"MANAGING EXPERT SYSTEMS: FROM INCEPTION THROUGH UPDATING"

by
Rob R. WEITZ*
Arnoud DE MEYER**

N° 88 / 50

* Rob R. WEITZ, Assistant Professor of Decision Sciences and Information Systems, INSEAD, Fontainebleau, France

** Arnoud DE MEYER, Associate Professor of Technology Management, INSEAD, Fontainebleau, France

Director of Publication :

Charles WYPLOSZ, Associate Dean for Research and Development

Printed at INSEAD, Fontainebleau, France

Managing Expert Systems:

From Inception Through Updating

Rob R. Weitz

Arnoud De Meyer

INSEAD
Department of Technology Management

Fontainebleau, France

October 1987

Note

The idea for this paper was sparked by a project we are working on with Digital Equipment Corporation, Europe. The project entails the implementation of an expert system for computer board diagnosis across multiple sites.

The current version of this paper does not yet include along with the framework presented, a particular example, or "instantiation" of that framework. The final version is upcoming and will include examples drawn from the DEC case study.

Rob Weitz

Arnoud De Meyer

Abstract

This paper addresses the problem of managing the development and implementation of a large expert system in an organization. A traditional systems analysis and design methodology is used as a framework to highlight similarities and differences in managing large scale traditional computer based projects and large expert systems. As a non-technical, prescriptive guide, this article focusses on defining at each stage in the project, the tasks to be accomplished, resources required, impact on the organization, likely benefits and potential problems.

Introduction

Research in the field of Artificial Intelligence (AI) signals great promise for the next generation of hardware and software. At the present time, expert systems are arguably the most commercially successful product of Artificial Intelligence research; they have crossed the threshold of the laboratory and are beginning to make their presence felt in real world applications. To date however, the management of their introduction into the workplace and their impact once there remain largely unexplored. This paper presents a prescriptive guide for managing large scale expert systems from inception through implementation and maintenance. It is motivated by experience with the development of a large expert system by an international company, the growing literature of cases describing expert systems in practice, and the belief that management and organizational considerations (as opposed to technical wizardry) must remain paramount for the success of such systems to be achieved. The thrust of this paper is to focus awareness on 1) the processes and resources required for an expert system project, 2) the costs and benefits of such an undertaking, and 3) organizational and task changes likely to result from the introduction of the system. This paper is <u>not</u> a technical guide for building expert systems; technical concerns are expressed here only insofar as they are inextricably linked to the management of such systems.

Expert Systems and Information Systems

A great deal of experience has been gathered regarding the design, implementation and maintenance of "traditional" computer based systems, from both technical and management viewpoints. The pitfalls, players and positive practices have been identified in a large body of existing research, and are well described in the information systems literature. The management of expert systems does not lie completely independent from this previous computer based project management experience. Indeed, keynote speakers at recent, major AI conferences have repeatedly emphasized that expert systems will become invisible - that is, they will be part and parcel of traditional systems. This paper builds on the existing groundwork by taking as a framework a traditional systems analysis and design (SA&D) methodology and adapting it for application to expert systems.

Expert Systems

Expert systems (ES's) are computer programs for solving difficult, "fuzzy" problems in domains where human expertise is normally associated with a great deal of training and experience. Application areas to date include such areas as medical diagnosis, chemical analysis and computer system configuration. Expert systems are typically characterized by:

performance levels at, or exceeding, those of experts in the problem domain,

the utilization of large amounts of domain specific knowledge,

the ability to use incomplete or uncertain information,

the capacity to explain their behavior (a kind of self-knowledge),

symbol manipulation, that is "reasoning" about objects, as opposed to numerical manipulation (which typifies traditional computer programs).

Expert systems are built in an iterative, incremental fashion via repeated interviews between one or more domain experts and a "knowledge engineer". Briefly, the task of the knowledge engineer is to elicit the expert knowledge, map the knowledge into a suitable structure and actually code the knowledge using appropriate software and hardware. The process tends to be tedious and time consuming and has in fact been referred to as the "knowledge engineering bottleneck". A good overview of experts systems in general, and of the building process is provided by Hayes-Roth et al. (1983).

For ES's the importance of early development of working prototypes is stressed. The prototype is incrementally improved and its capabilities expanded by repeated trials with the domain expert and actual use in a test environment. In fact, several versions of the prototype are typically successively developed, until a sufficiently evolved version is realized for possible release.

Building and implementing an expert system is both time consuming and resource intensive. While improved software environments have helped speed development, and recent experience has suggested some guidelines for easing the overall process somewhat, existing verifiable applications of expert systems suggests that, for large systems, the time required to go from prototype to implementation is typically on the order of person-years, with costs measured in at least the tens of thousands of dollars. For a case in point, see Linden (1982). Clearly, larger systems (as measured by amount of knowledge captured by the system and by number of users of the system) require more effort than smaller systems. In any case, it seems that the value of automated expertise supports the magnitude of these efforts.

A Traditional Systems Analysis and Design Framework

Creating a large computer based system for more than personal use is a complex task requiring technical skills, creativity and good management of resources. A guide for this process has been established through experience and while details vary from source to source, the overall thrust is fairly standardized in the information systems literature. One such systems analysis and design framework is provided in Figure 1, and is due to Lucas (1982). As described previously, this paper adapts the traditional systems analysis and design procedures for use with expert systems.

Figure 1

A Traditional Systems Analysis and Design Framework

(from Lucas, 1982)

Inception

Feasibility Study

Systems Analysis

Design

Specifications

Programing

Testing

Training

Conversion/Installation

Operations

The bulk of the rest of this paper traces this outline as it applies to expert systems; fundamental variations of the traditional systems analysis and design process are noted, and described or referenced.

