Organizational Knowledge Management

DAVID B. PARADICE
JAMES F. COURTNEY
Texas A&M University

Significant payoff potential exists for organizations that effectively exploit artificial intelligence technologies to manage organizational knowledge. The fundamental concepts of knowledge management, knowledge acquisition, knowledge representation, and knowledge utilization are examined in light of the unique characteristics of organizational problems. Key areas where these technologies can be applied in organizational settings are discussed, including document processing, organizational learning, information systems management, and group decision support systems. Managerial issues regarding worker resistance, competitive advantage, and cost of adopting knowledge management technology are also considered.

Introduction

Knowledge about how to run an organization is an enormously valuable corporate asset. Currently, this knowledge is embedded in managers' minds, and is difficult to replace as managers come and go. If this knowledge could be externalized, managed, used, and shared, the payoff to the organization could be immense.

Information systems containing this knowledge could reduce managerial training time and cost, lead to better managerial decisions, reduce the time lost awaiting routine authorizations, reduce the number of managerial layers in an organization, accelerate response times in customer service areas, and assist in scheduling who should attend important meetings as well as determining an appropriate time and place for the meeting. As information systems utilize more sophisticated knowledge management techniques, the system begins to take a more active role in the operations of the firm. Thus, employees can shift their attention from the administrative details of the firm's operations (which do not generate revenue) to providing a better product or service (which does generate revenue).

The situation regarding organizational knowledge may be compared with organizational treatment of data a few years ago. The

Manuscript originally submitted January 25, 1989; Revised April 25, 1989; Accepted June 21, 1989 for publication.

asset value of data was largely unrecognized, and data was often poorly managed, inaccessible, and unsharable. Only in the last ten to fifteen years have organizations begun to try to harness the information potential in organizational data files. Now, large-scale integrated, shared databases are becoming commonplace.

However, the potential asset value of corporate knowledge dwarfs the value of data in organizational databases. The value of organizational knowledge may be estimated in several ways. For example, organizational knowledge value is reflected in the cost to obtain patents and copyrights, which are means of protecting organizational knowledge. The salary and benefits paid to retain key personnel reflect the cost of retaining organizational knowledge. Training expenses are another organizational knowledge cost. All of these costs are incurred by organizations to advance the production and retention of organizational knowledge.

The problem, of course, is how to externalize and harness this knowledge, how to manage it, and how to make it accessible to appropriate people throughout the organization. Managing large knowledge bases is, at present, a very difficult problem —much more complex than data management. In this paper, we describe some of the aspects of knowledge management technology, which consists primarily of techniques and software for developing expert systems, and discuss how knowledge may be captured, managed, and shared in an organizational setting. Many of the techniques that we discuss are still in the laboratory stage of development. Nevertheless, these techniques show promise of payoffs heretofore unheard of in the managerial use of computers.

Figure 1: Steps in developing a knowledge management system

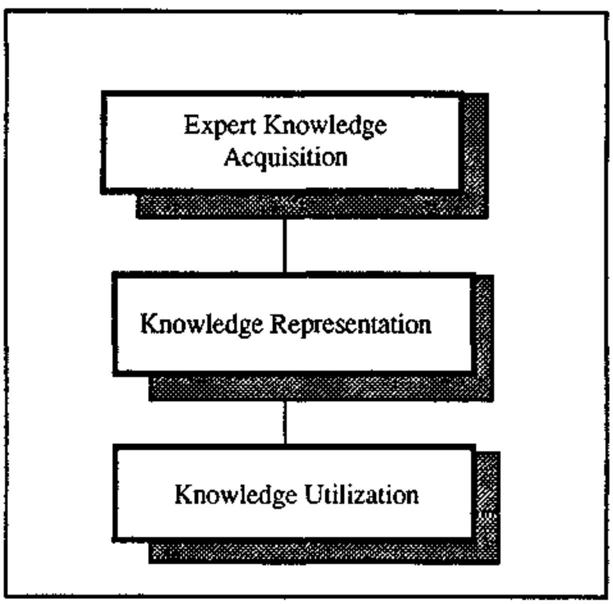

Fundamental Aspects of Knowledge Management

Understanding the potential of knowledge management technology requires an examination of three fundamental areas of knowledge work: (1) knowledge acquisition, (2) knowledge representation, and (3) knowledge utilization. (See Figure 1.)

