Working Knowledge

Thomas H Davenport
Laurence Prusak
Harvard Business School Press - 1998

Introduction
This book is a classic written by two of the world's leading experts on knowledge management. It provides a primer on knowledge management (KM). Starting with a definition of knowledge, the book goes into various aspects of KM. Davenport and Prusak make the simple but important point that knowledge is neither data nor information. Understanding what data, information and knowledge are and how individuals and companies get from one to another is essential to doing knowledge work successfully. The book, after explaining how knowledge markets work, covers the role of technology and behavioral issues in KM. The book also offers practical advice on how to implement KM in organizations.

Data
Data is a set of discrete, objective facts about events. Data can be viewed as structured records of transactions.

Data management is typically evaluated in terms of cost, speed, and capacity. How much does it cost to capture or retrieve a piece of data? How quickly can we get it into the system or call it up? How much will the system hold? Qualitative measurements are timeliness, relevance, and clarity. Do we have access to it *when* we need it? Is it *what* we need? Can we make sense out of it?

People gather data because it is factual and therefore creates an illusion of scientific accuracy. They think that if enough data is available, objectively correct decisions will become easy to make. This is false on two counts. First, too much data can make it harder to identify and make sense of a situation. Second, there is no inherent meaning in data. Data provides no judgment or interpretation. Data, by itself cannot tell us what to do. Data says nothing about its own importance or relevance. But all the same, data is important to organizations, because it is what gives rise to information.

Information
Information is a *message*. As with any message, it has a sender and a receiver. Information is meant to change the way the receiver perceives something, to have an impact on his judgment and behavior. It is data that makes a difference.

Information moves around organizations through hard and soft networks. A hard network has a visible and definite infrastructure: wires, delivery vans, satellite dishes,

post offices, addresses and electronic mailboxes. The messages these networks deliver include e-mail, traditional or "snail" mail, delivery-service packages, and Internet transmission. A soft network is less formal and visible and more ad hoc. When a colleague sends a note or a copy of an article marked "FYI", information is being transmitted through a soft network.

Quantitative measures of information management tend to include connectivity and transactions: How many e-mail accounts or Lotus Notes users do we have? How many messages do we send in a given period?

Qualitative measures help track informativeness and usefulness: Does the message give us some new insight? Does it help make sense of a situation and contribute to decision making or problem solving?

We transform data into information by adding value in various ways.
. *Contextualizing:* To know for what purpose the data was gathered
. *Categorizing:* To know the units of analysis or key components of the data
. *Calculating:* To analyse the data mathematically or statistically
. *Correcting:* To remove errors from the data
. *Condensing:* To make the data available in a more concise, user friendly form

Knowledge
Knowledge is broader, deeper, and richer than data or information. Information becomes knowledge, through:

. *Comparison:* How does information about this situation compare to other situations?
. *Consequences:* What implications does the information have for decisions and actions?
. *Connections:* How does this bit of knowledge relate to others?
. *Conversation:* What do other people think about this information?

One of the reasons that we find knowledge valuable is that it is closer than data or information to action. Knowledge must be evaluated by the decisions or actions to which it leads. Better knowledge must lead to improved productivity or lower cost and must facilitate wiser decisions.

Knowledge develops over time, through experience, which provides a historical perspective from which to view and understand new situations and events. Knowledge born of experience recognizes familiar patterns and can make connections between what is happening now and what happened in the past. Experience changes ideas about what *should* happen into what *does* happen.

By helping us deal with complexity, knowledge provides value. It is tempting to look for simple answers to complex problems and deal with uncertainties by pretending they don't exist. Knowing more usually leads to better decisions than knowing less, even if

the "less" seems clearer and more definite. Certainty and clarity may seem convenient but they often come at the price of ignoring essential factors.

Unlike data and information, knowledge contains judgment. Not only can it judge new situations and information in light of what is already known, it judges and refines itself in response to new situations and information. Knowledge can be likened to a living system. Knowledge grows and changes as it interacts with the environment.

Knowledge sometimes works through rules of thumb which, help us to solve quickly to new problems that resemble problems previously solved by experienced workers. Those with knowledge see known patterns in new situations and can respond appropriately. They don't have to start from scratch every time. So knowledge allows people to deal with situations quickly, even some very complex ones that would baffle a novice.

Knowledge Markets

Knowledge is exchanged, bought and bartered. Like markets for goods and services, the knowledge market has buyers and sellers who negotiate to reach a mutually satisfactory price for the goods exchanged. There are also brokers who bring the buyers and sellers together. Knowledge market transactions occur because the participants believe that they will benefit from them in some way.

Knowledge initiatives that ignore market dynamics and the social and political realities will fail. If the political reality of an organization allows hoarders of knowledge to thrive, then there is no incentive for people to share their expertise. Knowledge exchange will be minimal. If it is considered a sign of weakness or incompetence within the culture of an organization to admit you can't solve a problem on your own, then the social cost of "buying" knowledge will be too high. Once again, the knowledge market won't operate well.

