | Idea for system  | Idea for system   |
| Feasibility Study | Cost/Benefit;<br>Mgmt. approval  | Working prelim.<br>version                                 | R & D;<br>Set boundaries  |
| Systems Analysis  | Details of<br>current procs.     | Work with user<br>group                                    | Listen to<br>single expert                                      |
| Design            | Create new<br>system             | Iterative<br>enhancement                                   | Begin to code<br>single expert                                  |
| Specifications    | Detailed design                  | Evolving prototype,<br>or prototype as<br>model for system | Expert critique;<br>new rules added                             |
| Programming       | Code specs.                      | Code specs.  | Add rules   |
| Testing           | Debugging;<br>Acceptance testing | Debugging;<br>Acceptance testing                           | Continued<br>enhancement;<br>Problems when<br>exposed to others |
| Training          | For users                        | For users  | For non-experts?  |
| Conversion        | Scheduled                        | Gradual evolution  | Gradual   |
| Operations        | Maintenance &<br>enhancements    | Maintenance &<br>enhancements                              | Evolving expertise  |

Table 1: System Life Cycle Comparisons

### 6.1. Traditional Systems Analysis and Design

Usually someone in the firm has an idea for how a system might improve information processing or solve a problem. The systems staff conducts a feasibility study showing various costs and benefits (including intangible benefits) and management is asked to approve developing the application. During systems analysis a design team of users and analysts document the details of present information processing procedures.

The team creates the design and specifications for a new system. Programmers



develop necessary code or modify packages. The team tests the system, trains users and converts to the new application. During maintenance, the information services department makes changes and enhancements to the system, an activity that consumes on the average 50% of the discretionary resources of the department.

## 6.2. Prototyping

What is different in the development process for an expert system based on the experience with the underwriting expert? A new approach to conventional design called prototyping has received a great deal of recent attention [L 85]. Table 1 shows that a prototype differs significantly from traditional development approaches. Basically, the prototype is a working model of a system or some part of it.

In the conventional approach, the user does not receive feedback until too late in the process; often the user does not understand the system, but yet agrees to its specifications. With a prototype, the user has something concrete to which he or she can respond; the prototyping process also tends to draw the user into the design process. It is hard to avoid user involvement with a prototype.

## 6.3. Expert Systems Design

The design of an expert system comes closest to the prototyping process. Some of the major differences between prototyping and ES design result from the current R&D nature of expert system development. So few of these systems exist in a business environment that a clear methodology has not emerged for their design. Based on experience with the underwriting expert system, six issues will be addressed, below: reliance on a single expert, staffing, program structuring, user training, systems maintenance, and accountability.

Based on the experience with the underwriting expert, one significant difference with other design approaches is the reliance on a single expert. ES designers have generally been unanimous that one cannot currently design an expert system if it is necessary to adjudicate among different experts with differing opinions. In conventional design, most organizations try to include as many potential users in the design process as feasible.

An issue of major concern to businesses weighing the potential costs and benefits of pursuing ES technology is one of staffing. DP departments do not currently have staff who are familiar with AI languages or with the process and techniques of knowledge engineering. A commitment to ES development will require substantial investment either in outside consultants, training for current staff members or the hiring of new staff trained in ES development. However, this expertise is itself still rare, making it difficult to locate and then keep a staff of knowledge engineers. A final point is that the maintenance of ES's is

more akin to the learning and maturation of a human expert than to the maintenance and enhancement of traditional DP systems. Staffing over the lifetime of the usefulness of an ES, which is assumed to be rather lengthy, poses an additional problem.

With an ES the designers are using a different type of language than COBOL or even fourth generation, nonprocedural languages. There is relatively little experience with rule-based languages in the business environment. One advantage of these languages is purported to be the ease with which rules can be added as designers acquire new knowledge. Experience with the underwriter shows that some new knowledge can be significant enough to require restructuring of parts of the program even though it is written in Prolog, a rule-based language.

With an ES training is a potential problem; what happens when other experts are asked to evaluate the system? If it was designed with a single expert, will it necessarily be robust enough to pass tests designed by colleagues? How does one approach the training of users? How does the system respond when a nonexpert user presents a situation beyond its domain of expertise? Designers of commercial systems are used to developing applications with extensive error checking, sometimes well over half of the code exists to detect user errors. Will expert systems designers be able to detect and prevent user errors?

Maintenance is a major problem in the information processing industry. Expert Systems will add to this problem, not lessen it. ES's will probably become members of a class of enduring systems, applications so expensive or so complex that one could not envision replacing them completely. Good examples of this type of systems are those used by airlines for making passenger reservations. Two U.S. carriers have invested over \$200 million on their systems [GS 84]; it is inconceivable that these systems will ever be totally replaced. Instead, they will grow and be modified over time.

Expert Systems will have to receive careful maintenance. New rules will be discovered so there will have to be a staff to tend the system and incorporate new knowledge. The domain and nature of the problem addressed by the ES may change as well. The care of an Expert System will therefore probably be more demanding than the maintenance of a conventional application.

A final issue relates to the notion of auditing and accountability. Expert systems will presumably become more and more involved in making and/or assisting in the making of complex, high-risk decisions. In many applications, e.g. insurance underwriting, drug synthesis, etc., there are corporate and legal responsibilities for maintaining complete records about the decisions made and the environment in which they were made. Increasingly, lawsuits are won (or lost) with the help of such documentation. The recording and maintenance of such information about a decision as when, how, why, and under what assumptions about the "state of the world" was a decision made will increasingly become part and

parcel of the expected behavior of expert systems as they begin to become involved in the making of crucial organizational decisions.

