The Knowledge Sharing Challenge

The Sensemaking White Paper

Core members: Mark Stefik (chair), Michelle Baldonado, Daniel Bobrow, Stuart Card, John Everett, Giuliana Lavendel, David Marimont, Paula Newman, Dan Russell, and Steve Smoliar (FXPAL).

Contributors: Eytan Adar, Annette Adler, Todd Cass, Jean-Pierre Chanod (XRCE), Francine Chen, Richard Crouch, Drew Dean, Paul Dourish, Bill Janssen, Ron Kaplan, Dan Larner, Lauri Karttunen (XRCE), Karin Petersen, Maia Pindar, Eric Saund, Hinrich Schuetze, Ralph Sprague, Corbett Williams (DSG/XISS), and Annie Zaenen (XRCE).

Revised by JF: March 7, 2002

Supplementary Materials

Here is a brief guide to the appendices:

- To understand the difference between knowledge technology and other computing infrastructure – see Appendix A.
- For an overview of current (October 1999) knowledge sharing companies and products – see Appendix B.
- For the state of the art in knowledge-sharing technology – see Appendix C.
- If you want a perspective on how corporations create and use knowledge – taking into account the Internet opportunity – see Appendix D.
- If you want concrete examples of solutions and their value propositions – see Appendix E. These value proposition are intended to capture key pain points – identifying key customer needs for solutions.
- If our “model” of knowledge creation and use differs from your favorite model or just want to sample other models – Appendix F.

Appendix A: Knowledge and Information Technologies at Three Levels[1]

The computer industry is not all of a kind. One might say that practically *every* computer system product is in service of information processing or knowledge sharing – and this conclusion might easily be reached by reading the product literature. There is a very wide range of products and functions.

It is the nature of computer information systems that new and higher-level systems are built out of and integrated with components of lower-level systems that are built and often purchased separately. For example, a system for managing online sales for an online bookstore involves not only the web server and a database of order forms and catalog information, but also back-office systems for order processing, inventory, shipping, and billing. An online bookstore may hire out the web site construction from one firm and interface it to legacy systems for inventory and shipping. The bookstore buys its computers and its database software from different companies. The equipment for networking the computers comes from a third company. An application server provides load balancing and state server functions, and simplifies interoperability with backend systems. These different systems work together to support electronic commerce.[2]

Level Number	Competencies & Disciplines	Companies & Products	Industry Pattern
3 **Sharing and interpreting knowledge**	linguistic, community profiling, social sciences, information sciences, cognitive sciences, content analysis	knowledge management companies, Google, IBM document miner, Lotus Notes Applications	Many small players and point solutions now, but the future and major growth belongs to big players.
2 **Exchanging information**	computer sciences, computer engineering, signal processing	Lotus Notes infrastructure, Jini, data miner, Exchange, Verity, AltaVista, Real Audio, Egroups, ERP Applications, workflow, Application Servers, Net Meeting	Highly innovative; new companies. rapid growth
1 **Managing data**	network connectivity, databases, operating systems, printers, disks	NT, Solaris, Oracle, Cisco, 3COM, Hewlett-Packard, Xerox	Established; relatively mature; consolidating

Figure A-1. Three broad levels of systems for information management.

[1] This section draws heavily from a corresponding section from the white paper for the CSL Enterprise Initiative of May 20, 1999 available online at http://amberweb.parc.xerox.com/Get/File-5162/Draft-1999-05-20.doc. In essence, the sensemaking white paper is about level three or knowledge technology, and the enterprise initiative white paper is about level two or information technology.

[2] From an engineering perspective, each level of technology addresses a different class of concerns. For enterprise-scale applications, there are important level-two issues involving system evolvability, maintainability, reliability, survivability, and heterogeneity. Those issues are discussed in the enterprise initiative white paper.

A conventional way to illustrate the functions of computer systems is in terms of layer cake diagrams, where systems at the top level are built out of systems at lower levels. No two information systems are exactly the same and each has its own characteristic components and levels. Nonetheless, there is value at looking broadly at computer information systems in three levels.

The first and bottom level in Figure A-1 corresponds to the basic computing systems on which practically every information system is built. At this level are the computers themselves, the operating systems that carry out operations for the higher level systems, file systems, databases, and the software for exchanging packets of information on the network. The part of the computer industry producing products at this level is the most mature of the three levels and has largely consolidated. Although computers are built by a dozen or so main companies, they are mostly powered by processors built by (or at least compatible with) Intel, Sun, Compaq or IBM. Most of the operating systems are either Windows/NT or a version of UNIX. There are a handful of major providers of databases and network communications gear.

The second level in the figure represents systems built on top of the first level. The dynamic growth of software companies in the 1980s and of so-called Internet companies in the 1990s populate this layer. Like the first layer, the second layer supports "communication." However, communication at this level corresponds to human readable messages and documents. It includes electronic mail, web pages, streaming videos and audio, and shared Lotus notes – built using the computers, local area networks and transport protocols of the first level. Some segments of the industry at level two have consolidated according to the dynamics of the market such as Office Suites and browsers. However, many companies and product groups at this level are new, highly innovative, small to midsize, and poised to grow rapidly.

The third and topmost level in the figure is built on top of systems at the second and first levels. Although there have been systems developed at this level for several years, products tend to be custom built, categories are more elusive, and the industry is highly fragmented. The technologies characteristic of this level are less mature for "industrial strength" applications, involving for example, linguistic technologies for content analysis, summarization, message routing, and data mining.

Some companies offer products at all three levels. For example, IBM offers personal computers, workstations, mainframes, databases, and operating systems at level one. It offers enterprise resource planning systems, Lotus Notes, electronic commerce systems and other systems at level two. It also has products for data mining and document mining at level three. Most of IBM's profits are made at the first level and the company is consolidating its position in higher levels for future growth. There is much competition for market share for established products at the first level. Profit margins are thin where products have become commodities. By adopting a multiple-level strategy, IBM is positioning itself for expansion as the higher levels mature. Other companies like Microsoft, Sun and Hewlett-Packard also operate at multiple levels, and use their positions at the lower levels to move into the upper levels as those markets develop.

One way to look at the three levels has to do with the degree of computer competence in making use of digital information. Viewed in this way, level one is about handling *data* (network communications, files, data bases). Level two is about distributing human-usable *information* (e-mail, web pages, Notes). Level three is about deeper computer interpretation of information or perhaps *knowledge* (content analysis, summarization, and intelligent consolidation, extraction, and routing of information).

As a business enterprise, Xerox has information systems as well as internal projects and products at all three levels. The core of the Xerox business in copiers and printers is squarely at level one. At a smaller scale, Xerox products and initiatives in information systems – such as DocuShare, and various DSG Internet initiatives are at level two. Internal Xerox projects such as Eureka and external products such as linguistic and visualization components offered by Inxight are at level three.