Knowledge buyers are usually people trying to resolve complex problems by looking for insights, judgments, and understanding. They seek knowledge to make a sale or accomplish a task more efficiently; to improve their judgments and skills and help them make better decisions. In short, they want knowledge to do their work more effectively.

Knowledge sellers are people with an internal market reputation for having substantial knowledge about a process or subject. Although virtually everyone is a knowledge buyer at one time or another, not everyone is necessarily a seller. Some people are skilled but unable to articulate their tacit knowledge. Others have knowledge that is too specialized, personal, or limited to be of much value on the knowledge market.

A knowledge seller is typically motivated by one or more of three factors: work, reciprocity, repute, and altruism.

Knowledge sellers will spend the time and effort needed to share knowledge effectively if they expect the buyers to be willing sellers when they are in the market for their knowledge.

Knowledge sellers usually want others to know them as a knowledgeable person with valuable expertise. Having a reputation for knowledge sharing makes achieving reciprocity more likely. Being known as a knowledge seller makes one a more effective knowledge buyer. Having a reputation as a valuable knowledge source can also lead to job security, promotion, and all the rewards and trappings of an internal guru. Although sellers may not receive cash directly, they may receive a higher salary or bonus from sharing knowledge with others. In many consulting firms, consultants' bonuses are tied to demonstrated knowledge generation and transfer. In any organization, however, the value of reputation in the knowledge market will depend on the political and social structures of the organization.

Altruism may also motivate knowledge sharing. After a certain age, some people have an urge to pass on what they have learned to others. Firms can encourage this tendency by formally recognizing mentoring relationships, giving managers time to pass on their knowledge, and understanding that experienced employees *have* valuable knowledge to foster mentoring. Many firms ignore the contribution that older workers can make to their younger colleagues.

Knowledge brokers play an essential role in the knowledge market. Because they have broad, boundary-spanning interests, many view them as unfocused or undisciplined, or even "nosy" or "gossipy." For example, firms often do not realize the importance of librarians. Making knowledge connections mainly by talking to people, they are sometimes criticized for spending their time "chatting" rather than doing "real work." Since they are facilitators of other people's success, their contribution may not be visible to managers who think in terms of traditional productivity. It is hard to measure the profit but easy to identify their cost to the company. One of the first things firms do when they cut costs is to close the corporate library.

Without trust, knowledge markets will not work well.

1. Trust must be visible. The members of the organization must actually see people get credit for knowledge sharing.

2. Trust must be ubiquitous. If part of the internal knowledge market is untrustworthy, the market becomes asymmetric and less efficient.

3. Trustworthiness must start at the top. Trust tends to flow downward through organizations. If top managers are trustworthy, trust will seep through and come to characterize the whole firm. If they cynically exploit others' knowledge for personal gain, distrust will spread throughout the company.

Informal markets play an important role in the buying and selling of knowledge. Probably the best knowledge market signals flow through the informal networks of practice that develop in organizations. Within these webs, people ask each other who knows what-who has previously provided knowledge that turned out to be reliable and useful. If the person you approach doesn't know an appropriate seller, she probably knows someone else who does know.

Informal networks engender trust because they function through personal contact and word of mouth. A recommendation that comes from someone we know and respect within the firm is more likely to lead us to a trustworthy seller with appropriate knowledge than would a cold call based on the organizational chart or corporate phone directory. Such informal networks are also dynamic. Because they consist of people more or less continually in communication with one another, they tend to update themselves as conditions change. People share information about who has left the company or moved to new projects, who has recently become a surprisingly useful sources of knowledge, and who has become unexpectedly reticent.

The main disadvantage of informal networks is that, they are not readily available to all who need them. Their viability depends on chance conversations and local connections that sometimes work well but not so on other occasions.

Making knowledge markets more efficient

Three factors in particular often cause knowledge markets to operate inefficiently in organizations: the incompleteness of information about the knowledge market; the asymmetry of knowledge; and the localness of knowledge.

Incompleteness. Much of the interest in KM arises when firms realize they do not know where to find their own existing knowledge. Knowledge transactions may also be inhibited by uncertainty about what the likely return on shared knowledge will be.

Asymmetry. There often exists abundant knowledge on a subject in one department of an organization and a shortage somewhere else. A certain amount of asymmetry must exist in any market. But too much asymmetry can create serious problems in transferring knowledge.

Localness of Knowledge. People usually get knowledge from their neighbors. The knowledge market depends on trust, and individuals generally trust the people they know. Face-to-face meetings are often the best way to get knowledge. Reliable information about more distant knowledge sources is usually not available. Also, mechanisms for getting access to distant knowledge tend to be weak or nonexistent. People will buy whatever knowledge the person in the next office may have rather than try to discover who in the company may know more.