<u>Knowledge Acquisition.</u> Knowledge acquisition involves the process of obtaining knowledge about a problem area from people, documents, books, and other sources. Typically, knowledge acquisition activities focus on "extracting" knowledge from an expert in a particular problem area. Knowledge acquisition in organizational systems is a multi-dimensional problem in that several "experts" are usually involved (Kim & Courtney, 1988).

While expertise may be easily identified in areas such as mathematics, engineering, and medicine, expertise in business is more difficult to identify. This situation is caused by the inherent characteristics of business (Paradice & Courtney, 1986). Mathematics, engineering, and, to a lesser extent, medicine, are fields that have well-formulated relationships (Basden, 1983 ??). Mathematics and engineering have infallible laws that govern relationships between components of the field. In medicine, certain symptoms indicate with high probability the existence of specific medical problems.

Business, however, has few such "laws." Although a price reduction may lead to increased sales, the impact of a price reduction depends on many other aspects of the business environment such as brand name recognition, relative market share, market maturity, and perhaps social attitudes. Because business relationships are so dynamic, identifying an expert that "knows" all the relationships is practically impossible. Consequently, one of the first problems in knowledge acquisition in business is identifying a suitable business functional area where knowledge management technology can be applied and an "expert," or a group of experts, in that area can be identified.

Expertise is developed in many ways, and often there are a combination of factors which are relevant in identifying someone as an expert. In highly technical areas, one may become an expert through specialized training. Many entry-level MIS employees today are "experts" in artificial intelligence technology when compared to co-workers that are not recent graduates. These experts will soon lose their expertise, however, if they do not use it in their job because the field is evolving so rapidly.

Experience provides another means of becoming an expert. One may develop expertise over a long period of time working in an area. This type of expertise is characteristic of engineers that must work with complex systems, such as engineers in an electric generating plant. The complexity and scope of the plant requires years of daily experience to assimilate the interrelationships between the many plant components.

Expertise need not always require years to develop, however. Persons that handle unique, crisis-oriented events may become experts rather quickly. A labor dispute arbitrator, for example, may be called to handle a critical situation that has no precedent. Yet, if successfully handled, the arbitrator can likely earn a reputation as an expert arbitrator.

As these examples show, expertise is more than years of experience. Although job experience and its resultant attributes such as job title, seniority, and company position can be a good starting point for identifying experts, one must also carefully consider the characteristics of the area of expertise. Experts can come in many forms.

But, even when experts have been identified, the problem of successfully extracting their knowledge must still be resolved. One characteristic of "expertise" is that experts "just know" what needs to be done. Novices, regardless of the area being studied, tend to follow specific rules to determine a course of action. Experts, however, have little trouble determining a course of action, but can have a much more difficult time explaining how particular choices are made. The rules over which novices labor are "second nature" for experts, sometimes to the point that the expert has difficulty verbalizing the rules (Hart, 1985). So, a second problem in knowledge acquisition is successfully extracting the knowledge from the expert.

Now, suppose that experts can be identified and their knowledge successfully extracted. A third problem is combining the knowledge of several experts. Frequently, knowledge is spread among many people in an organization. And, while the head of each functional area may be an "expert" in that area, his knowledge may also be biased by his or her perspective. Certainly, many of us are familiar with cases where two (or more) experts have analyzed a situation and reached conflicting conclusions. Because many important business decisions are made by groups, the problem of combining expertise in a business environment may be particularly acute.