Figure A-2 gives a view of the "enterprise" by the Gartner Group. This view reflects trends that are now shaping competitive enterprises – the move towards computer networks for enabling electronic commerce, streamlining internal operations, managing supply chains with partners, and increased use of outsourcing. Both trends – acquisitions and outsourcing – create demands for interconnecting enterprise information systems as enterprises tie together their distributed groups, outsourcers and customers.

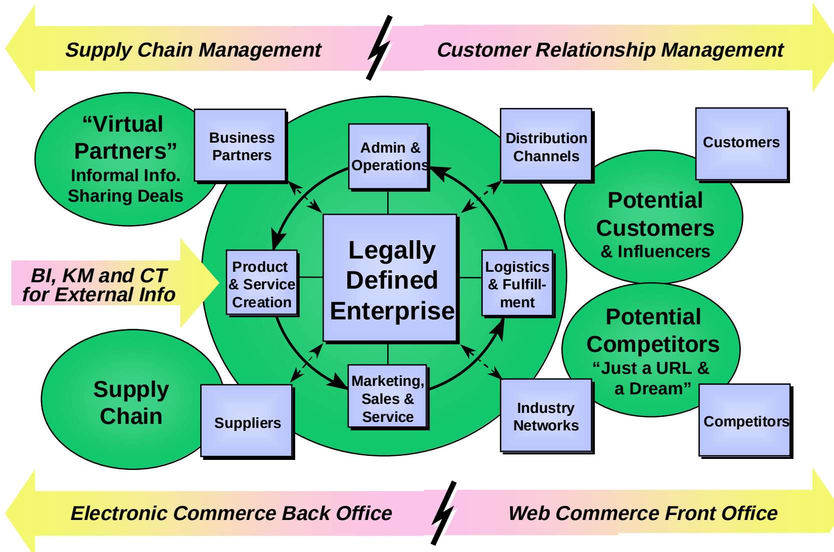

BI = Business Intelligence, KM = Knowledge Management, CT = Collaboration Technologies

GartnerGroup

Figure A-2. Operations in a Generic Virtual Enterprise.

Users of knowledge sharing systems work in organizational contexts -- individuals are members of work groups and larger organizations, affiliated in different organizational relationships as customers and partners. From an enterprise perspective, these divisions may relate in different roles, such as partners, customers, and outsourcers. Restated, the users of knowledge-sharing technology are working in varying relationships of collaboration and competition. They form different subcommunities and operate in overlapping and nested contexts.

Particular individuals need access to information that is part of their own "legally defined" enterprise, or that are part of an enterprise that contracts their work. Because virtual enterprises now contain multiple trust boundaries, the capabilities for sharing, understanding, and using document-centric knowledge need to be mediated by capabilities for properly controlling access and availability of documents. For this reason, level-three technology for knowledge sharing must be integrated with level-two and level-one network infrastructure technology. Put simply, practicality requires that "sharing the knowledge" be combined with "keeping the secrets."

In summary, new technologies at level three are enabling for the knowledge-sharing opportunity. Success at level three will depend on integration technologies at levels one and two in order for the systems to have appropriate properties for scaling, reliability, and security. Level three today is populated by point solutions and many small players. However, the major growth of the future will be dominated by big players.

Appendix B. Knowledge-Sharing Products

This section discusses companies and products related to knowledge management (KM). Here we take a broad view of KM and cover products that manipulate information on which future KM products could be based. For example, one of the categories of products listed below is "business process software and services," which includes Enterprise Resource Planning and Customer Relationship Management, among others. While most KM companies currently do not consider the information collected and produced by these products within their purview, the information is a rich source for future KM products.

The data for our analysis come from a number of different sources. We have compiled a list of over 400 companies, beginning with the companies mentioned in the presentations and consultant's reports in the sensemaking background readings and Knowledge Horizons web sites. News items in the technology trade press added many more companies to this list. Finally, visiting the companies' web sites enabled us to find related companies **--** suppliers, partners, customers, and the like.

The appendix is organized as follows. First we present a list of categories related to our broader view of KM. Next we examine a few interesting companies in more detail (and note that the descriptions of Autonomy, Verity, Intraspect, OpenText, Documentum, and Plumtree are taken largely from their web sites). We conclude with an evaluation of the major trends and propose a framework for thinking about KM products and services.

Categories of products and services

We divide products and services related to KM into five categories: unstructured information management, structured information management, enterprise information portals, and business process software and services. Our goal here is not just to cover products and services that identify themselves with KM, but all those that deal with the information involved in business processes.

Unstructured information management

We use the term "unstructured information management" to refer to document, content, and knowledge management. The unstructured information is almost always text, but speech, images, video, and graphics is sometimes included and will probably be more often as times goes on.

Products in this area include a combination of the following functionality:

- Search and retrieval
- Navigation and browsing
- Ability to access information from heterogeneous sources, including file systems, email, databases, corporate intranets, the web, and repositories like Lotus Notes and Microsoft Exchange.
- Content extraction, topic identification, categorization, summarization, and indexing.
- Organizing information by automatic hyperlinking and creation of taxonomies
- User profiling by tracking what a user reads, accesses, or creates
- Use of profiles to route information to users likely to be interested, to locate expertise, and to facilitate formation of communities interested in related topics or other forms of collaboration
- Personalization of "information workspaces" to facilitate access to relevant information
- Many products in this area emphasize their use of XML as a way to represent structured extracted from unstructured information. Both companies that sell only XML-related products and industry consortia that promote the use of XML by setting and promoting XML standards are playing an increasingly prominent role.

Structured information management

Structured information usually resides in relational databases, and as such is organized in tables with labeled rows and columns. Products and solutions to store, use, and analyze structured information are more mature and in wider use than are those for unstructured information. Relational database technology itself makes it possible to organize and retrieve structured information. Report generation organizes retrieved information into a convenient form. On-line analytical processing is a standard technique to analyze information in relational databases. Decision support systems provide information tailored to a specific business process. Data mining is a more recent development aimed at a more exploratory style of analysis. Another recent development is the increased use of data warehouses, which tie together disparate databases and enable users to treat them as if they were a single database.

Enterprise information portals

Enterprise information portals provide a uniform means of access to all the information relevant to an enterprise, both structured and unstructured, and from a variety of sources. Typically this is via a browser that can be customized for individuals or for groups, including the entire enterprise. According to Forrester Research, the typical enterprise portal includes the following technologies: content-management servers, search tools, automated classification engines, content catalogs, extensible Markup language tools, and directories. As such, EIPs are squarely in the category of unstructured information management – except that they provide access to structured information also.