The hoarding of knowledge arises due to various reasons. The not-invented-here mentality is a refusal to buy knowledge. A variation is the class barrier, an unwillingness to give knowledge to or accept it from people in the organization who have relatively low status. A barrier may also be established by an executive who has the power to enforce a corporate orthodoxy by banning subjects that threaten it. Trade barriers can also arise when companies lack a good knowledge transfer infrastructure like an effective computer network or communications system or meeting places.

Electronic knowledge markets create convenience and choice. But due to variable quality and a lack of personal contact, trust and commitment are reduced. In the electronic home shopping industry, the result is a lot more browsing than buying. Online knowledge, is often ignored or treated with suspicion unless it has been evaluated and edited by a respected online broker.

A firm's investment in knowledge exchange is another signal that it genuinely values knowledge. Putting highly regarded people in full time knowledge-enabling jobs, holding well-attended fairs and forums, and giving people time to learn and exchange knowledge are signals of commitment.

Codification

Codification aims at putting organizational knowledge into a form that makes it accessible to those who need it. It attempts to make knowledge as organized, explicit, portable, and easy to understand as possible.

Since the purpose of codification is to put knowledge in a usable form, the corporation needs some idea of what uses it has in mind. The definition of usefulness should not be too narrow, however.

Finding the sources of the knowledge we want to codify is obviously essential. Once found, someone must evaluate the knowledge to assess its usefulness and importance to the organization, and to determine what kind of knowledge it is. The rich, tacit, intuitive knowledge of a seasoned expert, developed and internalized by the knower over a long period of time, is almost impossible to reproduce in a document or database. Rules-based, explicit knowledge is easier to document.

Some knowledge that is quite complex and initially tacit can be externalized and embedded in a company's products or services. The knowers can use their expertise to develop a process or product that contains at least some of what they know. Any manufacturing process, is constructed from what was once the knowledge of individuals. In theory, this embedded knowledge is independent of those who developed it and therefore has some organizational stability. An individual expert can disappear without bringing the process to a halt or reducing the company's stock of embedded knowledge. In practice, however, it is difficult to locate the dividing line between knowledge that is fully embedded in a process and the tacit, human knowledge

that keeps the process going.

Some forms of knowledge, like patents are already codified and explicit. A patent represents knowledge that is protected by being publicly described and connected with an owner. By definition, patented knowledge is knowledge that can be explicitly expressed. Similarly, reports and other structured documents are examples of knowledge that has already been made explicit.

Codifying knowledge is an essential step in leveraging its value in the organization. Codification represents or embeds knowledge in forms that can be shared, stored, combined and manipulated in a variety of ways across the organization. The challenge is to codify knowledge and still leave its distinctive attributes intact, putting in place codification structures that can change as rapidly and flexibly as the knowledge itself.

One way to deal with this problem is that instead of trying to turn knowledge into a "code", we can often *encode* the stories themselves so as to convey meaning without losing much of its value. For example, managers can prepare a video that tells the story of an important business event, such as how a key sale was made. Knowledge is more likely to be absorbed if it is delivered with feeling, and is placed in a context or frame that the audience can relate to.

Knowledge maps

A knowledge map points to knowledge. Developing a knowledge map involves locating important knowledge in the organization and then publishing some sort of list or picture that shows where to find it. Knowledge maps typically point to people as well as to documents and databases.

The main benefit of a knowledge map is to show people in the organization where to go when they need expertise. Rather than making do with accessible but imperfect answers, the employee with a good knowledge map has relatively easy and quick access to knowledge sources that would otherwise be difficult or impossible to find.

A firm's organizational chart is a poor substitute for a knowledge map. Most organizational charts are hierarchical, describing formal reporting structures with far more detail at the top than the bottom. But key knowledge may exist anywhere in the company. Effective knowledge seekers almost always need to cross departmental boundaries and ignore reporting structures to get what they need.

At least initially, knowledge mapping should focus on a clearly defined need. The information needed to create a knowledge map is usually in a fragmented and undocumented form. Every employee has a little piece of the map in her head, knowing about her own expertise and where she goes to get certain questions answered. Creating an organizational map is a matter of combining these individual "mini-maps." Organizations that develop knowledge maps often use surveys that ask employees what

knowledge they have and where they get the knowledge they need to do their jobs.

Technology can play a major role in constructing knowledge maps. On-line Yellow Pages can be made accessible to people across the organisation. They allow users to search by topic or key word, making it easy to locate and compare potential knowledge sources. Most important, an electronic map can be revised more frequently than a printed one. Since successful knowledge transactions depend so heavily on trust and compatibility, personalizing the entries can make the map more effective. In many companies, Knowledge Yellow Pages show an image of the person listed. A few organizations include a brief video clip.