Knowledge Representation. The second major area of knowledge work is knowledge representation (Dahl, 1983; McCalla & Cercone, 1983). How should knowledge be represented or organized for storage in a computer system? There are many approaches currently being investigated. One popular approach represents knowledge in the form of IF-THEN rules. This approach operates as follows: IF some condition is true, THEN some action should occur. For example, IF more than 300 units are sold, THEN discount the price by 10%.

Another business example of IF-THEN rules is illustrated in Table 1. This example involves rules to determine pricing policy based on profit margin and demand. The example uses both factual knowledge (e.g., Price is 50.00.) and process knowledge (the IF-THEN rules). Process knowledge is a critical category of knowledge. Economic and financial data are meaningless without knowledge of how to interpret them. The appropriate measures of raw materials in a product are useless in the absence of knowledge of how to combine the raw materials. When quality control data indicate that a process is out-of-control, real power lies in the diagnostic knowledge which leads to the source and ultimately to the resolution of the problem. As evidenced in these examples, factual knowledge is usually only part of an organizational knowledge base. Process knowledge is needed to completely specify what is known in organizations.

Table 1: An Example Knowledge Base for a Firm's Pricing Policy

FACTUAL KNOWLEDGE
Price is 50.00.
Cost is 45.00.
Demand is 1121.
Margin is (price - cost).

PROCESS KNOWLEDGE

1. If margin is high and demand is weak, then price-policy is decrease-price.

2. If margin is normal and demand is steady, then price-policy is maintain-price.

3. **If margin is low and demand is strong, then price-policy is increase-price.**

4. If margin greater-than 25, then margin is high.

5. **If margin is less-than 10,then margin is low.**

6. If margin is-not high and margin is-not low, then margin is normal.

7. **If demand greater-than 1100,then demand is strong.**

8. If demand less-than 900,then demand is weak.

9. If demand is-not strong and demand is-not weak, then demand is steady.

In Table 1, factual knowledge may be inserted into the IF-THEN rules to determine which pricing policy applies. Rules activated by the current set of facts are shown in boldface type in Table 1. Symbolic or textual data used in these rules is similar to numbers used in mathematical equations. The ability to handle symbolic (non-numeric) data in this manner makes expert systems appropriate for a different class of problems than purely quantitative approaches to problem-solving.

The expert system component that manipulates rules to reach conclusions is called the "inference engine," because it derives inferences from factual and process knowledge. Factual knowledge and rules constitute the "knowledge base" of a rule-based expert system.

A decision maker using an expert system with the knowledge base in Table 1 could ask questions such as: "What is price-policy?". With the present facts (and a crude explanation facility), the system could respond with "Price-policy is increase-price." Next, one could ask "Why?". The system would use the rules activated by the current facts to generate an explanation such as: "Price-policy is increase-price because margin is low and demand is strong." When asked for more information about margin, the system would respond "Margin is low because margin less-than 10." More complex systems represent knowledge at many levels and consequently are capable of more sophisticated explanations.

Another popular approach represents

knowledge as a "frame" (Winograd, 1975). Conceptually, a frame is similar to a preprinted form containing blanks, but each blank on the form has an associated "process" explaining how the blank should be filled. Frames are best suited for representing objects. Table 2 shows an example of a frame for a job shop knowledge base.

The frame in Table 2 illustrates how knowledge about a drill in a job shop could be represented. The drill is a type of machine. In our simple example, there are drill presses and hand drills. A particular drill is either one or the other, and in this example, a drill is assumed to be a drill press unless specified differently. Each drill has an identification number that must be specified. The slot (blank) for the identification number indicates that if an identification number is not specified, then a process called "create unique identification" should be invoked. The location of the drill is defined by the cell number, and the cell number must be a valid one for this shop. SHOP would be another frame type. When a repair is required, the "parts needed" slot contains a list of drill parts needed. Here, DRILL-PARTS is another frame type.