Business process software and services

Software and services to manage business processes has existed for a number of years.

Enterprise resource planning (ERP), for example, manages and coordinates financial, human resources, and manufacturing processes within an organization; a typical example is coordinating factory operations with purchasing. A recently development is the extension of ERP to manufacturing execution systems, which is used to control and get information about factory operations. Formerly, ERP systems scheduled a customer's order in a manufacturing plant but could not communicate with the manufacturing plant's process control systems (i.e. MES).

The integration of MES with ERP is an example of enterprise application integration: integrating legacy enterprise applications and databases so they can interoperate.

Customer relationship management includes “front office” processes like sales and marketing. Putting the necessary information online enables building profiles of customers that in turn enables more narrowly targeted, and presumably more effective, targeting. Related products include software for managing problem resolution and call centers, and networks for product support.

Business intelligence (including customer profiling) is a marketing term that refers to tools businesses can use to analyze massive amounts of information stored on databases. The tools aim to help firms improve marketing efforts or boost sales.

Businesses use supply-chain management (SCM) software to track goods and services from suppliers through manufacturing to distributors and sale. SCM involves managing processes that take place within a business, as well as processes that involve outside suppliers, manufacturers, distributors, and sales channels. The software is often customized to meet the needs of different types of businesses.

Business-to-business (B2B) communication for more general purposes is an area of increasing activity as well. This includes products and sites for B2B electronic commerce; a particularly interesting development is vertical industry portals, which provide access to information and markets for a specificy industry.

Others: enterprise asset management, enterprise skills management, intellectual capital asset management, digital marketing, -- the growing list includes most business processes that are well-structured enough to benefit from being put online.

Interesting companies

IBM

Transition

IBM is negotiating gracefully a transition from hardware to software and services. Global Services, where most of IBM's Knowledge Management activities reside -- they are part of "Professional Services" -- is a corporate star with 1998 revenues of $28 billion, and a pattern of 20% annual growth.

Services foremost

Knowledge Management belongs to Services, in IBM parlance, and the phrase KM is rarely used officially by executives. Larry Prusack, the designated spokesman in KM matters -- with an information resources rather than MIS background -- is officially a "managing principal" in the IBM Consulting Group, and is also known as Executive Director of IBM's Institute for Knowledge Management Prusack promulgates a people-based, humanistic, sensemaking view of KM activities.

However, other important components to IBM's KM effort are contributed by technology-based Lotus Consulting and Database Management Group.

ERP angle

It is significant that several of IBM's professed goals pertain to ERP, and are complementary to Knowledge Management, such as gaining "insight and foresight from both the enormous underutilized stores of data that organizations already possess, and the sea of information that pervasive computing will generate…mining databases in retail or insurance for patterns and insights…."

Lotus Notes

IBM has the good fortune, or rather had the foresight, of having secured the leading software in which KM applications are deployed. This is an enormous advantage which IBM has enjoyed in a rising tideswell since 1994.

IBM labels Lotus/Domino "software and collaboration" tools, whose new installations totaled more than 14 million in 1998. In 1994 only 1.35 million seats were sold: now installed seats total almost 35 million, worldwide.

Database technologies

The IBM Intelligent Miner for Text offers web crawler, concept extractor, summarizer, clusterer and categorizer functions "providing a context that enables the creation of tacit knowledge or insights." Intelligent Miner and the companion Visual Warehouse were recently augmented by the acquisition of KnowledgeX, an application that diagrams information stored in various sources, including the internet, displays retrieved documents as icons end graphically links related contents.

KM and the web

As described above, IBM's strategy relies on the strength of its Global Services, where KM consulting is practiced, on the prevalence of Lotus/Domino, and on the advanced technologies of the Database Management Group -- in this order.

A relatively new development concerns Knowledge Portals., an ultimate refinement of the approach which began with Intranet Web Portals, evolved into Information Portals, and is now "the fundamental building block of a Knowledge Management Infrastructure," according to the Technology Director of IBM's Knowledge Management Institute.

Knowledge Portals

Knowledge Portals have two interfaces:

Knowledge Producer for "mapping" while gathering, analyzing, adding value, and sharing information;

Knowledge Consumer that facilitates communication and dissemination throughout the enterprise.

It is interesting that other key IBM features such as the Knowledge Map (view from the top of the firm's knowledge assets), Lotus Notes and Teamroom are being integrated with the Knowledge Portal repository, adding collaboration and expertise elements to the mix.

Microsoft

Microsoft to date has not fulfilled its promise to spring from groupware to the KM market ; it was to be based on a close integration of Outlook and the Office Suite. It remains to be seen whether the new management called to formulate Microsoft's Internet strategy -- i.e. Rick Belluzzo, formerly the no. 2 man at Hewlett-Packard and Silicon Graphic's CEO - will be able to provide the necessary functionality and integration on the Web within the available time window.

To date, Microsoft's position, as defined by analysts, has been "bottom up," insisting that its tightly woven software can support KM solutions: this is a software only strategy, with no mention of services. It could be a monumental weakness, but it is the Microsoft way. It is also possible, in all fairness, that such haphazard approach is due to restraint imposed by the current anti-trust proceedings against the company.

The big five consulting companies

Andersen Consulting began as a technology-based system integrator. It has successfully built on this background with its Knowledge Exchange project based on Lotus Notes, but has experienced problems of an organizational and cultural nature (technology first!). By installing Notes internally all over the firm, Andersen gained invaluable experience in matching content with KM tools.

Internal KM is based on the four concept s of Integration, Socialization, Utilization, Automation, all flowing into a massive Knowledge Database., with security layer. Customer site practice is similarly mapped as Strategy, People, Process, Technology.

Andersen has over 50,000 employees worldwide, and spends heavily on research, especially KM oriented. As experienced by other management consulting firms, its KM practice is growing, and routinely embedded in new consulting agreements.

Ernst & Young is the most adventurous of the former Big Five as far as KM. is concerned, and leans heavily on its "thought leadership" process for its "Knowledge Based Business.". It is an early adopter, and its Centers for Business Knowledge (CBKs), based on early internal experiences, are well established internationally. Internal experiences, people-based, are fundamental to the firm's approach.; "community of practice," "community of interest, "best practices" are familiar terms. It is interesting that a proposed merger with KPMG failed last year because of cultural differences. (Lots of sensemaking!)

ERP is a growing revenue source, as is KM: health services are a strong presence among customers. Over 70,000 employees worldwide. Has an impressive list of KM customers, from American Express to Shell, and has (or had) an agreement with Xerox.