Inception

Inception refers to the realization that a computer based system can be used to facilitate some procedure. The procedure may be one already existing in the organization, or one that is being proposed. At this stage the envisioned system is naturally somewhat murky in its details, though the overall goals - to improve some process, reduce costs, ease a production bottleneck, for example, are more clear. Several variations of possible systems are likely entertained at this point. The question proposed here is, why should the envisioned system (or one of the envisioned systems) be an expert system?

It should be kept in mind that expert systems are not cheaper or easier to build than conventional systems; it is more realistic to assume the contrary. Therefore, there should be strong, positive reasons for proposing an expert system solution. Simply stated, an expert system should be proposed if there is a good correspondence with the task under consideration and the characteristics of expert systems. Criteria for tasks well suited for expert systems have been enumerated elsewhere (see Bobrow et al., 1986 for example). These criteria are briefly summarized below.

Expert systems are used to replace or assist experts. (A rough definition of an expert is someone who can solve difficult problems more quickly and with less effort than a non-expert. Typically, expertise is acquired through lengthy training and experience and is limited for an individual to a particular field of knowledge.) Insofar as technical criteria are concerned, for applying an expert system to a problem, the task should clearly

Entail reasoning about symbols, that is, "things". Mathematical calculations, no matter how complex, may be handled with traditional programs.

Involve a well circumscribed, limited field of knowledge. Expert systems are not general problem solvers.

Be difficult, but not too difficult. As a general guideline the task should take a human expert somewhere between minutes and hours.

Given the above necessary conditions, an expert system is more strongly recommended for tasks that

Have semi-structured solution processes. Expert systems were created in an attempt to model problems that do not have algorithmic or "step-by-step" solutions. (Traditional, decision tree programs were suitable for these type of problems.) Expert systems are well suited for tasks where a body of loosely structured knowledge exists; this technology was designed for capturing the heuristics or "rules of thumb" for arriving at good solutions.

Require explanation to the user of the system's reasoning in solving the problem.

Involve reasoning with uncertain or incomplete information.

The following sampling of appropriate tasks for expert systems should help clarify these criteria: medical diagnosis, mechanical/electronic troubleshooting, credit rating, insurance adjusting, portfolio management, equipment design.

The above criteria are purely technical requirements, and define the minimum considerations for further pursuing the process of thinking about applying expert system technology. Related to the above technical concerns are more organizational requirements for an expert system to be envisioned as a solution. These include that

Human experts exist and are willing and able to participate in the project.

Human expertise is scarce and/or expensive.

Management is likely to commit the required resources to the project. (This implies that the task is viewed as an important one.)

At the inception stage, some indication that these non-technical concerns are satisfied must be sensed, or at least their negation must not be evident. It should be clear at the conclusion of the feasibility study whether or not these requirements are indeed satisfied. The next section discusses the feasibility analysis and will elaborate on these considerations and others.

Prototype

This stage does not exist in traditional SA&D; its primary purpose is to demonstrate technical feasibility of the project. Prototype in this sense means a scaled down (in scope, power or both) version of the fully envisioned system. This formal prototype stage is suggested because 1) it may be hard to know what is feasible when it comes to automating expertise without attempting a working, trial version, 2) given the newness of the technology (and the hype that surrounds it) a prototype can help set realistic expectations, 3) a demonstration can serve to garner and solidify enthusiasm for the project and 4) a phased approach involves less risk. Finally, due to the incremental nature of building expert systems and the software development tools created to support this development, prototypes may be built rather quickly.

The prototype stage begins with a formal proposal to proceed with the development of a demonstration prototype. The proposal for the prototype is a brief report which includes the areas described in the previous "Inception" section, plus an outline of how the envisioned system will work, what it will do, and the costs and benefits expected. The resources and time frame required should be broadly sketched for the entire project, and firmly specified for the demonstation prototype phase. Criteria for judging the success of the prototype should also be detailed. (The prototype should be evaluated with the goal of demonstrating technical feasibility in mind, high level proficiency in the task domain not being intended.) This proposal is a rough-cut version of the feasibility study described in the next section. The purpose of the proposal is to garner enough interest and commitment to proceed with the prototype; the decision to proceed with the full blown project will be taken when the prototype is ready for demonstration. At that time the full feasibility study should be prepared.

As discussed previously, all expert systems are built through incremental improvements of working, sub-expert versions of the system. Each of these working versions may be termed "prototype". The prototype proposal specifies a fixed date at which the system, such as it is, will be accessible for demonstration purposes. This time period is on the order of weeks or months, and depends on the resources and time constraints relating to the project.

Feasibility Study

As the prototype is realized, the full nature and scope of the envisioned system become more clear. The decision to proceed with the full project will depend on the performance of the demonstration prototype and the results of a formal feasibility analysis. This analysis is comprised of a comparative evaluation of all conceivable alternative systems. (At minimum a single new system is proposed and evaluated with respect to the existing procedures.) The ultimate purpose of the feasibility study is to select the best of all possible alternatives, but this part of the analysis serves many other purposes. Among these, it helps solidify ideas, provides a common source of reference, and serves as a focus for gaining commitment from users and management.

The contents of a standard study report are outlined below. Following each directive is a description of how that directive applies to an expert system feasibility study. The report should:

Describe the current system including a rationale for changing it.