This example illustrates that knowledge can be represented at different levels: drills are a specific instance of the more general concept of machine. Successful knowledge work generally requires knowledge representation at many levels. Knowledge representation may be generally described along a continuum anchored by the concepts of surface knowledge and deep knowledge (Chandrasekaran & Mittal, 1983).

Surface knowledge is knowledge of relationships between objects without full understanding of why the relationship exists. For example, surface knowledge is reflected in knowing that if one turns a key in an automobile's ignition, the car starts. Novices typically work with surface knowledge. Deep knowledge, on the other hand, is knowledge of causality. Knowing that the key in the car's

Table 2: An Example Frame from a Job Shop Knowledge Base

Generic DRILL Frame					
Specialization of: MACHINE			Last maintenance:		
				range:	Date
Types:					
	range:	(One of: Press, Hand)	Down time for parts:		
	default:	Press		range:	Time
Identification:			Parts needed:		
	range:	Machine-identification-number		range:	(List of: DRILL-PARTS)
	ensure:	Non-null			
	if-needed:	(Process: Create-unique-identification)	Down time for service:		
				range:	Time
Power:			Total down time:		
	range:	Horse-power		range:	Time
				ensure:	(Total down time ≥ Down time for parts + Down time for service)
Maximum Bit:					
	range:	Inches-in-diameter			
Shop location:				if-needed:	(Process: Down time for parts + Down time for service)
	range:	Cell-number			
	ensure:	(Contained in: SHOP Cells)			

ignition completes a circuit containing a starter, the distributor, the spark plugs, and so on, and knowing how these objects interact reflects deep knowledge. Experts typically work with deep knowledge. Sometimes the distinction between these levels is difficult to define because the levels are complementary. Generally, if one is unable to explain why a relationship holds, one is working with surface knowledge. Recognition of deep knowledge helps to explain how experts are able to deduce solutions to problems that are not immediately obvious. The ability to recognize and work with causal relationships allows experts to reach conclusions which novices may not have rules to support.

A characteristic of artificial intelligence programming languages is the ability to support these knowledge representation approaches. Some languages work entirely with IF-THEN rules. Other languages are frame-based, while some support frames, rules, and other techniques as well. Most artificial intelligence languages also support knowledge representation at different levels (also known as support for generalization and specialization). Although artificial intelligence languages are not absolutely required for representing knowledge, they do provide a number of characteristics that facilitate knowledge work.

Knowledge Utilization. The third important facet of knowledge work, utilization, can only occur after knowledge has been acquired, organized, and stored in the machine. Then it is available for use by appropriate people throughout the organization. Reflecting back on the examples in the previous section, this knowledge might relate to the pricing policy that should be used in the current situation, or how to diagnose a problem with an inoperative drill press. Sharing knowledge such as this can be very helpful to employees who have just moved into a new position and can lead to better performance, as well as more uniform enforcement of company policy.

Knowledge use must be tempered with the realization that there are very few certainties in the business domain. Knowledge-based systems in business should emphasize this aspect by reporting some type of "confidence level" associated with system outputs. This approach is already used in medical diagnostic systems to reflect that a medical diagnosis is rarely an absolutely certain conclusion. For example, some of the indications of appendicitis may be associated with other abdominal problems. Thus, in the absence of absolute certainty, a system to utilize knowledge should emphasize the likelihood of a conclusion being correct. As another example, an expert system to predict interest rate changes might indicate that under present conditions it is "highly likely" that the Federal Reserve Board will raise the discount rate and that interest rates will increase.

There are many potential applications of knowledge-based systems in organizations. In the section that follows, we discuss some of the areas where development is currently underway.

Applications of Knowledge Management Technology

Some of the applications of knowledge management technology currently under development include document processing and electronic mail, hypertext, knowledge-based decision support systems, organizational information systems management, and group decision support systems. These applications are briefly described next.

Document Processing and Electronic Mail

Today's computer systems provide the capability to create and send documents and messages from one workstation to others. Word processing systems provide document creation and editing and electronic mail systems allow users to communicate at their convenience, as detailed messages can be stored indefinitely waiting for a response. The next wave of technology, however, will provide even greater capabilities by incorporating knowledge management techniques into word processing and electronic mail systems.