Organizational knowledge maps are political documents too. If knowledge is genuinely important to an organization and those who have it are recognized and rewarded, then the knowledge map will be a picture of status and success as well as a knowledge locator. If politics plays no part in a KM initiative, it only means that the organization perceives nothing of value is at stake.

Knowledge transfer

Knowledge is transferred in organizations whether or not the process is managed. But knowledge transfers often tend to be local and fragmentary. We discuss a business problem with someone down the hall because she is conveniently close and we feel comfortable with her. Not necessarily because she is the best person to consult on the subject. We rarely try to find the person in the company who has the deepest knowledge of the subject. We merely hope to get good enough information from someone nearby.

Spontaneous, unstructured knowledge transfer is vital to a firm's success. Although the term "KM" implies formalized transfer, one of its essential elements is developing specific strategies to encourage such spontaneous exchanges. This is particularly necessary for organizations whose primary role is to create knowledge.

Transferring knowledge through personal conversations is being threatened by the move to “virtual offices,” where workers are encouraged to work at home or at a customer site. While these arrangements offer benefits such as greater employee flexibility, it also lowers the frequency of informal knowledge transfer. Firms that initiate virtual office programs should at least encourage workers to be in the office on the same days, identify ways to make up for lost interaction and encourage workers to stay in touch through virtual interaction.

Informal knowledge transfer is endangered by a wrong notion of what is and isn't "real" work. An employee who dutifully reads and answers e-mail messages is supposedly hard at work. On the other hand, an employee who reads a book at his desk is looked at with suspicion. A company that claims to value knowledge but discourages reading effectively sends a message that knowledge is not much valued after all. Indeed, the

availability of "slack" time for learning and thinking may be one of the best metrics of a firm's knowledge orientation.

Conversations should never be seen as a spare time activity. They must be encouraged to flourish. But conversations may not necessarily ensure that an innovation developed in one part of the world will be adopted in another. So we also need to consider more formal and deliberate ways of sharing knowledge within organizations.

Tacit knowledge transfer generally requires extensive personal contact. The "transfer relationship" may be a partnership, mentoring, or an apprenticeship, but some kind of working relationship is usually essential. Companies committed to transferring tacit knowledge often set up formal mentoring programs and make passing on knowledge to young employees an explicit part of the job description of skilled senior staff. As a general rule, the more rich and tacit knowledge is, the more technology should be used to bring people together to share that knowledge directly. It's not a good idea to try to contain or represent the knowledge itself using technology.

Effective knowledge transfer involves both transmission and absorption, followed by some change in behavior, or the development of some new idea. It is fairly common for someone to understand and absorb new knowledge but not put it to use for a variety of reasons. Not respecting or trusting the source of the knowledge is an important one. Pride, stubbornness, lack of time, lack of opportunity, a fear of taking risks are others. Our self-esteem is based on what we know and how we've done things in the past. We are likely to resist when someone points out a better way of doing work.

Velocity & Viscosity

Velocity is the speed with which knowledge moves through an organization. How quickly and widely is it disseminated? How quickly do the people who need the knowledge become aware of it and get access to it? In recent years, IT has played a major role in enhancing velocity.

Viscosity" refers to the richness (or thickness) of the knowledge transferred. How much of what we try to communicate is actually absorbed and used? To what extent does the original knowledge get pared down? Does what was absorbed have resemblance to what we tried to transmit? Viscosity is influenced by a number of factors, especially the method of transfer. Knowledge transferred by means of a long apprenticeship or mentoring relationship is likely to have a high viscosity. The receiver will gain a tremendous amount of detailed and subtle knowledge over time. Knowledge retrieved from an on-line database or acquired by reading an article will be much thinner.

Velocity and viscosity are often at odds. Enhanced velocity may come at the price of reduced viscosity. The core of knowledge transfer is about striking an optimum balance between these two factors.

Organising knowledge management

Few organizations have many workers who are skilled at framing and structuring their own knowledge. Even fewer of them have the time or inclination to sit down and feed it into a database. Organizations need people who will extract knowledge from associates, put it in a structured form, and maintain or refine it over time. Employees in dedicated roles with specific responsibilities must therefore play an important role in coverting data and information into knowledge.

On the other hand, KM will not take off if it is solely the responsibility of a small-or even a large-staff group. Ultimately, managers and workers who do other things for a living have to do the bulk of the day-today KM activities.

Chief Knowledge Officer (CKO) positions require a blend of technical, human, and financial skills. A good CKO combines an orientation to technology-based, explicit knowledge with a feel for the cultural and behavioral factors that impede or enable knowledge. Faith in the virtues of knowledge must be combined with a hard-nosed business sense.