> As the saying goes, "If it ain't broke, don't fix it!" Shortcomings in the current system and improvements provided by the suggested system should be clearly specified. Expert system proposals based on rationales of selecting a "high tech, state-of-the-art solution" should be rejected.

Explain why the particular solution is proposed, as opposed to some other alternative; already existing, similar systems both inside and outside the organization should be referenced if possible.

Why is an expert system appropriate? The previous section describes problems and environments which suggest an expert system solution.

Lay out the goals, scope and objectives of the envisioned system.

It is important to set realistic expectations for the proposed system. Terms such as "artificial intelligence" and "expert systems" tend to spur the imagination. Bear in mind that this technology typically cannot do what a human expert cannot do and that, in any case, the breadth of knowledge encompassed by the system will be limited. Specify the bounds of the proposed system; i.e., what are not the goals, objectives and expected capabilities of the system.

Describe in a detailed way what the proposed system will do. Include in this description exactly how the system will work, who will use it, how the task itself and related tasks will change.

The following questions should be addressed. Will "experts" themselves be using the system (i.e., doctors using a medical diagnosis system) or will "technicians" use the system? What skills will be required of the users? Will the users be new employees or current workers already doing the task under the present system? If current workers, how will their input into system design, and their goodwill in general be solicited? Will the system "de-skill" their task, and if so how do they feel about it? How will job responsibilities and lines of communication change? The overall technical specifications should be outlined; i.e., how the system will interface with existing information systems, data needs, input/output mechanisms, performance requirements, etc.

Detail the resources required, when they will be required and from whom they will be required.

For the proposed project, allocations should be included for knowledge engineers, domain experts, programmers, hardware, software, training, and travel expenses. Note that the participation of domain experts is crucial, while their time is typically at a premium. (Hence the value of automating their expertise.) Provisions for re-specifying their job description to include work on the new system should be detailed. Commonly, programmers are required to write user and system to system interfaces. A fuller description of all the players required for such a project and how they differ from those involved in developing a traditional system is given at the end of this section. As with traditional systems, the choice of hardware and software is one of making trade-offs. Indeed, much of the criteria (price, speed, compatibility, reputation of supplier, for example) are the same. Hardware selections fall into three categories: microcomputers, mini's/mainframes, and workstations. Special purpose expert system development software (typically called "shells" or "environments") exist for each category of hardware and ease the programming task considerably. It is typically the knowledge engineer's responsibility to choose appropriate hardware and software for the project. (For more details on hardware and software selection see Harmon and King, 1985, and Gilmore and Howard, 1986. For more on the specifics of the knowledge engineering process see Hayes-Roth et al., 1983.)

Provide a cost-benefit analysis. This analysis should include both tangibles and intangibles; typically some estimates are required, but the process should be supported with a sensitivity analysis.

Benefits commonly ascribed to expert systems include: preservation and dissemination of scarce expertise, relieving experts of tedious tasks and thereby allowing them to concentrate on more difficult/more interesting problems, speedier solutions and more consistent problem solving. Important costs include: personnel (i.e., knowledge engineer and expert), software and hardware (perhaps specifically for expert system development and use), and those expenses associated with training, maintenance and updating.

Define, as best as possible, evaluation criteria and agree on a how the success of the system will be measured. The evaluation criteria should be utilized during the periodic reviews.

Expert systems will typically be evaluated on a host of criteria. These include the quality of the solution, the speed of solution, the manner in which the solution is reached (transparency), breadth of knowledge, explanation and help facilities, user satisfaction, and the ease with which the system can be transported, modified and updated. The relative importance placed upon each of these criteria is determined by the project's goals and objectives. The criteria should be established by the ultimate users of the system and the managers of the function in question, in conjunction with the developers of the system.

Specify a timetable for the project, including periodic reviews and expected status of the system at each review.

As mentioned above, expert systems are developed incrementally. The first review should occur when a "somewhat competent" prototype system has been developed and some results have been obtained regarding its performance. (This review should take place roughly six months from the start of the development work.) Feedback from this review should direct further technical work on the system, while any bureaucratic difficulties which have surfaced should be aired and attended to. Attendees at the meeting should include those involved with development, representatives of the ultimate users of the system, and those responsible for the direction of the project. At minimum, one more review will be required (on the order of 6 months after the first meeting) to decide to either release, delay or kill the project.

Assign ownership of the system over the course of it's lifetime.

In many applications, updating of the knowledge encompassed by the system is a large, ongoing job due to the nature of the task in question. Consider, for example, a diagnosis expert system for some large piece of machinery. Frequent updates will be required if new models of the machine are produced, parts (or part numbers) change, and new faults (and procedures for finding them) appear as the machine ages. Responsibility after release of the system may or may not rest with the developers, but in any case this responsibility and the mechanisms for undertaking the maintenance task should be made explicit.

Many of these points will be elaborated upon in the subsequent sections.

Some comments about the players involved in expert systems development and how they compare with those for traditional systems are in order. In traditional systems, three major players are identified: users, the information systems (IS) department, and management. Much has been written of the responsibilities of each of these agents over the course of the life of the information system project (see Lucas, 1982 for example). Briefly stated, ideally the user participates in the design process and may in fact originate the project idea. The user is best equiped to understand the workings of the present system and therefore to provide input for system specifications, and good test examples. Clearly user satisfaction with the system is an important criterion for success. The IS department reviews the feasibility study, designs the system, specifies system alternatives and the trade-offs they imply (languages, "off-the-shelf" or special purpose programmes, use of service bureaus, batch versus time sharing mode, hardware, etc.), handles the required programming, documentation, training, conversion and maintenance. Management responsibilities are overall approval and direction of the project, and providing commitment in terms of resources and recognition.