Thousands of documents are currently stored in computer-readable form as a result of widespread use of word processing and electronic mail systems. These documents contain a tremendous wealth of knowledge about the organizations they support. They range in complexity from brief memoranda and letters to reports analyzing various problems and opportunities faced by the organization. The knowledge contained in these documents represents the highest level of thinking by the most knowledgeable people in the organization. It is perhaps the most valuable of all organizational assets other than people themselves.

Of course, within an organization there is both "private" and "public" information stored in documents. Private organizational information refers to personal information, such as employee evaluations, or confidential organizational information, such as contract information. This type of information would not be generally available under normal circumstances, and knowledge management technology utilization is not a reason to make it generally available now. Techniques from database management can be applied to document systems to protect the information in these types of documents. Public organizational information, however, relates to the policies of the organization and the actions it should take to ensure success. This type of information is generally available in standard operating procedures, or via advisory actions from managers in the organization, and should be shared by appropriate people throughout the organization.

Knowledge acquisition systems could scan organizational documents and automatically extract kernels of organizational knowledge for use later. For example, some large corporations have volumes of "standard operating procedures." The technology already exists to encode the knowledge in these documents into artificial intelligence systems so that any employee could receive guidance on an organizational policy question at any time.

There are several potential benefits to the organization that has such systems. Less time would be spent by employees seeking answers to standard policy questions. Costs of maintaining standard operating procedures manuals would be reduced, because the on-line manual would be accessed by all. Guidelines for filing various organizational forms (insurance claims, grievances, etc.) could provide assistance to employees so that there are fewer filing errors. A significant benefit could be the reduced organizational cost of judgment errors made by employees. The knowledge system could provide advice concerning which actions are consistent with organizational policy.

A knowledge management technology approach provides a significant advantage over conventional information systems or database approaches. Specifically, in either a conventional information systems or database approach, all possible alternatives must be catalogued in the system in order to respond to an information request. A knowledge management approach, however, would contain an inferencing mechanism capable of deducing answers to information requests without having the answer explicitly stored in the system. For example, a user may want to know whether a specific expense is covered by his insurance policy. While a database approach would need to have the expense specifically listed as one that is covered, the knowledge management technology approach would ask the user questions

Research is currently being pursued to develop systems that are capable of constructing models of the relationships between significant components of an organization's environment (Billman, 1989). For example, a system could scan a database and deduce relationships between concepts such as sales volume, price, and advertising expenses. Once the system has developed a model of these relationships, the model is available for a number of uses. The model can be used to diagnose a managerial problem such as declining sales volume by examining underlying variables causing changes in sales volume (Ata-Mohammed, Courtney, & Price, 1988; Courtney, Ata-Mohammed, & Paradice, 1987) and can be used in an advisory mode to counsel in future managerial decision problems (Paradice & Courtney, 1987). It can be used to train new managers in various aspects of the organization's operations. The model could also be used to predict the impact of strategies in some situations, or it can serve as a basis for evaluating the organization's current condition based on data stored in the organizational database. As organizations become more complex, the cost of manually performing these functions increases dramatically.

Information Systems Management

Computer-based systems long ago reached a level of complexity that makes managing these systems a most complicated task. Some of today's systems contain millions of lines of code contained in thousands of routines, whose interaction is no longer totally known by one or a few employees. Because systems evolve over periods of years, there may be dozens of persons that have developed system components who are no longer working with the system on a routine basis. These persons may have moved up in the company, retired, or changed jobs. In any case, their expertise is not readily accessible, and the likelihood is slight that anyone has complete knowledge of the interactions of the various modules in a complex system.