KPMG (Peat Marwick), a relative latecomer to the discipline, follows the principles of Value-Based KM, and originates from the firm's large IT consulting business; this is typical of the large consulting organization from the financial arena (the former Big Five),where IT is the key driver.

There are three KM core operations in the U.S., U.K. and Netherlands -- the U.S. operation being the most interdisciplinary, including e-commerce, etc. Data warehousing is a specialty, as for all the financial consulting firms. While only about 100 consultants are dedicated to KM, another 1000 can b called upon, and every business transaction has at least a 10% KM content by design. Total staff numbers over 85,00, in 155 countries.

PricewaterhouseCoopers (PWC) defines KM as the discipline of leveraging intangible assets to generate business value. Its KM operation shares many of the characteristics mentioned above for other hard line consulting firms, moving at a slow pace from technology to a broader view of KM.

Its recently launched global effort to provide "consulting, implementation, and content" for customers focuses on ERP, and has two main thrusts, as shown by recent significant moves. It acquired the Ventrix knowledge support system, which provides a digital infrastructure to capture and distribute knowledge through the enterprise -- and will also help support SAP/R3, PeopleSoft, Oracle applications. PWC has also opened a Global Knowledge Center near Chicago, dedicated to data warehousing and mining; a similar facility is next to be opened in Europe. Data warehousing knowledge-based solutions seem to be a focus.

PWC is the largest professional services organization in the world, with over 140,000 people distributed in 152 countries. KM is part of the Management Consulting practice, which is 25,000 strong; among these are 1,000 data warehousing consultants.

Autonomy

Autonomy bills itself as a provider of KM technology products that “automate the handling of vast amounts of unstructured information on the Internet and across the enterprise.” These products categorize, hyperlink, personalize and search large volumes of information (using natural language queries when appropriate). Searching is facilitated by aggregating information from a variety of sources (including documents, email, news feeds, and web sites) into directories for easy navigation. The underlying technology is statistical pattern recognition.

Autonomy claims that its technology can “derive meaning in a piece of text.” As a result, the technology can also be used to profile users by analyzing the ideas in the documents they read or produce. Such profiles can be used “to deliver personalized information, create communities of interest and, in knowledge management applications, identify colleagues with useful expertise.”

Verity

Verity provides “knowledge retrieval solutions” that “turn unstructured text-intensive information into usable and shareable knowledge.” According to their web site, their products

- Index, search and retrieve information on Web and file servers distributed across the enterprise and stored in virtually any format.
- Include powerful search and navigation facilities including document clustering, automatic summarization and query by example.
- Enable indexing and retrieval of information from relational databases and popular repositories such as Lotus Notes and Microsoft Exchange.
- Actively monitor information sources across the Internet and intranet and automatically filter and deliver highly personalized information to individuals and groups, based on their content and delivery preferences.
- Organize information stored on intranets, extranets, Internet Web sites, file servers, and networked repositories into hierarchical directories of categories.
- Classify documents automatically into one, or more such directories based on business classification rules, file system hierarchies, document metadata, and URLs.

Verity emphasizes the importance of unstructured information to business:

> While relational databases have simplified access to structured information, more than 80% of corporate knowledge is unstructured in form — residing in file systems and Web servers — not in databases. Today, business knowledge is found in many forms, from bids, proposals, and resumes, to project plans and patents — in a wide variety of formats, including word processing documents, spreadsheets, presentations, e-mail messages, SGML documents, news feeds, Adobe Acrobat files, and Web pages.

Intraspect

INTRASPECT provides Collaborative Knowledge Management software which enables people to make faster, better decisions by connecting them with a Group Memory of relevant information and company expertise. Intraspect enables users to capture and share knowledge with customers, suppliers, and partners, while working on line.

Intraspect's flagship product Intraspect Knowledge Server is a collaborative environment for collecting, finding, and re-using information. We deliver a single, personalized interface to any digital content such as reports from enterprise applications, documents, and electronic mail organized into secure, topic-specific folders. Intraspect encourages and facilitates participation and captures how people use information to make decisions.

Knowledge-enabled solutions are the path to higher productivity and improving organizational learning. Key to delivering knowledge-enabled solutions is the concept of the corporate "group memory." Intraspect Knowledge Server enables:

- the capture and storage of information from across your enterprise, including external sources such as web documents, and a variety of internal sources such as desktop applications, enterprise reports, and legacy databases.

- the organization of information into its working context. Users e-mail directly into public and private workspaces. Context-specific attributes of the information are automatically captured and displayed on retrieval.
- collaboration and knowledge sharing. Inter-company participation is made easy through close integration with any e-mail client or browser. Comments and discussions are user controlled, supporting the emergence of self-supporting communities of practice.
- sophisticated search and reuse of captured knowledge. Intraspect's unique combination of powerful, sophisticated search, database, and contextual mapping technologies gives users fast, relevant results. Users can find not only documents, but teams, projects and expertise.
- the use of Knowledge Agents. Users can subscribe to documents, URLs, discussions, topics, or specific locations in the group memory and be notified via e-mail or a web page when something changes, or something new has been added. In addition, users can create targeted, persistent searches and receive notification whenever a new "match" is detected in the group memory.

Knowledge-enabled solutions make knowledge work more productive by improving the way information is acquired, distributed, shared, and reused in an organization. People should be able to easily find colleagues, create a shared workspace, contribute to the project, and collaboratively discuss the evolving work. They should be able to make the results of their work available online, without spending extra effort to prepare it for publication. They should be able to easily discover and reuse their colleagues' knowledge, without relying on others to anticipate their information needs and keep them informed. IKS is designed to meet *all* of these needs, supporting the work practice of knowledge workers while building a corporate *group memory*. The group memory is the repository of the information used in an organization, and meta-data about how it is used to solve business problems.

Intraspect products provide the following knowledge management services:

- Information collection: enables end-users to import any object that is visible in a client web browser; upload any files visible to the client file system; send in any e-mail message, with attachments, from any SMTP-compliant e-mail application
- Information Sharing: provides on-line workspaces for collecting and organizing information, with versioning and access control. Integrates personal and public spaces into enterprise object repository to support individual usage and simplify wide-area collaboration. Enables integrated commenting to allow collaborative review of content, or for taking personal notes.
- Information Organization: provides hierarchical organization of information, with single copies of objects in multiple folders to support multiple contexts of information (re)use. Identity of objects is independent of their location in hierarchy, to support dynamic reorganization.
- Integrated Discussion: enables discussions to be held within the workspaces where documents are shared. Threads and parses messages automatically for attachments. Routes existing discussion lists to the IKS. Compatible with all SMTP mail clients, and independent of mail server, to lower barriers to participation in on-line work.
- Discovery and search: indexes automatically for full text search of all collected documents including binary formats, e-mail messages and remote web pages. Integrates attribute search with full-text search
- Universal Subscription: subscribes to any object, including workspaces, individual documents, and persistent searches. Issues notification of changes to subscribed objects delivered asynchronously to end-users.