Expert systems projects include additional players. First, the domain expert is not necessarily a typical user of the system. If the domain expert is in fact just an "average user", this user's involvement in the development of the system will, in any case, be much more intimate that of typical user involvement in traditional systems. Second, the position of knowledge engineer requires skills that are not necessarily found in IS departments. This role may be filled by grooming in-house personnel, or through external services. These services may be obtained via independent consultants, expert system software companies or expert system hardware companies.

In traditional system development, the IS department has the responsibility for choosing outside vendors and services. The situation is analagous here, with many of the trade-offs (such as reputation of the company, cost, type of service provided) being similar. However, the IS department must be knowledgeable about ES technology, and the market in order to evalulate these trade-offs. Further, the IS department will have to work with the external organization to promote a smooth interface between existing systems and procedures and the new system.

If an outside organization is contracted for developing the system, arrangements for transfering ownership to the IS department (and what that ownership entails) should be specified.

Analysis, Design, Specifications, and Programming

In developing traditional information systems, each of the steps of analysis, design/specification and programming is separate and distinct; one does not proceed to a subsequent step until finished with the previous one. In the analysis phase, the current system is studied; transactions, data volumes, decisions made, etc., are detailed. Design/specification of the new system involves enumerating hardware and software, input and output files, media, procedures for use, security and error control. Writing and testing the program follows the design phase.

For expert systems, some analysis, design and detailing of specifications is required for the prototype proposal, and a large part of these tasks plus some programming is completed as part of the prototype system and feasibility report. These processes however are far less differentiated in expert system development than for traditional systems. In fact, these processes more closely resembles a paradigm for decision support system development. (See Sprague and Carlson, 1982 and Keen and Scott Morton, 1978 for example.)

The incremental, iterative procedure for building ES's has already been described. This is a function of the fact that experts simply cannot sit down and fully and completely specify their own problem solving behavior. Expert systems are constructed as a collaborative and iterative effort between one or more knowledge engineers and one or more domain experts. The knowledge engineer is experienced in eliciting knowledge from an expert, structuring knowledge such that in can be programmed, and then coding the knowledge. Working from the first prototype, the knowledge engineer can improve, refine and expand the capabilities of the system through further interaction with the expert. This loop: eliciting, structuring and coding knowledge is traversed many times before a suitable for release version is produced. Here, programming may be considered to be a function of the analysis, design and specification of the system and visa versa.

At the beginning of the development of the system, infrequent contacts between the knowledge engineer and the expert will suffice; as development proceeds, more frequent and more lengthy meetings will be required. In any case, it is crucial that the expert has sufficient time and incentive to work with the knowledge engineer on the system.

Testing

The project report which includes the feasibility study should specify a timetable for release of the system, first as a trial or test version, and later as the "product" version. The test version receives limited distribution and is used to fine tune the system for general release. For each release, the system should pass the minimum requirements set up in advance in the project report.

Disagreement and uncertainty concerning the evaluation procedure is likely to exist. Knowledge engineers may be most concerned with technical criteria, and users with quality interfaces, while managers worry about how the system will help solve "business" problems. (Even a purely technical evaluation based on the quality of solutions provided by the system may be difficult to conduct. For example, experts themselves may have differing opinions as to what constitutes a "right" answer.) While disagreement may exist as to how to evaluate the system, the time to resolve such differences is prior to development. (See Hayes-Roth et al., 1983 for a good discussion of evaluating expert systems.) In any event, expert systems must be evaluated under multiple criteria. Possible evaluation criteria include whether use of the system:

Results in better solutions.

Results in faster solutions.

Results in more consistent solutions.

Provides greater worker satisfaction.

Is easy to use.

Reduces training time (for the human users).

Reduces dependence on scarce individuals.

Has resulted in greater insight into the problem.

Reduces the number of extremely bad solutions.

These criteria, should be easy to derive from the goals of the system as stated early in the project. The more difficult part is measuring how the system fulfills these criteria. While some data collection procedures may already exist, (time or quality standards for diagnosis tasks, as an example), data collection may have to be initiated as part of the project to allow for a "before and after" study. While the level at which the system satisfies each criteria may be flexible, other requirement may be both easy to determine and rather fixed, for example:

Time to repair must be reduced by half.

"Grade 1" technicians must be able to do the task (with the aid of the system) as opposed to the "grade 5's" who do it now.

Calls to the resident expert for help on difficult problems must be all but eliminated.

The number of people doing the task should be reduced by 10%.

During testing, the system should be pushed to discover its limits, in terms of the range of its knowledge, but also on very traditional dimensions: security, back-up procedures, and error handling, for example.

User and/or management satisfaction with the system can be assessed via formal questionaires or interviews. Acceptance will be predicated on not just how well the system does what may be narrowly defined as it's task, but also on it's ease of use, intuitiveness of the interface, appropriate help facilities, speed, and on how well it helps the user in doing the task as (s)he sees it. A more fundamental question regards not the system per se, but rather the individuals for whom the system is designed. Knowledge workers, may or may not want a machine looking over their shoulders while they do their jobs, particularly if they view the system as skill or prestige reducing. (Doctors are the classic case in point.) This possibility should be addressed early in system development, so that disaffection due to resentment does not show up among users at the evaluation stage.

As with traditional systems, both tangible and intangible outcomes will likely need to be measured; release of the system is, of course, contingent on the benefits outweighing the costs.