Knowledge management systems will be capable of supporting management of other (possibly knowledge-based) systems. Computers will be used to solve problems associated with using computers, in this case, the problem of managing software development and maintenance (Minch & Paradice, 1987 ??). That maintenance is a large cost associated with a system's overall lifetime cost has long been known. Much of this cost is related to inefficiencies that exist in the maintenance process due to incomplete knowledge of interactions within the system.

For example, suppose a modification is required in a calculation that affects a separate calculation in another routine. If the programmer is well trained, testing will expose unanticipated results and the programmer will begin searching for the cause of the results, eventually resolving the problem. In other cases, the change will be put into production and the impact of the change will be found while the system is actually being used. In either case, completion of the modification requires more time and effort than originally scheduled (or budgeted).

The technology exists for representing the relationships between modules in a complex computer-based system. Systems should help manage systems. By constructing systems to manage systems, maintenance costs could be decreased because fewer (new) problems are introduced by the maintenance process.

Group Decision Support Systems

Because many important organizational

decisions are made by groups, a significant amount of research is being pursued to determine the most effective ways to utilize knowledge management technology to support group decision-making processes (Applegate, Chen, Konsynski, & Nunamaker, 1987; Nunamaker, Applegate, & Konsynski, 1987). The support comes in many ways.

Sophisticated systems are being developed that enhance capabilities for groups to brainstorm, share ideas, and work collaboratively. In these settings, the software can be used to develop an agenda for meetings, identify key issues, bring conflicting opinions and assumptions into the open, support decision analysis and alternative development, and automate voting and alternative ranking procedures. Group decision support software can usually operate in a manner that provides anonymity for participants during decision-making phases such as brainstorming, opinion surfacing, and alternative ranking. Anonymity encourages creativity and participation by persons that otherwise may be intimidated or afraid to voice their opinions.

Managerial Issues Related to Knowledge Management Technology

Organizational change is a phenomenon that requires careful management. Movement to knowledge management technologies will require managers to carefully manage the workers' resistance to changes in the work environment, the diminishing returns of using knowledge management technology as competitors also adopt the technology, and the costs associated with the technology.

Worker Resistance

All organizations experience some degree of resistance to changes in daily operations. One can expect resistance to incorporating knowledge management technology from managerial and staff personnel. Typically, reluctance to embrace organizational change is due to misunderstanding or misconception about the purpose of the change. When organizational changes are driven by new technologies, resistance may be caused by ignorance of the function and scope of capability of the technology.

Knowledge management technologies suffer from many misconceptions. Being developed from the research in artificial intelligence, many foresee corporate use of knowledge management technology as a means to replace employees. Certainly, there will be job skills that information systems may assume as knowledge management technology evolves. Most of these skills, however, are mundane and contain rather low task satisfaction content. Ultimately, these systems are computer-based information systems, and thus their strength comes in routine tasks such as filing, retrieving, and processing information. Even examples wherein the system may deduce an answer to an information request is still quite uninteresting compared to human-executed activities such as design, planning, and negotiation. These latter tasks are unlikely to be totally automated soon. Personal computers now permit financial analysts to examine dozens of scenarios in an afternoon in order to plan daily investment strategies, whereas twenty years ago the computations alone may have required several weeks of work. Thus, workers' fears of being replaced by knowledge management technology are generally unfounded. The technology will free workers to perform more creative tasks by supporting the workers, not replacing them.

Managers will reap the same benefits. A common managerial complaint is that too much time is spent "fighting fires." Knowledge management technology can be used for fire fighting by embedding knowledge into systems to handle routine managerial matters. Another common managerial complaint is that managers often feel too removed from the daily operations of the organization. As knowledge management technology is used to process routine managerial tasks, the organizational hierarchy will begin to flatten (fewer middle managers are needed). Therefore, management will be closer to the operational organizational level.

Another fear is that computers will constantly monitor employees (staff and management), detecting and reporting poor performers. This "Big Brother" application of the technology is unlikely, as it is economically an ineffective use of the computing resource. Additionally, there are few applications of knowledge work which lend themselves to this type of analysis. Knowledge work consists of planning, design, diagnosis, and related judgment tasks. The quality of performance in these tasks (in most cases) can only be determined after the fact. Monitoring knowledge work on a "rate of production" basis would be very difficult.