The actual designation used, also reveals a lot. If a company calls the position "CKO," it is likely that a primary focus of the role is capturing and leveraging structured knowledge, with information technology as a key enabler. Managers with these roles often come from technology-oriented backgrounds, though they also typically have experience in cultural and organizational change.

If the designation is "Chief Learning Officer" (CLO), the key focus of the job may have more to do with training and education than with capturing and leveraging structured knowledge. It's also likely to involve the human resources function more than the information systems group. Most of these executives have responsibility for executive development at a minimum, and sometimes all employee training.

CKO roles are particularly appropriate in firms where knowledge is a critical business resource. Professional service firms have this characteristic. CKOs are also appropriate in businesses where knowledge is embedded in the products sold, or critical to the services offered to customers. This is a likely situation, for example, in the computer industry.

Still, creating a CKO role is not for every firm. Even in companies where KM is quite popular, there may be circumstances that dictate against establishing a CKO position. The organization may have such a decentralized organizational structure that a central knowledge role would be inappropriate.

Technology

Information technology has a key role to play in knowledge management. Technologies

designed for managing data are structured, typically numerically oriented, and address large volumes of observations, and do processing without substantial human intervention. On the other hand, knowledge technologies deal most frequently with text rather than numbers, and text in relatively unstructured forms, such as clauses, sentences, paragraphs, and even stories. Knowledge technologies are also more likely to be employed in an interactive and iterative manner by their users.

There are various types of knowledge technologies. Some technologies involve participation by broad groups in the use of knowledge; others involve only a few individuals. Some knowledge tools effectively require that the user be something of an expert on the topic. Others assume that the user is a more passive participant in the knowledge process. Some knowledge-work environments allow time for search, synthesis, and reflection. A good example is an academic researcher. Others, require real-time or near real-time performance. A good example is a doctor.

In the past, repositories contained largely competitive intelligence, market knowledge, or external technical, legal, or commercial knowledge. Now, however, many firms are creating repositories of internally sourced, structured knowledge. They are creating repositories of internal product knowledge, marketing knowledge, customer knowledge, etc.

Lotus Notes and Intranet-based Webs are the two leading toolsets for managing knowledge repositories today. Notes excels at database management, discussion-group creation and management, and replication of databases for remote disconnected use in the field. Notes is particularly appealing in professional services because work in that industry often involves travel to the client site, and the replication feature in Notes allows a remote employee to quickly download all new items added to databases of interest and then to peruse them off-line.

The Web is ideal for publishing information across multiple types of computer platforms, for multimedia databases, and for displaying knowledge that is linked to other knowledge through hypertext links. The Web is a very intuitive technology, and deals easily with audio, graphic, and video representations of knowledge. The hypertext structure of the Web makes it very easy to move from one piece of knowledge to another. Most Web-based repositories are smaller and easier to negotiate than those built in Notes. Intranet Webs are therefore the easiest way to get KM started.

If a company plans to use Web technology for knowledge management, a Web browser and server software alone do not suffice. A complex suite of tools is normally necessary to capture the information, store it, and allow broad access. Hypertext Markup Language (HTML) publishing tools for producing Web documents, a relational database system for storing them, text search-and-retrieval engines, and some approach to managing the "metaknowledge" that describes and facilitates access to the knowledge available, are needed.

Another requirement for search-and-retrieval knowledge management is the development of an on-line thesaurus. Knowledge is unwieldy to structure. Searchers will be looking for knowledge using terms that can't always be anticipated. The idea behind a thesaurus is to connect the terms by which the knowledge has been structured, with the terms employed by the searcher.

The underlying technique for both Web and Notes-based knowledge repositories is text search-and-retrieval. While this technology has been around for decades, it has both strengths and shortcomings for knowledge management. On the positive side, the knowledge itself typically has plenty of meaningful context that was created by the original author of the article. However, the knowledge in textual databases is indexed on the basis of keywords and their proximity in the text. These are relatively shallow aspects of the knowledge. So it is often difficult to extract knowledge in search queries on this basis.

Notes or the Web work well for broad knowledge domains when there is no right answer to a problem or when there are many different answers scattered around the organization. The use of these tools requires substantial user time (to search the database and read the retrieved knowledge) and intelligence (to synthesize and interpret the retrieved knowledge). Not all KM environments are blessed with these conditions.

Early on in the life of KM initiatives, a "let a thousand flowers bloom" technology strategy may be helpful. Later on, however, the sharing of knowledge across organizational boundaries will be easier with a single, broadly employed toolset.

Some organizations have concentrated knowledge domains rather than a community of expert users. This is the best situation for expert systems. The user normally needs to engage in a dialogue with the system, entering information about the problem or situation, a process that takes time. Expert systems, which are typically structured in a set of rules, can perform very complex reasoning. However, it can be difficult to extract knowledge from experts either because they do not know what they know, or because they don’t want to surrender the knowledge. For this reason the rules governing the expert system must be carefully specified. Expert systems have another limitation. These highly structured systems are difficult to maintain or add knowledge to. So the knowledge domain needs to be fairly stable.