Training, Conversion/Installation

Training, conversion and installation are similar to that for traditional systems. Appropriate documentation should be developed by the knowledge engineer in conjunction with the IS department if necessary. Training procedures are set up in cooperation with the users and management Hopefully, conversion may be phased in, with the date for conversion agreed upon by all those involved. As with traditional systems, the question of compatibility of hardware and software is important. With expert systems, compatibility may focus on the use of AI languages or specialized workstations which need to interface with more conventional hardware and software. Smooth methods of accessing company data, and generally interfacing with existing systems is frequently a necessity. These issues should have been planned for in the design phase.

Operations

Assuming an outside agent has been involved in developing the system, the arrangement should include provisions for ongoing maintenance and of the system, in the traditional sense (i.e., information "hot-line", software revisions, etc.). Maintenance duties which are the responsibility of the IS department should be clarified.

Particular attention must be paid to the updating function as expert systems work in knowledge intensive areas, and problem solving knowledge in practical applications frequently changes. Updating may be required for knowledge proper, for data the system uses, or both. Some assessment of the extent of updating should be made early in the project, based on the nature of the task. For the XCON system, Digital Equipment Corporation's ES for configuring their Vax family of computers, some eight individuals are employed full time on the updating, that is, collecting and encoding task. This huge effort is due to the fact that the system must constantly be modified to work on new products (McDermott, 1984). An estimate of the frequency of knowledge updating and the assignment of updating responsibilities should be part of the project report.

New knowledge (and alterations to the system independent of the knowledge base) will likely be suggested by users of the system, particularly if they are experts. Some formal mechanism should be established to capture these suggestions, and system performance in general. Unfortunately the knowledge encapsulated in expert systems cannot be updated simply by adding knowledge. (Automated updating, that is, updating performed entirely by users suggesting new knowledge to the system, is an important area of research in AI. Of course, systems which learn from their own mistakes, another research area, would solve this problem.) Aside from having to put the new knowledge into a form the system can understand, the new knowledge must be tested for its affects on the existing knowledge. This is to say that individuals familiar with the details of structure and operation of the system must be involved in the updating task, in addition to domain experts. If the system was developed by knowledge engineers outside the organization, updating will have to be performed with the assistance of these knowledge engineers. Alternatively, as part of the project sufficient expertise must be brought in-house (i.e., to the IS department). Again, depending on the nature of the task, its projected future requirements, and the data involved, updating may be a considerable job.

Summary

Figure 2 provides an outline of the life cycle of an expert system, as described in the preceeding pages. These phases have been defined, and a prescriptive guide toward managing the resources and responsibilities required over the course of this life cycle was presented.

The task of managing the development and implementation of a large scale expert system is in many was similar, but in substantive ways different than that for large, traditional computer based systems. This paper has served to highlight the similarities, describe the differences and provide a rationale for the contrasts.

References

Bobrow, D.G., Mittal, S. and Stefik, M.J., "Expert Systems: Perils and Promise", Communications of the ACM, vol. 29, no. 9, 1986.

Gilmore, J. and Howard, C., "Expert System Tool Evaluation", Proceeding of the Second Annual Conference on Expert Systems Tools and Applications, Avignon, France, 1986. (Authors' address: Artificial Intelligence Branch, Georgia Tech Research Institute, Atlanta, Georgia 30332, USA)

Harmon, P. and King, D., Expert Systems: Artificial Intelligence in Business, John Wiley and Sons, 1985.

Hayes-Roth, F., Waterman, D. and Lenat, D., Building Expert Systems, Addison-Wesley, 1983.

Keen, P. and Scott Morton, M., Decision Suport Systems: An Organizational Perspective, Addision-Wesley, 1978.

Linden, E., "Intellicorp: The Selling of Artificial Intelligence", High Technology, March 1985.

Lucas, H., Information Systems Concepts For Management, McGraw-Hill, 1982.

McDermott, J., "R1 Revisited: Four Years in the Trenches", AI Magazine, Fall, 1984.

Sprague, R. and Carlson, E., Building Effective Decision Support Systems, Prentice-Hall, 1982.

Figure 2

Life Cycle of An Expert System

Inception + Prototype Proposal

Approval of Prototype Development

Working Versions (Incremental Improvements)

Demonstration Prototype + Feasibility Report

Approval of Project

Development of System for Test Release

Approval of Test Release System (Evaluation Criteria Passed)

Test System Released

Evaluation of Test System Performance

Improvement of Test System For Final Release

Release/Conversion (perhaps phased in)