Competitive Advantage

A question frequently asked is whether firms without knowledge management technology may be able to compete solely on the basis of efficient use of their human resources. It is unlikely that a organization without knowledge management technology will be able to compete in the long run with an organization exploiting this technology. The economics of the situation preclude successful competition. For any task assumed by knowledge management technology, the cost of continuing to perform the task is negligible. Once the system is built, assuming it has been designed and implemented in a reasonable fashion, maintenance cost is unlikely to exceed the cost of retaining an employee to perform the task. The employee must be paid wages, which may increase but will rarely decrease, and must be given employee benefits. If the employee is promoted or otherwise moves out of this position, there is a cost to train a new employee to perform the task. As knowledge management technology supports more tasks, the cost discrepancy between manual and automated approaches must grow. As one includes processing speed and reliability into the analysis, the benefit of the knowledge management technology approach grows greater still.

Knowledge management technology is a technology which, once exploited by an organization, must be adopted by competitors in order to remain competitive. However, one should expect the greatest advantage to occur in the first utilization, with diminishing returns afterwards. As competitors adopt the technology, the competitive advantage gained by the first organization using the technology diminishes.

Costs

The cost of adopting knowledge management technology is substantial for projects that exceed the capability of microcomputer hardware and software. Although not absolutely necessary for knowledge management systems development, the new class of networked artificial intelligence workstation computers should be considered as the basis for these systems. These systems provide a system design, development, and implementation environment with capabilities tailored to knowledge management systems. Conservative purchase costs for these systems begins in the $60,000 range for the hardware, with an additional $50,000 for a sophisticated software development environment.

Although universities are supplying more graduates each year with the technical knowledge required to build knowledge management systems, technical expertise alone is not nearly enough to support serious utilization of knowledge management technology. Because knowledge management systems rely heavily on extracting and encoding business domain knowledge, there must be commitment from the business "experts" to the system development effort. A technical person with business experience is probably a better candidate for the development effort than an entry-level employee with the technical skills, because the learning time for acquiring technical knowledge will probably be shorter than the time needed to understand the business environment. The business trained technical person will be better able to work with the business expert. Some training will be provided with the hardware and software acquisition, but advanced training may run as much as $2,000 per class.

In addition to the costs already mentioned, Cuppello and Mishelevich (1988) advise budgeting $50,000 for consulting, $1,000 for literature, and $20,000 as a contingency reserve. They estimate that the first exploration of knowledge/expert system development in an organization can easily cost as much as $250,000.

Summary

Knowledge is an enormously valuable organizational asset. Future competitive pressures are likely to demand effective management of this asset. These demands, in turn, will require the acquisition of knowledge from documents or experts, organization of that knowledge with techniques such as frames and IF-THEN rules, and access to knowledge via expert systems supported by networks of workstations.

At present, organizational knowledge management is in its infancy and knowledge-based systems often require expensive, special-purpose workstations. Personnel experienced in artificial intelligence systems are in short supply and system development costs are high.

However, in the long run, it seems clear that organizations that learn to manage knowledge effectively will make better decisions, reduce managerial training costs, retain managerial expertise that might otherwise be lost as managers come and go, provide more uniform and more effective policies and procedures, and free managerial time for creative activities still beyond the capabilities of computers. Some areas in which benefits may be realized in the near term include document management with hypertext systems, intelligent electronic mail systems, knowledge-based decision support systems, information systems management, and group decision support systems. As with any organizational change, the implementation of these applications is likely to be met with resistance which must be managed and responded to when it has a legitimate basis.

Moreover, knowledge processing is a new frontier, and as such there are many obstacles yet to be cleared. Effective ways of acquiring knowledge to store in the new systems must be developed. Efficient ways of representing the knowledge in the system must be designed and implemented. Appropriate areas and methods for utilizing computer-based knowledge must be identified.