Companies with focused knowledge environments can also look at constraint-based systems, which are suited for situations with high levels of data but normally less quantitative data than that required by neural networks. Like expert systems, they are suited for relatively narrow problem domains, such as product configuration or pricing. Constraint-based systems capture and model the constraints that govern complex decision making. Because constraint-based systems are usually object-oriented

underneath (rather than rule-based), they are easier to modify than expert systems. There are no complex interactions to understand and modify.

If the time available is less and users are smart, the KM tools described above will be less appropriate. Take customer support or "help desk" applications, for example. The customer is on the telephone in real time. In this situation there are a couple of options. If the users are capable of understanding problems, but not normally of solving them or classifying their symptoms, case-based reasoning (CBR) may be the best bet. CBR applications require someone to input a series of "cases," which represent knowledge about a particular domain expressed as a series of problem characteristics and solutions. Then when a customer analyst is presented with a problem, its characteristics can be compared against the set of cases in the application, and the closest match can be selected. CBR is a branch of artificial intelligence that is most commonly found in the customer service and support processes in firms.

If plenty of time is available and a user is highly qualified such as a Ph.D. in statistics, neural networks may be the right way of turning data into knowledge. A neural network is a statistically oriented tool that excels at using data to classify cases into one category or another-say, whether a loan customer is likely to default on a loan, or pay it back. As these systems "learn", their classification becomes more accurate with more cases. Neural networks require a lot of data and a high-powered computer. They can yield very accurate classifications of cases even with many interrelated variables. Because setting up the analysis and interpreting results can be very tricky, these systems require a very knowledgeable user, at least to set up the initial model. Subsequent data (for example, a month's new scanner data in a consumer products firm) may be analyzed with the same model, so converting the data into knowledge can happen faster and with less expertise. Still, in order to make decisions based on the recommendations of neural networks, it's very helpful to know how they work.

However, neural networks are something of a "black box". It's not easy to explain why they did what they did. A particular case will be classified in a particular fashion according to nodes and variable weightings, and is therefore difficult to interpret. Some new neural networking tools, hide the complexity from the user and are able to explain to some degree why the system did what it did. Smart businesspeople nevertheless may not like them because of difficulties in interpretation.

Technology cannot create new knowledge. The technological support for knowledge creation may improve in the future, but is still insignificant today. But if the appetite, the skills, and the attention to knowledge are already present in an organization, technology can expand access and ease the problem of getting the right knowledge to the right person at the right time. Technology can also raise the motivation to share KM. When people see their company investing time and money on its Web site, for example, they may take KM more seriously.

Much of the energy in KM has been spent on treating knowledge as an "it," an entity separate from the people who create and use it. The typical goal of this type of project is to take knowledge embodied in documents-memos, reports, presentations, articles, and so forth-and put it into a repository where it can be easily stored and retrieved. A somewhat less structured form of accumulated knowledge is the discussion database, in which participants record their own experiences on an issue and react to others' comments. The metaphor of a library is useful for conceptualizing knowledge repository projects.

Another type of project concentrates on providing access to knowledge or facilitating its transfer among individuals. In knowledge access projects, the focus is on finding the person with the knowledge one needs, and then transferring it. "Knowledge Yellow Pages" might best symbolize the purpose of knowledge access projects.

The last type of project attempts to establish an environment conducive to knowledge management. Within this category, there are examples of projects intended to measure or improve the value of knowledge capital, efforts to build awareness and cultural receptivity, initiatives attempting to change behavior as it relates to knowledge, and attempts to improve the KM process.

Measuring the outcome

What constitutes success in KM? Economic returns from knowledge may not be easy to quantify. So we must rely on more general indications of success. Some of the attributes that can be used to define success in KM are:

- Growth in the resources attached to the project, including staffing and budgets.
- Growth in the volume of knowledge content and usage (for example, the number of documents or accesses for repositories, or participants for discussion database projects).
- The likelihood that the project will be sustaining beyond a particular individual or two, that is, the project is an organizational initiative, not an individual project.
- Comfort throughout the organization with the concepts of knowledge and KM.
- Some evidence of financial return, either for the KM activity itself or for the larger organization. This linkage need not be rigorously specified and may be only perceptual.

Improving the chances of success

The following factors can contribute to the success of a knowledge project.

- A knowledge-oriented culture
- Technical and organizational infrastructure
- Senior management support
- A link to economics or industry value
- A process orientation

- Clarity of vision and language
- Some level of knowledge structure
- Multiple channels for knowledge transfer
- Linking performance appraisal system to knowledge sharing.