Maintenance/Updating

INSEAD WORKING PAPERS SERIES

1986

86/01	Arnoud DE MEYER	"The R & D/Production interface".
86/02	Philippe A. NAERT Marcel WEVERBERGH and Guido VERSWIJVEL	"Subjective estimation in integrating communication budget and allocation decisions: a case study", January 1986.
86/03	Michael BRIMM	"Sponsorship and the diffusion of organizational innovation: a preliminary view".
86/04	Spyros MAKRIDAKIS and Michèle HIBON	"Confidence intervals: an empirical investigation for the series in the M-Competition" .
86/05	Charles A. WYPLOSZ	"A note on the reduction of the workweek", July 1985.
86/06	Francesco GIAVAZZI, Jeff R. SHEEN and Charles A. WYPLOSZ	"The real exchange rate and the fiscal aspects of a natural resource discovery", Revised version: February 1986.
86/07	Douglas L. MacLACHLAN and Spyros MAKRIDAKIS	"Judgmental biases in sales forecasting", February 1986.
86/08	José de la TORRE and David H. NECKAR	"Forecasting political risks for international operations", Second Draft: March 3, 1986.
86/09	Philippe C. HASPESLAGH	"Conceptualizing the strategic process in diversified firms: the role and nature of the corporate influence process", February 1986.
86/10	R. MOENART, Arnoud DE MEYER, J. BARBE and D. DESCHOOLMEESTER.	"Analysing the issues concerning technological de-maturity".
86/11	Philippe A. NAERT and Alain BULTEZ	"From "Lydiametry" to "Pinkhamization": misspecifying advertising dynamics rarely affects profitability".
86/12	Roger BETANCOURT and David GAUTSCHI	"The economics of retail firms", Revised April 1986.
86/13	S.P. ANDERSON and Damien J. NEVEN	"Spatial competition à la Cournot".
86/14	Charles WALDMAN	"Comparaison internationale des marges brutes du commerce", June 1985.
86/15	Mihkel TOMBAK and Arnoud DE MEYER	"How the managerial attitudes of firms with FMS differ from other manufacturing firms: survey results", June 1986.
86/16	B. Espen ECKBO and Herwig M. LANGOHR	"Les primes des offres publiques, la note d'information et le marché des transferts de contrôle des sociétés".
86/17	David B. JEMISON	"Strategic capability transfer in acquisition integration", May 1986.
86/18	James TEBOUL and V. MALLERET	"Towards an operational definition of services", 1986.
86/19	Rob R. WEITZ	"Nostradamus: a knowledge-based forecasting advisor".
86/20	Albert CORHAY, Gabriel HAWAWINI and Pierre A. MICHEL	"The pricing of equity on the London stock exchange: seasonality and size premium", June 1986.
86/21	Albert CORHAY, Gabriel A. HAWAWINI and Pierre A. MICHEL	"Risk-premia seasonality in U.S. and European equity markets", February 1986.
86/22	Albert CORHAY, Gabriel A. HAWAWINI and Pierre A. MICHEL	"Seasonality in the risk-return relationships some international evidence", July 1986.
86/23	Arnoud DE MEYER	"An exploratory study on the integration of information systems in manufacturing", July 1986.
86/24	David GAUTSCHI and Vithala R. RAO	"A methodology for specification and aggregation in product concept testing", July 1986.
86/25	H. Peter GRAY and Ingo WALTER	"Protection", August 1986.
86/26	Barry EICHENGREEN and Charles WYPLOSZ	"The economic consequences of the Franc Poincare", September 1986.
86/27	Karel COOL and Ingemar DIERICKX	"Negative risk-return relationships in business strategy: paradox or truism?", October 1986.
86/28	Manfred KETS DE VRIES and Danny MILLER	"Interpreting organizational texts.
86/29	Manfred KETS DE VRIES	"Why follow the leader?".
86/30	Manfred KETS DE VRIES	"The succession game: the real story.
86/31	Arnoud DE MEYER	"Flexibility: the next competitive battle", October 1986.
86/31	Arnoud DE MEYER, Jinichiro NAKANE, Jeffrey G. MILLER and Kasra FERDOWS	"Flexibility: the next competitive battle", Revised Version: March 1987
86/32	Karel COOL and Dan SCHENDEL	Performance differences among strategic group members", October 1986.

86/33	Ernst BALTENSPERGER and Jean DERMINE	"The role of public policy in insuring financial stability: a cross-country, comparative perspective", August 1986, Revised November 1986.
86/34	Philippe HASPESLAGH and David JEMISON	"Acquisitions: myths and reality", July 1986.
86/35	Jean DERMINE	"Measuring the market value of a bank, a primer", November 1986.
86/36	Albert CORHAY and Gabriel HAWAWINI	"Seasonality in the risk-return relationship: some international evidence", July 1986.
86/37	David GAUTSCHI and Roger BETANCOURT	"The evolution of retailing: a suggested economic interpretation".
86/38	Gabriel HAWAWINI	"Financial innovation and recent developments in the French capital markets", Updated: September 1986.
86/39	Gabriel HAWAWINI Pierre MICHEL and Albert CORHAY	"The pricing of common stocks on the Brussels stock exchange: a re-examination of the evidence", November 1986.
86/40	Charles WYPLOSZ	"Capital flows liberalization and the EMS, a French perspective", December 1986.
86/41	Kasra FERDOWS and Wickham SKINNER	"Manufacturing in a new perspective", July 1986.
86/42	Kasra FERDOWS and Per LINDBERG	"FMS as indicator of manufacturing strategy", December 1986.
86/43	Damien NEVEN	"On the existence of equilibrium in hotelling's model", November 1986.
86/44	Ingemar DIERICKX Carmen MATUTES and Damien NEVEN	"Value added tax and competition", December 1986.