These are technical issues. Knowledge utilization has many interesting and complicated operational issues to be resolved as well. For example, who owns the knowledge in a knowledge-based system? If an employee is dismissed or retires, should the employee be compensated for the knowledge left behind in the system? If a company goes into bankruptcy, can the knowledge in the company's knowledge-based system be distributed to creditors as payment? Who is liable for an erroneous conclusion made by a knowledge system: the system vendor, the workers that extracted the experts' knowledge, or the experts? These are just a few of the issues to be resolved. As yet, there is no precedent providing guidance on these issues.

Research is being pursued on many fronts. Some organizations are developing their own specialized systems that meet a specific need. Other organizations are developing systems for the marketplace. Academics are pursuing fundamental research issues regarding design and development of these systems. For example, knowledge management technology systems are typically prototyped iteratively, thus their development does not follow the traditional life-cycle development approach.

The next few years will have a number of exciting developments in the knowledge-work technology area. The payoff to organizations promises to be great.

References

Applegate, L., Chen, T. T., Konsynski, B. R., & Nunamaker, J. F. (1987). Knowledge management in planning. Journal of Management Information Systems, 3(4), 20-38.

Ata-Mohammed, N., Courtney, J. F., & Paradice, D. B. (1988). IEEE Transactions on Systems, Man, and Cybernetics **(more citation data needed)**

Basden, A. (1983). On the application of expert systems. International Journal of Man-Machine Studies, 19 (5), 461-477.

Billman, B. (1989). Automated discovery of causal structure in managerial problem domains. Unpublished doctoral dissertation, Texas A & M University, College Station, Texas.

Chandrasekaran, B., & Mittal, S. (1983). Deep versus compiled knowledge approaches to diagnostic problem solving. International Journal of Man-Machine Studies, 19, 425-436.

Conklin, J. (1987). Hypertext: An introduction and survey. Computer, 20 (9), 17-41.

Courtney, J. F., Ata-Mohammed, N., & Paradice, D. (1987). A knowledge-based DSS for managerial problem diagnosis. Decision Sciences, 18 (3), 373-399.

Cuppello, J. M., & Mishelevich, D. J. (1988). Managing prototype knowledge/ expert system projects. Communications of the ACM, 31 (5), 534-542.

Dahl, V. (1983). Logic programming as a representation of knowledge. IEEE Computer, vol, 106-111.

Hart, A. (1985). Knowledge elicitation: Issues and methods. Computer-Aided Design, 17 (9), 455-462.

Kasper, G. M. (1985). The effect of user-developed DSS applications on forecasting decision-making effectiveness in an experimental setting. Journal of Management Information Systems, 2 (2), 26-39.

Kim, J. D., & Courtney, J. F. (1988). Application of knowledge acquisition techniques to decision support systems. Journal of Management Information Systems, vol, pages.

McCalla, G., & Cercone, N. (1983). Approaches to knowledge representation. IEEE Computer, vol, 12-18.

Minch, R., & Paradice, D. B. (1987). Meta-level systems for resource management. Proceedings of the Decision Sciences Institute National Conference (pp. 268 - 270), Boston.

Nunamaker, J. F., Applegate, L. M., & Konsynski, B. R. (1987). Facilitating group creativity with GDSS. Journal of Management Information Systems, 3 (4), 5-19.

Paradice, D. B., & Courtney, J. F. (1986). Controlling bias in user assertions in expert DSS for problem formulation. Journal of Management Information Systems, 3 (1), 52-64.

Paradice, D. B., & Courtney, J. F. (1987). Causal and non-causal relationships and dynamic model construction in a managerial advisory system. Journal of Management Information Systems, 3 (4), 39-53.

Winograd, T. (1975). Frame representation and the declarative/procedural controversy. In D. G. Bobrow and A. Collins (Eds.), Representation and understanding, (pp. 185 - 210). New York: Academic Press.