Intraspect's other products are

- <u>Collaborative Knowledge Portal:</u> Intraspect provides customers the ability to build and customize how they "view" information based on what is important to their business. A Collaborative Knowledge Portal enables users to access consistent information, combine it with other data, analyze and share relevant results, and collaborate around it in order to increase productivity and perform day to day jobs more effectively. The Collaborative Knowledge Portal will be the window for sharing, collaborating, and re-using information.

- Collaborative Business Intelligence (CBI): CBI delivers the results of Business Intelligence analysis and enterprise reports in the context of a Collaborative Knowledge Management environment, where they are combined with unstructured information, distributed to interested parties, and acted upon by decision makers. The result is extending value and use of existing enterprise applications and the ability to make better, faster decisions.

OpenText

Open Text's provides collaborative knowledge management solutions based on Livelink, their flagship product. Livelink uses the corporate intranet to provide a collaborative environment, enabling project teams to form, and to capture, share, and reuse corporate knowledge. Livelink provides enterprise document management, information retrieval, workflow and project collaboration, and enterprise group scheduling capabilities. Livelink is designed to handle complex business processes over the Web.

Three distinct Workspaces – personal, enterprise, and project-- provide users with multiple views of the enterprise and its activities for improved knowledge access, sharing, and navigation. From each Workspace, users have access to Livelink's powerful services. The UI to each workspace is browser-based.

Livelink is based on a set of knowledge management services:

- **Knowledge Library:** Livelink's Knowledge Library provides document management services for accessing, storing and managing millions of documents. It applies document management functionality to a secure, central repository where you can store any type of object including documents, folders, search queries, channels, URLs, workflow maps, discussions, tasks, reports, etc. in an organized, hierarchical structure. The Knowledge Library can span multiple locations, allowing users across the globe to access, store, browse, search and manage documents from a single secure Web address.
- **Virtual Team Collaboration:**. Livelink's Virtual Team Collaboration services transform your intranet into a highly collaborative environment for sharing knowledge and ideas. Users and workgroups can collaborate on any type of project, whether ongoing processes such as quality assurance or short term events such as organizing a press conference.
- **Information Retrieval:** Livelink's Information Retrieval services provide a comprehensive knowledge retrieval solution for finding, organizing and navigating through large quantities of corporate data residing in a variety of distributed and diverse network environments and data repositories. Search for information in documents, files and objects uses sophisticated search tools such as intelligent relevance ranking, thesaurus, proximity and Boolean searches.
- **Business Process Automation:** The ability to capture, route and track information makes it possible to manage and improve business processes.

Documentum

The Documentum product family includes the Documentum 4i Web application environment and the Documentum Innovation Application Series.

Documentum 4i dynamically manage all facets of content creation and delivery for both corporate and e-business portals. Utilizing developments in XML content and process integration, Documentum 4i provides aplatform for development and deployment of globally distributed portals.

The Documentum Innovation Application Series creates a corporate portal to business know-how and business-critical content, so you can apply knowledge for innovation of new products and processes securely and reliably. Built upon scalable server and repository technologies, these Web applications dynamically serve up and personalize Web content within and between organizations. They automatically connect knowledge workers to the right people and processes, at the right time, in the right format and business context. They provide the capability to codify domain expertise into the Web applications to meet specific challenges of applying knowledge in vertical industries

Documentum 4i is a Web application platform for corporate portals that dynamically aggregates business knowledge, processes and content within and between organizations. It delivers capabilities for building, deploying and administering corporate portals. Documentum 4i captures, configures, and manages business knowledge -- the understanding of the relationships between people, information, and processes that interact with each other to accomplish a business objective. It provides (1) the ability to manage content and to automate the processes with which the content is associated; (2) the ability to configure business knowledge; (3) a robust Web architecture that scales to the demands of the Internet; (4) tools for rapid application development and deployment. Documentum 4i includes the following key capabilities:

A **Web architecture** that integrates with industry-standard Web technologies and scales to meet the demands of the Internet. This helps to deliver fast-payback applications by leveraging existing developer know-how and ensuring rapid global deployment.

Capabilities for **aggregating relevant content** with the appropriate business functions for creating personalized views that improve productivity by keeping users in sync with changing content and processes.

Functionality that **automates the business processes** associated with routing, approval, distribution, notification and auditing. By automating the tasks associated with using global content, Documentum 4i accelerates time-to-result.

Web content management services that manage content storage, versioning, publishing, security, tracking and archiving. By automating Web content management, Documentum 4i can ensure content integrity by eliminating the errors associated with manual content management tasks.

Tools for **developing and configuring applications** that reflect dynamic roles, content, and processes. By replacing many common programming tasks with simple configuration, Documentum 4i lowers the total cost of application ownership.

Plumtree

Plumtree Software provides software for corporate portals. In one simple, personalized Web page, employees, partners, and customers can see all the information relevant to their work, and all the productivity tools and e-commerce services that they use every day. In one place, employees, partners, and customers can review product and market news, analyze key performance metrics, launch productivity tools, and complete e-commerce transactions. Plumtree claims that this makes corporate portal users more productive,and gives management a simple, consistent way to keep everyone informed about new products and strategies.

Plumtree Software specializes in several applications: field support for professional services, high technology and other industries; product research and exploration for life sciences, chemicals, and oil and gas; customer service for businesses with complex products; and competitive intelligence for companies in fast-moving industries.

To bring together everything important to users' work in one Web page, the Plumtree Corporate Portal scans databases, Groupware systems, Web pages, documents on file servers, and other types of content.

Hosting dynamic applications such as on-line reports, calendars, e-mail, and commerce services creates a one-stop shop for users to go about almost all of their daily tasks: analyzing customer trends, checking schedules, viewing revenue- or project-related performance metrics, and buying or selling products. Combining all the information relevant to users' work with the ability to act on that information enables organizations to get more done.