1987

87/01	Manfred KETS DE VRIES	"Prisoners of leadership".
87/02	Claude VIALLET	"An empirical investigation of international asset pricing", November 1986.
87/03	David GAUTSCHI and Vithala RAO	"A methodology for specification and aggregation in product concept testing", Revised Version: January 1987.
87/04	Sumantra GHOSHAL and Christopher BARTLETT	"Organizing for innovations: case of the multinational corporation", February 1987.
87/05	Arnoud DE MEYER and Kasra FERDOWS	"Managerial focal points in manufacturing strategy", February 1987.
87/06	Arun K. JAIN, Christian PINSON and Naresh K. MALHOTRA	"Customer loyalty as a construct in the marketing of banking services", July 1986.
87/07	Rolf BANZ and Gabriel HAWAWINI	"Equity pricing and stock market anomalies", February 1987.
87/08	Manfred KETS DE VRIES	"Leaders who can't manage", February 1987.
87/09	Lister VICKERY, Mark PILKINGTON and Paul READ	"Entrepreneurial activities of European MBAs", March 1987.
87/10	André LAURENT	"A cultural view of organizational change", March 1987
87/11	Robert FILDES and Spyros MAKRIDAKIS	"Forecasting and loss functions", March 1987.
87/12	Fernando BARTOLOME and André LAURENT	"The Janus Head: learning from the superior and subordinate faces of the manager's job", April 1987.
87/13	Sumantra GHOSHAL and Nitin NOHRIA	"Multinational corporations as differentiated networks", April 1987.
87/14	Landis GABEL	"Product Standards and Competitive Strategy: An Analysis of the Principles", May 1987.
87/15	Spyros MAKRIDAKIS	"METAFORECASTING: Ways of improving Forecasting. Accuracy and Usefulness", May 1987.
87/16	Susan SCHNEIDER and Roger DUNBAR	"Takeover attempts: what does the language tell us?, June 1987.
87/17	André LAURENT and Fernando BARTOLOME	"Managers' cognitive maps for upward and downward relationships", June 1987.
87/18	Reinhard ANGELMAR and Christoph LIEBSCHER	"Patents and the European biotechnology lag: a study of large European pharmaceutical firms", June 1987.
87/19	David BEGG and Charles WYPLOSZ	"Why the EMS? Dynamic games and the equilibrium policy regime, May 1987.
87/20	Spyros MAKRIDAKIS	"A new approach to statistical forecasting", June 1987.
87/21	Susan SCHNEIDER	"Strategy formulation: the impact of national culture", Revised: July 1987.
87/22	Susan SCHNEIDER	"Conflicting ideologies: structural and motivational consequences", August 1987.
87/23	Roger BETANCOURT David GAUTSCHI	"The demand for retail products and the household production model: new views on complementarity and substitutability".

87/24	C.B. DERR and André LAURENT	"The internal and external careers: a theoretical and cross-cultural perspective", Spring 1987.
87/25	A. K. JAIN, N. K. MALHOTRA and Christian PINSON	"The robustness of MDS configurations in the face of incomplete data", March 1987, Revised: July 1987.
87/26	Roger BETANCOURT and David GAUTSCHI	"Demand complementarities, household production and retail assortments", July 1987.
87/27	Michael BURDA	"Is there a capital shortage in Europe?", August 1987.
87/28	Gabriel HAWAWINI	"Controlling the interest-rate risk of bonds: an introduction to duration analysis and immunization strategies", September 1987.
87/29	Susan SCHNEIDER and Paul SHRIVASTAVA	"Interpreting strategic behavior: basic assumptions themes in organizations", September 1987
87/30	Jonathan HAMILTON W. Bentley MACLEOD and J. F. THISSE	"Spatial competition and the Core", August 1987.
87/31	Martine QUINZII and J. F. THISSE	"On the optimality of central places", September 1987.
87/32	Arnoud DE MEYER	"German, French and British manufacturing strategies less different than one thinks", September 1987.
87/33	Yves DOZ and Amy SHUEN	"A process framework for analyzing cooperation between firms", September 1987.
87/34	Kasra FERDOWS and Arnoud DE MEYER	"European manufacturers: the dangers of complacency. Insights from the 1987 European manufacturing futures survey, October 1987.
87/35	P. J. LEDERER and J. F. THISSE	"Competitive location on networks under discriminatory pricing", September 1987.
87/36	Manfred KETS DE VRIES	"Prisoners of leadership", Revised version October 1987.
87/37	Landis GABEL	"Privatization: its motives and likely consequences", October 1987.
87/38	Susan SCHNEIDER	"Strategy formulation: the impact of national culture", October 1987.
87/39	Manfred KETS DE VRIES	"The dark side of CEO succession", November 1987
87/40	Carmen MATUTES and Pierre REGIBEAU	"Product compatibility and the scope of entry", November 1987
87/41	Gavriel HAWAWINI and Claude VIALLET	"Seasonality, size premium and the relationship between the risk and the return of French common stocks", November 1987
87/42	Damien NEVEN and Jacques-F. THISSE	"Combining horizontal and vertical differentiation: the principle of max-min differentiation", December 1987
87/43	Jean GABSZEWICZ and Jacques-F. THISSE	"Location", December 1987
87/44	Jonathan HAMILTON, Jacques-F. THISSE and Anita WESKAMP	"Spatial discrimination: Bertrand vs. Cournot in a model of location choice", December 1987
87/45	Karel COOL, David JEMISON and Ingemar DIERICKX	"Business strategy, market structure and risk-return relationships: a causal interpretation", December 1987.
87/46	Ingemar DIERICKX and Karel COOL	"Asset stock accumulation and sustainability of competitive advantage", December 1987.