Culture is clearly one of the most important conditions for the success of a knowledge project. It is the hardest factor to build from scratch, and has several different components. Employees must be bright and intellectually curious. They must be willing and free to explore. Their knowledge-creating activities should be given credence by executives. Failure should not be penalized heavily.

Knowledge projects are more likely to succeed when the requisite technology and organization infrastructure is available. Of the two, technological infrastructure is easier to put in place. Building an organizational infrastructure means establishing a set of roles and structures from which individual projects can benefit. Many companies find this difficult to do. Some firms, have been able to establish multiple levels of new roles, from chief knowledge officers to knowledge project managers to knowledge reporters, editors, and knowledge network facilitators.

Successful KM projects benefit from some degree - though not too much-of a knowledge structure. Because knowledge is naturally fluid and closely linked to the people who hold it, its categories and meanings change frequently. This means that too much structuring may be difficult.

Successful knowledge managers realize that knowledge is transferred through multiple channels that reinforce each other. So repositories must be backed by opportunities for face to face interaction.

Successful KM requires an unusual combination of human, technical, and economic skills. These attributes must be present not only in a firm's overall KM effort but also in individual projects. It's often difficult to inculcate them all into a project team, but it can be done.

KM should start with a recognized business problem that relates to knowledge. Customer defections, poorly designed products, losses of key personnel, or a lower "win rate" for service engagements are all business problems that might be traced to poor knowledge management. Attacking these problems, identifying their knowledge component, and using the business value of solving them as justification for knowledge efforts are all good ways to build momentum.

The most important factors in deciding where to start are the importance of the specific knowledge domain to the firm and the feasibility of the project.

Best practices can be a starting point but they must not be viewed as the only form of

knowledge worth collecting and sharing. Moreover, firms should not underestimate the difficulties involved in transferring importing best practices across the organization or from one organization to another. Best practices may be contextual and specific to an organization. Finally, best practices-oriented knowledge management programs deal only with articulated and documented practices. More tacit knowledge about how work is done may not be documented as a best practice. Indeed, broader KM initiatives are usually required to incorporate certain kinds of complex expertise into organizational knowledge.

To begin KM with a focus on organizational learning would be a good idea, but firms rarely do so. Depending on the organization's philosophy, the concepts and approaches involved may include:

- Thinking about the organization as a "system."
- Building and facilitating communities of learning and practice.
- Focusing on issues of personal development and "mastery."
- Creating less hierarchical, more "self-organizing" organizational structures.

Conclusion

Knowledge must always serve the broader aims of the organization. Otherwise it becomes at worst a liability and at best a distraction. We shouldn't learn anything without relating it to practice. A healthy tension between knowledge and action is the key to organizational success.

What makes knowledge valuable to organizations is ultimately the ability to make better the decisions. For this reason, a few organizations are taking a decision-oriented approach to knowledge management. They are attempting to monitor and track "who knows what when" to determine how knowledge is reflected in specific decisions.

Knowledge managers should spend some time assessing their organization's culture before launching KM initiatives. An excessive focus on technology is the most common pitfall in KM. But many firms continue to default to technology because it's easier to buy, implement, and measure.

There is a definite need for KM as a separate organizational function. As Davenport & Prusak mention, all engineers in the organization should be creating and using new product development knowledge. But not all engineers can or will do a good job at writing down what they know. Every person should reflect on life, but not everyone can write stories or poems about their musings. KM will not succeed if there are no workers and managers whose primary jobs involve extracting and editing knowledge from those who have it, facilitating knowledge networks, and setting up and managing the knowledge technology infrastructure. At the same time, KM should not be restricted to a small team at the headquarters. To make a real impact, KM should spread across the organization in reasonable time.

A *Case* in *Point: Javelin* Development *Corporation*

Javelin Development Corporation, a real but disguised engineering and Construction company, developed a plan to make knowledge available across projects in hopes of reducing construction time and costs. The idea was to apply existing design solutions to new situations. The centerpiece of the initiative was an on-line knowledge "warehouse" that engineers could draw from as they developed their designs. A year after implementation began, less than 5 percent of the planned features were in place and support for the initiative seemed to be fading.

We can analyze these disappointing results in terms of knowledge market inefficiencies. Chief among them was the lack of a clear price paid to individuals who shared their knowledge. Having been through a period of layoffs and fearing that more were coming, employees saw their unique knowledge as a source of job security and felt that sharing it would weaken their position. Like many engineering cultures, Javelin's also valued the creation of new knowledge over the re-use of existing designs. Although management supported knowledge sharing in a general way, its actions did not communicate assurance that sharing knowledge was genuinely important and would be rewarded. For instance, employees were expected to learn on their own time, not during office hours, a company norm that implied that acquiring knowledge wasn't "real work." The knowledge initiative had verbal support, but managers did not back it up with a sufficient investment of money and personnel. Some designated knowledge facilitators spent only 10 percent of their time on the project. No one created a mechanism for evaluating knowledge sharing in performance evaluations. As a result of all these signals, trust in the genuineness of corporate commitment to knowledge exchange remained low:

In addition, Javelin's knowledge warehouse was a bust as a marketplace. Potential sellers felt they gained little from adding to the stock of on-line knowledge. Potential buyers did not like the organization of the warehouse content. Project designers had favored a rather loosely structured organization so that knowledge would not be forced into old categories. But the engineers who were the intended users of the system favored a hierarchical system that would make it easy for them to find just the information they needed to solve a specific problem.