The Plumtree Corporate Portal has the following features:

- Personalized: enables users to personalize their portal pages to feature only the content, applications, and services relevant to their work. Mobile users can choose to receive updates via electronic mail or other media in lieu of the Web.
- Comprehensive: provides a complete view of both the content and application services available to an organization. The system is extensible, so it can integrate access to new content and services to meet customers' specific requirements.
- Organized: assembles content in browseable logical categories regardless of its physical location, providing an at-a-glance view of important information without requiring users to issue searches.
- Automated and Distributed: Portals that rely exclusively on users to submit links die of neglect. Portals that categorize information without user involvement focus efforts on technology rather than the ways in which people think of their business. The Plumtree Portal strikes a balance, automating repetitive tasks, but also

involving different business experts to approve portal updates. As a result, keeping the portal up-to-date becomes a byproduct of standard business practices

- Secure: secures access to each embedded application and to entire categories of information, as well as to individual document links. As a result, administrators can use the portal to broker access to sensitive content and tools.
- Scalable: scalable and fault-tolerant, with support for any number of Web servers to support an increasing number of users, and any number of job processing servers to support large numbers of documents. If any Web or job processing server fails, users experience no interruption in service. As a result, large organizations have the confidence to deploy the portal to a large number of users on an intranet, extranet, or on the Internet
- Location-transparent: Content and tools are hosted by their native systems, reducing maintenance costs and simplifying deployment

Major trends

Standalone and client/server applications move to the internet

In both enterprise applications and other PC software, the software model is moving from standalone (e.g. office productivity suites) and client/server (business process) to the internet.

This transition has received much attention lately in business process software, where IBM, Siebel, and Oracle have been leaders in this trend. Companies like SAP and PeopleSoft have been criticized for being slow to join this trend, but now even they are changing their strategies as their sales suffer.

The transition in office productivity software is proceeding more slowly but very recently has begun to pick up speed. Sun’s recent purchase of StarOffice, companies like myDesktop.com, myWebOS.com, myInternetDeskTop.com have all been in the news lately. Even Microsoft recently announced plans to make Office available via an application server platform.

Most companies and products in document, content, and knowledge management seem to have already successfully negotiated this transition. In some cases, the products arose so recently that they never existed in another other form.

One way to look at this trend is that software is becoming a service rather than a product. Instead of buying multiple copies of a software product, users will lease the services the software provides from an application service provider.

Personalized products and services

Both in the business and consumer markets, the marketing of many new products and services emphasizes the advantages of personalizing the product or service to a particular user or set of users. The proliferation of web sites named “myX.com,” where X equals SAP, DeskTop, etc. is an indication of this trend. While the personalizing is often done manually through more advanced user interfaces, an increasing trend is for software to “learn” a user’s preferences by observing the user’s behavior. Products for portals and for document, content, and knowledge management are leaders in this area.

This is part of a major trend sweeping not just the software business but business in general: in more and more industries, a customer is able to order "made-to-order" products. Dell has been a leader here; Dell's internet site enables a customer to personalize a computer by specifying a set of options from a long list. Then Dell manufactures that computer and ships it to the customer. Precise control of and rapid communication along the supply chain, enabled by the internet and Dell's intranet and extranet, makes this personalization possible. Now, supply-chain management software is making possible this level of personalization in other industries, like auto manufacturing (e.g. Ford, GM, and Toyota).

Convergence of document, content, and knowledge management

Products and services for document, content, and knowledge management are looking more and more alike.

Business process software extends its scope

More business processes are being handled with enterprise applications (provide list).

Enterprise application integration

An important trend in business process software is enterprise application integration (EAI). Until now, a company's enterprise applications may have come from different vendors, with different applications communicating and storing data in incompatible formats. Now, many companies are offering EAI software and services to integrate these legacy applications both so they can interoperate and so the underlying data can be viewed and analyzed with a single application.

Business-to-business interaction

Just as the business processes internal to a firm are going online, so are the business processes between firms. Electronic procurement and supply-chain management are two examples of this trend; vertical industry portals that coordinate activity among a group of firms are another. Many industries are developing "dialects" of XML to facilitate interfirm communication; Microsoft's Biztalk is an a framework for creating such dialects.

Outsourcing

In 1998, companies outsourced 15% of all manufacturing. In 2000, they will outsource more than 40% (Hambrecht & Quist quoted in Business Week of 4 October 1999).

Business and competitive intelligence

The recent spate of new offerings in business and competitive intelligence is a significant development because BI and CI are quintessential knowledge management processes. Central to both BI and CI is collecting (from heterogeneous sources), organizing, analyzing, and using knowledge about one's customers and competitors. The development of this market may provide a useful paradigm for other forms of KM.

Portals

Information portals provide uniform access to all of an organization's information. The most common kind is an enterprise information portal (EIP), but portals for vertical industries exist as well and serve to coordinate business activity among the individual enterprises in the industry.

As enterprise applications cover more and more business processes, as enterprise application integration ties the various applications together, and as these applications move the the internet, it has become possible to design user interfaces to access all of an enterprise's information. A number of companies now provide software and services to create EIPs, and more and more companies are using them.

Portals also provide access to an organization's unstructured information; Plumtree's portal products, for example, have many of the features associated with products for document, content, and knowledge management.

Bandwidth and the burden of abstraction

When high bandwidth is expensive or impossible to obtain, the sender of a message economizes on bandwidth by either compressing or abstracting the message. Now that the cost of high bandwidth is plummeting rapidly, economizing on bandwidth is less likely, so more messages are transmitted without an initial stage of abstraction. For example, when the telegraph was the only means of rapid long-distance communication, a reporter covering a news event would transmit a very concise, high-level version of the story. Now it is common for a video of the news event to be transmitted instead.

Because messages must be interpreted to be used, the burden of interpretation or abstraction has shifted from the sender to the recipient. Moreover, the number and size of messages an individual recipient has increased greatly, also because bandwidth has become so cheap. Thus it has become necessary to automate the process of abstraction and interpretation (whereas in the expensive-bandwidth scenario, a human could perform the abstraction because no one was sending a large volume of messages). This drives the market for document, content, and knowledge management products.

Raising the level of structure in information and process

A key development in many recent products is their ability to manipulate unstructured information and to enable the performance of unstructured processes.

Until recently, much of the information involved in business processes was highly structured and stored in relational databases. Often these data were collected to support highly routine or structured processes like purchasing and payroll. The techniques for analyzing such data consisted of querying, report generation and online analytical processing. A more recent addition to this list is data mining, which can be viewed as supporting an unstructured process, since the typical data-mining user doesn't know in advance what kinds of structure he or she may discover.

Workflow is an example of a structured process. Workflow modeling is an example of raising the level of structure in a process. Examples from knowledge management include capturing best practices and converting tacit to explicit knowledge, both of which involve raising the level of structure in a process. Raising the level of structure in a process is equivalent to the conversion of process knowledge from tacit to explicit form.

The typical unstructured information that recent products manipulate is text. These products go beyond conventional document management, which made no attempt to extract the content of the documents automatically. Instead, the operations on text supported by the newer products both extract and manipulate the text in ways that can be described as raising the level of its structure – converting what's "tacit" in the text (absent a human reader) to explicit form, for example by part-of-speech tagging or parsing. At higher levels, content extracted from a collection of texts can be used to raise the level of structure in the collection, for example by hyperlinking or creating a taxonomy.