<u>1988</u>

88/01	Michael LAWRENCE and Spyros MAKRIDAKIS	"Factors affecting judgemental forecasts and confidence intervals", January 1988.
88/02	Spyros MAKRIDAKIS	"Predicting recessions and other turning points", January 1988.
88/03	James TEBOUL	"De-industrialize service for quality", January 1988.
88/04	Susan SCHNEIDER	"National vs. corporate culture: implications for human resource management", January 1988.
88/05	Charles WYPLOSZ	"The swinging dollar: is Europe out of step?", January 1988.
88/06	Reinhard ANGELMAR	"Les conflits dans les canaux de distribution", January 1988.
88/07	Ingemar DIERICKX and Karel COOL	"Competitive advantage: a resource based perspective", January 1988.
88/08	Reinhard ANGELMAR and Susan SCHNEIDER	"Issues in the study of organizational cognition", February 1988.
88/09	Bernard SINCLAIR-DESGAGNé	"Price formation and product design through bidding", February 1988.
88/10	Bernard SINCLAIR-DESGAGNé	"The robustness of some standard auction game forms", February 1988.
88/11	Bernard SINCLAIR-DESGAGNé	"When stationary strategies are equilibrium bidding strategy: The single-crossing

88/12	Spyros MAKRIDAKIS	"Business firms and managers in the 21st century", February 1988
88/13	Manfred KETS DE VRIES	"Alexithymia in organizational life: the organization man revisited", February 1988.
88/14	Alain NOEL	"The interpretation of strategies: a study of the impact of CEOs on the corporation", March 1988.
88/15	Anil DEOLALIKAR and Lars-Hendrik ROLLER	"The production of and returns from industrial innovation: an econometric analysis for a developing country", December 1987.
88/16	Gabriel HAWAWINI	"Market efficiency and equity pricing: international evidence and implications for global investing", March 1988.
88/17	Michael BURDA	"Monopolistic competition, costs of adjustment and the behavior of European employment", September 1987.
88/18	Michael BURDA	"Reflections on "Wait Unemployment" in Europe", November 1987, revised February 1988.
88/19	M.J. LAWRENCE and Spyros MAKRIDAKIS	"Individual bias in judgements of confidence", March 1988.
88/20	Jean DERMINE, Damien NEVEN and J.F. THISSE	"Portfolio selection by mutual funds, an equilibrium model", March 1988.
88/21	James TEBOUL	"De-industrialize service for quality", March 1988 (88/03 Revised).
88/22	Lars-Hendrik RÖLLER	"Proper Quadratic Functions with an Application to AT&T", May 1987 (Revised March 1988).
88/23	Sjur Didrik FLAM and Georges ZACCOUR	"Equilibres de Nash-Cournot dans le marché européen du gaz: un cas où les solutions en boucle ouverte et en feedback coïncident", Mars 1988
88/24	B. Espen ECKBO and Herwig LANGOHR	"Information disclosure, means of payment, and takeover premia. Public and Private tender offers in France", July 1985, Sixth revision, April 1988.
88/25	Everette S. GARDNER and Spyros MAKRIDAKIS	"The future of forecasting", April 1988.
88/26	Sjur Didrik FLAM and Georges ZACCOUR	"Semi-competitive Cournot equilibrium in multistage oligopolies", April 1988.
88/27	Murugappa KRISHNAN Lars-Hendrik RÖLLER	"Entry game with resalable capacity", April 1988.
88/28	Sumantra GHOSHAL and C. A. BARTLETT	"The multinational corporation as a network: perspectives from interorganizational theory", May 1988.
88/29	Naresh K. MALHOTRA, Christian PINSON and Arun K. JAIN	"Consumer cognitive complexity and the dimensionality of multidimensional scaling configurations", May 1988.
88/30	Catherine C. ECKEL and Theo VERMAELEN	"The financial fallout from Chernobyl: risk perceptions and regulatory response", May 1988.
88/31	Sumantra GHOSHAL and Christopher BARTLETT	"Creation, adoption, and diffusion of innovations by subsidiaries of multinational corporations", June 1988.
88/32	Kasra FERDOWS and David SACKRIDER	"International manufacturing: positioning plants for success", June 1988.
88/33	Mihkel M. TOMBAK	"The importance of flexibility in manufacturing", June 1988.
88/34	Mihkel M. TOMBAK	"Flexibility: an important dimension in manufacturing", June 1988.
88/35	Mihkel M. TOMBAK	"A strategic analysis of investment in flexible manufacturing systems", July 1988.
88/36	Vikas TIBREWALA and Bruce BUCHANAN	"A Predictive Test of the NBD Model that Controls for Non-stationarity", June 1988.
88/37	Murugappa KRISHNAN Lars-Hendrik RÖLLER	"Regulating Price-Liability Competition To Improve Welfare", July 1988.
88/38	Manfred KETS DE VRIES	"The Motivating Role of Envy : A Forgotten Factor in Management, April 88.
88/39	Manfred KETS DE VRIES	"The Leader as Mirror : Clinical Reflections", July 1988.
88/40	Josef LAKONISHOK and Theo VERMAELEN	"Anomalous price behavior around repurchase tender offers", August 1988.
88/41	Charles WYPLOSZ	"Assymetry in the EMS: intentional or systemic?", August 1988.
88/42	Paul EVANS	"Organizational development in the transnational enterprise", June 1988.
88/43	B. SINCLAIR-DESGAGNE	"Group decision support systems implement Bayesian rationality", September 1988.
88/44	Essam MAHMOUD and Spyros MAKRIDAKIS	"The state of the art and future directions in combining forecasts", September 1988.
88/45	Robert KORAJCZYK and Claude VIALLET	"An empirical investigation of international asset pricing", November 1986, revised August 1988.
88/46	Yves DOZ and Amy SHUEN	"From intent to outcome: a process framework for partnerships", August 1988.

88/47	Alain BULTEZ, Els GIJSBRECHTS, Philippe NAERT and Piet VANDEN ABEELE	**"Asymmetric cannibalism between substitute items listed by retailers", September 1988.**
88/48	Michael BURDA	**"Reflections on 'Wait unemployment' in Europe, II", April 1988 revised September 1988.**
88/49	Nathalie DIERKENS	**"Information asymmetry and equity issues", September 1988.**