With uncertainty and skepticism about the value of offering or acquiring knowledge, lukewarm management support, and a marketplace poorly matched to the habits of potential buyers, the knowledge market at javelin could not function efficiently. The company's serious localness problem was perhaps best exemplified by the experience of a very senior executive who had recently joined the firm. In his previous position at another organization, he had been the primary champion for a very successful knowledge management initiative, yet the organizers of Javelin's knowledge project knew nothing of his interest and expertise. Overall, the company has not yet begun to see the benefits it hoped to get from its knowledge project.

A case in point: Microsoft's Knowledge Map

As we've noted, the knowledge map can refer to documents and structured knowledge, to people, or to both. The most elaborate knowledge maps for people can be quite complex because knowledge structures are complex, knowledge changes over time, subjectivity comes into play, and expertise involves power. One of the best examples of a people oriented knowledge map can be found at Microsoft, where the information systems group decided to map the knowledge of system developers. A 1995 pilot in an application-development group was successful, and full implementation is proceeding. The project, called Skills Planning "und" Development (or "SPUD"), is focused not just on entry-level knowledge but rather on that needed to stay on the leading edge of the industry.

The project objective is to improve the matching of employees to jobs and work teams. Microsoft also believes that once its IT employees have a better idea of what knowledge is required of them, they will be better consumers of educational offerings within and outside of the company. Eventually the project may be extended throughout Microsoft and into products and services for customers.

There are five major stages to the project:
1. Developing a structure of knowledge competency types and levels.
2. Defining the knowledge required for particular jobs.
3. Rating the performance of individual employees in particular jobs by knowledge competencies.
4. Implementing the knowledge competencies in an on-line system.
5. Linking the knowledge model to training programs.

The SPUD project uses a four-type knowledge structure to evaluate employee competency. Entry-level competencies come under the heading of foundation knowledge. Above the foundation level there are local or unique knowledge competencies-advanced skills that apply to a particular job type. A network analyst, for example, might need a fault diagnosis competency for LANs. The next level of knowledge is global, which applies to all employees within a particular function or organization. Every worker in the controller organization, for example, would be knowledgeable in financial analysis; every IT employee would have expertise in technology architectures. The highest level in the knowledge structure comprises the universal competencies for all employees in the company. Examples include knowledge of the overall business the company is in, the products' it sells, and the drivers of the industry.

Within each of the four knowledge competency levels there are two different categories. Explicit knowledge competencies involve expertise in specific tools or methods (for example, Excel or SQL 6.0) and change frequently with the marketplace. Implicit competencies, such as requirements definition, involve more abstract thinking

and reasoning skills. All told there are 137 implicit competencies and 200 explicit ones in the Microsoft knowledge structure. Within each type of knowledge competency there are also four defined skill levels: basic, working, leadership, and expert. Each skill level for each knowledge competency is described in several bullet points that make the level clear and measurable.

Each job in Microsoft IT has to be rated by a manager in terms of the forty to sixty knowledge competencies required to perform it. Workers are also evaluated in terms of the knowledge they have exhibited in their current jobs. The initial rating is built in an iterative fashion by the employee and his or her supervisor; eventually the entire work team participates.

Microsoft is using the employee rating process to build an on-line knowledge map that can be accessed company-wide. A manager building a team for a new project can query the on-line system and ask, "Give me the top five candidates who have leadership skill levels on 80 percent of the knowledge competencies for this job and who are based in Redmond [Microsoft's headquarters location in Washington State]." The system runs on an SQL Server and has a Web front end for easy intranet access around the world.

The system's knowledge types and levels are also linked to specific course opportunities inside and outside Microsoft. Ultimately, the Learning and Communications Resources group hopes to be able to recommend not only specific courses but even specific material or segments within a course that would be aimed at the targeted knowledge level.

Microsoft's knowledge map demonstrates that the company's management values knowledge and supports its exchange. Their commitment of time and money is a symbolic action that has value apart from the actual ability of the map itself. The map simultaneously makes knowledge easier to find and pro motes the idea that corporate knowledge belongs to the corporation as a whole, not to a particular group or individual. Since the success or failure of knowledge work depends so heavily on culture, this benefit of the knowledge map should not be underestimated.