Other forms of unstructured information include graphics, images, video, audio, and speech. Speech recognition works quite well in highly constrained situations, and the output is text, and from there raising the level of structure proceeds along a familiar path. Raising the level of other forms of unstructured information is a difficult research problem, and products for doing so have had little impact on the market.

Unstructured numerical data is another form of unstructured information; here statistical techniques are the primary means of raising the level of structure. Many problems in text and web analysis reduce to this form, and statistical techniques are widely used in those areas.

Enterprise information portals, like data mining, support unstructured processes. Users of portals may know what information they seek but the search process itself may be quite difficult to make explicit. Portals provide access to both structured and unstructured information.

As mentioned above, many interfirm processes like procurement and supply-chain management are now being put online. This requires that the structure of these processes be made explicit. The same thing happened during the development of enterprise resource planning software, when the processes involved in financing, payroll, and others had to be made explicit. Making a process explicit is equivalent to raising the level of its structure.

XML has a key role to play in representing structure in information used in both structured and unstructured processes. As a document format, it represents the explicit structure of the document's content. As a general data format, it represents the highly structured information used in interfirm processes. In the future, raising the level of structure in information may take the form of converting from one form of XML to a more highly detailed one.

We note that hypertext systems have dealt with the problem of add structure to unstructured collections of documents for decades. While hypertext systems do not ordinarily provide the facilities for interpreting individual documents, they have produced many techniques for organizing collections of documents and making them accessible via browsing and navigation. More recently, architectures for open hypermedia systems have been proposed that have much in common with architectures for portals.

One way to interpret how the use of structured information has evolved in business process software is the following: putting business processes online required that the level of structure in both process and information be raised. Once the structured information in relational databases became available, less structured processes involved in exploring the information, like data mining, became possible. Moreover, as enterprise application integration proceeds, uniform access to all the structured information on which business processes are based become possible, for example via data warehousing. This in turn enables raising the level of structure in higher-level business processes, since the structured information on which they are based is available, as well as making the structured information available for unstructured processes of broader scope.

This history has obvious relevance for unstructured information like text, graphics, audio, images, and video. One possible future is that structured processes will lead to a particular way to structure the information, in turn leading to the use of the structured information in unstructured processes. Already the term “knowledge warehouse” is being used.

Appendix C. Knowledge-Sharing Technology

In the market-to-technology model (see Figure 1), value propositions map onto business models on the one side and onto technologies on the other side. In this section, we take a detailed look at the slate of technologies that contributes to knowledge-sharing solutions. In particular, we focus on describing the state of the art for a number of the Level 3 contributing technologies, which we term *knowledge-sharing technologies*. These technologies, the immediate building blocks for the creation of solutions, are the subject of active research and development. Thus, in this section we will consider not only the current state of each technology, but also (where known) the centers of excellence for advancing the technology, the challenges that face researchers and developers in improving the state of the art, and the metrics that we can use to evaluate the technology as it progresses.

Think of this survey as a rapid tour through a sampling of knowledge-sharing technologies. Our goal is not to set forth an exhaustive accounting of technologies, but rather to ground the discussion of solutions and value propositions in the reality of the current technical state of the art. We structure our discussion by grouping knowledge-sharing technologies according to the domains of inquiry to which they belong. While this set of domains is necessarily incomplete, it provides us with some established landmarks in the terrain of technologies. In particular, we discuss knowledge-sharing technologies in the context of the following domains (note that all of these domains encompass far more technologies than we cover here): artificial intelligence, digital libraries, document management, human-computer interaction, information retrieval, multimedia, natural language processing, security, and the Web.

Knowledge-Sharing Technologies from Artificial Intelligence

The domain of artificial intelligence has a natural role to play in the arena of knowledge sharing. Knowledge representation technologies are used to build knowledge bases, which can in turn be used for improving precision and recall in information retrieval. Agent-based technologies have served as the basis for a number of knowledge sharing applications, ranging from information finding to information filtering to e-commerce negotiations. Finally, data mining—at the intersection of artificial intelligence, databases, and statistics—has recently emerged as a technology mature enough to be used in commercial settings. Many large challenges remain, including making knowledge representation and agent technologies work for useful real-world domains, investigating how to prevent overfitting in data mining, and engineering computationally efficient and scalable algorithms that will work on large focused document collections. While the question of whether or not these technologies will ultimately scale to domains that are very general and document collections that are extremely large (e.g., the Web) is likely to remain unanswered for some time, researchers believe that artificial intelligence techniques are likely to succeed for mission-critical, relatively small document collections.

Example: Imagine that a high school physics student wishes to learn more about the concept of electrical resistance. A search on the Web for "resistance" would return not only documents about this topic, but also documents with sentences like "*Gauge resistance to the proposal at the next meeting.*" Artificial intelligence technologies promise to ensure that only documents that truly pertain to the topic at hand would be returned.

Knowledge representation

Research in the area of explicit knowledge representation is focused on developing large-scale ontologies that organize knowledge of things-in-the-world (such as fuses, meetings, airplanes, and cities), inferential mechanisms for efficiently drawing conclusions from this knowledge, and tools for entering new knowledge without introducing inconsistencies. DARPA is funding a major effort in this area, the High Performance Knowledge Base project. Two integration teams, one led by SAIC and the other by Teknowledge, annually field systems designed to solve the challenge problems in "friendly competition" with each other. Technology development participants include Stanford, MIT, Carnegie Mellon, Northwestern, University of Massachusetts, George Mason, University of Edinburgh, SRI, ISI, the Kestrel Institute, and TextWise Inc. The HPKB program held its first annual evaluation in June 1998, at which time extensive tests based on the challenge problems were conducted. The systems both attained, on average, from 60% to 80% correct response rates. In general, centers of excellence in this area include Stanford, Northwestern, the University of Texas at Austin, and CYCORP. Current active topics of research include ontology reuse, knowledge base integration, context modeling, and knowledge base construction.

Technologies based on explicit knowledge representation are just beginning to appear in commercial products. The Web search engine AskJeeves makes use of simple knowledge representations of natural language queries. Lexeme is marketing semantic indexing technology now in use at Medstract.org, where it automatically parses incoming abstracts into a database of functional characteristics of genes and proteins. Excalibur Technologies' RetrievalWare system employs a semantic network containing 400,000 word meanings and 1.6 million word associations to aid retrieval of documents.

Over the next decade, it is likely that we will see an increasing number of knowledge-based systems. In contrast to expert systems, which attempt to provide a correct answer for a given situation, these new systems will provide plausible or approximate responses within a more tightly coupled human-computer interaction.