## 7. SUMMARY AND CONCLUSIONS

ES in business differ from many of the expert systems created to date. One application for business ES's is at the transaction level, for processing more sophisticated transactions than those amenable to DP systems. Here one often finds a high volume of transactions with each transaction being of relatively low value. However, all transactions must be processed and there will be many expert decisions made in a short period of time. The insurance underwriter is an example of such a system. We believe that ES's will begin to be seen at this organizational level first, as business take preliminary steps toward the acceptance of the new technology. If they are successful at this level, ES's will begin to be applied to higher-level, more critical decisions in the organization.

The design process for an Expert System is sufficiently different from conventional design that it will be difficult for traditionally-oriented staff members to develop an ES without assistance. There is a need for educating knowledge engineers and for a methodology for ES design. An approach that follows prototyping looks appealing based on the experiences reported in this paper.

Because there are few ES in business, no one has addressed the question of implementation for these systems. In the underwriting environment, many overall implementation questions arise. Will an underwriting expert support or replace the human underwriter? Will it be used for the thousands of policy applications that require modest human expertise, eliminating a layer of underwriters at the level below the senior lay underwriter in this study? These questions will naturally arise in other business areas where ES's are being introduced; they will need to be considered carefully if business is to seriously accept the introduction of this technology.

In light of the high development costs currently incurred in the development of ES's, strategies for multiple systems using a central or "core" knowledge base need to be considered in order to make business ES's more cost effective. For example, one additional use of the underwriting expert system would be to move some of the simpler underwriting into the field. Another possibility would be to develop an underwriting training system which would draw upon the same knowledge base of expertise, but could be used to generate test case applications and score the decisions made by the underwriting trainees.

In summary, the introduction of Expert System technology into a business environment presents new and challenging problems as well as opportunities. Our experience in developing the underwriting expert illustrates both of these facets of this technology transfer. The opportunities involved in computerizing

activities that have until now resisted automation, with the potential for an increase in speed, accuracy, and competitive edge, are very real and attractive. The problems are no less real: how to incorporate techniques of knowledge engineering into environments more comfortable with traditional systems analysis and design techniques; how to minimize the cost of the knowledge acquisition and knowledge engineering process; how to integrate ES's into the wide family of cooperating systems that are today the norm in most businesses; how to manage the overall human impact of this technology on the employees and their relationship to their job environment, their colleagues, and their careers. Clearly much more research is needed on the design and implementation of ES's if this technology is to become widely adopted in the business environment.

#### REFERENCES

[An 65] Anthony, R.N. Planning and Control Systems: A Framework for Analysis, Harvard University GSBA, Studies in Management Control, Cambridge, Mass., 1965.

[CJV 83] Clifford, J., Jarke, M. and Vassiliou, Y. "A Short Introduction to Expert Systems," IEEE Database Engineering Bulletin, vol. 8, no. 4, December 1983.

[DR 84] Duda, R.O., and Reboh, R. "AI and Decision-Making: The PROSPECTOR Experience", in W.Reitman (ed.), Artificial Intelligence Applications for Business, Ablex, Norwood, NJ, 1984, pp.111-147.

[Gea 77] Gelernter, H.L., Sanders, A.F., Larsen, D.L., Agarwal, K.K., Boivie, R.H., Spritzer, G.A., and Searleman, J.E., "Empirical Explorations of SYNCHEM," Science v.197, 1977, pp.1041-1049.

[GS 84] Gifford, D. and Spector, A. "The TWA Reservation System," Communications of the ACM, 27,7, July 1984, 650-685.

[H 85] Hewitt, C. "Implications of Open Systems", in M.Jarke (ed.), Managers, Micros, and Mainframes: Proceedings of the NYU Symposium on Integrating Systems for End Users, New York University, May 1985.

[JCV 84] Jarke, M., Clifford, J., and Vassiliou, Y. "An Optimizing Prolog Frontend to a Relational Query System", Proceedings ACM-SIGMOD Conference, Boston, Mass., pp.296-306.

[JV 84] Jarke, M. and Y. Vassiliou. "Coupling Expert Systems with Database Management Systems," in W. Reitman (ed.) Artificial Intelligence Applications for Business, Ablex, Norwood, N.J. 1984, pp.65-85.

[Kea 78] Kunz, J.C. et al. "A Physiological Rule-based System for Interpreting Pulmonary Function Test Rules," Heuristic Programming Project Memo HPP-78-19.

[L 85] Lucas, H.C., Jr. The Analysis and Design of Information Systems, 2nd ed., McGraw-Hill, 1985.

[M 82] McDermott, J. "R1: A Rule-Based Configurer of Computer Systems", Artificial Intelligence, vol. 19, no. 1, pp.39-88.

[P 84] Pople, H.E., jr. "Knowledge-Based Expert Systems: The Build or Buy Decision", in W.Reitman (ed.), Artificial Intelligence Applications for Business, Ablex, Norwood, NJ, 1984, pp.23-40.

[Sh 76] Shortliffe, E.H. Computer Based Medical Consultations: MYCIN, American Elsevier, New York, 1976.

[Si 60] Simon, H.A. The New Science of Management Decision, Harper & Row, New York, 1960.

[SJ 84] Sivasankaran, T., and Jarke, M. "Logic-Based Formula Management Strategies in an Actuarial Consulting System", Decision Support Systems, vol. 1, no. 3, 1985.

[VCJ 85] Vassiliou, Y., Clifford, J., and Jarke, M. "Access to Specific Declarative Knowledge by Expert Systems: The Impact of Logic Programming," Decision Support Systems, vol. 1, no. 2., 1985.