Agents

The technologies for agents, computational entities that perform work on a user's behalf, are under development at both research institutions and commercial enterprises. Agent technologies can be divided into technologies for multi-agent systems and technologies for autonomous interface/information agents. Advances in multi-agent systems have been in agent coordination, communication, and reasoning. Advances in autonomous interface/information agents have been in monitoring, user modeling, and agent mobility. For both classes of agents, the predominant substrate is rule-based. Stanford, MIT, the University of Michigan, the University of Washington, and the University of Maryland are just a few of the universities that are at the forefront of agent technology. On the commercial side, General Magic was one of the more recent companies to extend market awareness on the topic with its deployment of mobile agents. Interface agents from human-computer interaction have also emerged; notably, the Microsoft Bob interface could be understood as an agent-based system. Another example of an interface agent is Letizia (from the MIT Media Lab), which tracks Web browsing history and performs automatic exploration of links from the current Web page in the browser in order to suggest additional links to the user. Interface agents that use natural language to communicate with users (e.g. several interface agents on the market for e-commerce) could undoubtedly be improved with natural language technologies (to be discussed later), including full parsing and generation.

Data mining

In the last decade, data mining—a domain concerned with identifying patterns and anomalies in data sets—has established itself as an important domain. Data mining technologies typically employ statistical methods and are developed by researchers at the intersection of artificial intelligence, databases, and statistics. Techniques for classification and clustering, trend and deviation analysis, dependency derivation, and visualization are advancing rapidly, but improvements are still necessary in predictive power, accuracy, and scalability. In addition, usability remains a large barrier to the widespread use of data mining tools. Most systems currently require a significant degree of expertise on the part of the end user. Nevertheless, some data mining technologies have already appeared in the commercial arena. IBM's data mining tools are well known in this respect. Overall, centers of excellence for these technologies include IBM Almaden, AT & T Labs, and Microsoft Research.

Image Understanding

Conventions governing document appearance have evolved over centuries to visually convey meaning and purpose to people. While new genres of documents are emerging in electronic media, documents on paper continue to play the major role in knowledge sharing. Meaning is contained not only in textual content, but also in visual appearance. Visual structure conveys information through its elements – font, layout, graphics, highlights, etc. – reflecting such implicit information such as genre, logical structure, emphasis, and formality. Computational methods that extract and exploit this visual structure are primary components to automatic and semi-automatic tools for acquiring and assimilating knowledge. Computer vision technology underlies dual functions in knowledge sharing: (1) the conversion between paper and electronic media through Optical Character Recognition and its extensions, and, (2) enablement of interactive and automatic batch tools that exploit perceptual aspects of visual structure for selecting, organizing, interpreting, and sharing document material. Of these, the former is by far the more mature technology.

OCR systems for office documents are by now a commodity whose vendors compete on price, speed, and accuracy, and features. Accuracy declines rapidly with degradation in image quality, making paper-to-text conversion sometimes a hit-or-miss undertaking. Steady but slow progress continues in incorporating aspects of document layout in OCR output. This includes formal document structure of articles and reports, table detection, and exportation of tabular data to spreadsheets. Batch forms conversion systems are gaining widespread acceptance but often involve a built-in human verification/correction step. In domains with highly constrained variations in document layout, font, and textual content, technologies that exploit these strong prior constraints promise to improve accuracy rates dramatically.

Systems that exploit perceptually evocative aspects of the visual structure of documents present an important opportunity for leading-edge research. Genre classification plays a significant role in the organization and evaluation of information and information sources. Fusion and creative mixing of document material from diverse sources requires tools that can segment salient text, annotation, and graphic elements from one another to make them available for independent evaluation and re-purposing in new documents, possibly across different media. Knowledge creation activities often include repeated cycles through the process of informally sketching and outlining, authoring in electronic editing tools, perusing formatted drafts, then returning to informal markups, annotations, and drafting on paper. To build tools supporting work with both scanned documents and complex electronic documents such as web pages, a critical achievement will be the ability to cope sensibly with mixtures of formal text and less-formal image content including handwriting, diagrams, and graphics.

Knowledge-Sharing Technologies from Digital Libraries

Knowledge sharing technologies from the area of digital libraries have origins in library and information sciences, in information retrieval, in multimedia, and in human-computer interaction. A digital library can be seen as an example of a knowledge-sharing arena, and thus the technologies are often particularly relevant to the value propositions we have articulated. Several of the recent technologies that have emerged from this area have been Level 2 technologies, including interoperability middleware and metadata inference/processing, but we consider them here because of their immediate relevance to our value propositions. Progress also has been made on higher level technologies, including resource discovery and annotation. These technologies, though relatively new, are already starting to appear commercially.

Example: An environmental scientist wishes to search for scientific articles on pollution. He does not want to search the Web at large, as he is interested only in reading vetted scientific articles. Accordingly, he searches a federated set of digital environmental science libraries. These libraries each have their own distinct access methods, but interoperability "glue" binds them together so that he needs to issue only one query that is automatically translated into the form each of the libraries requires. In addition, an algorithm for resource discovery ensures that only the most useful individual libraries are searched at all, thus saving him time and even money (since some of these digital libraries use different economic models, including pay-per-use).

Interoperability middleware

Information finding in an age of multitudinous online information sources can be both overwhelming and time consuming. The variety of query interfaces, query languages, and source vocabularies is staggering. Even if a user knows the eccentricities of each desired information source, it can be tedious to visit each source in turn and to then reconcile all of the resulting search hits. State-of-the-art interoperability technologies address this problem by introducing software proxies and mediators that can adhere to a common query protocol on one side and a source-specific query protocol on the other side. Beyond protocol translation, these proxies can translate expressions in a common query language into the form required by the particular information source. The Xerox AskOnce product has this capability, as does the Stanford InfoBus. Yet challenges still remain for this technology, most notably semantics and robustness. To avoid semantic mismatches between the common query language and the target query languages, most current interoperability technologies ensure that the common language is close to a lowest common denominator language. Robustness is an issue when proxies represent sources that frequently change their program interfaces. These issues aside, the potential benefits of interoperability are significant enough that these technologies are likely to become a staple of knowledge sharing solutions in heterogeneous environments.

Metadata

Metadata technologies range from technologies for inferring metadata from content to technologies for translating metadata and processing metadata. A number of metadata standards are important in digital libraries, including USMARC, the Dublin Core, Z39.50 Bib-1, and FGDC. The Stanford Digital Library project has initiated work on translating from one metadata standard to another for the sake of search interoperability. On the Web, RDF (the Resource Description Framework) provides for metadata interoperability, with a particular emphasis on facilities that enable automated processing of Web resources. As usual, syntactic metadata translation is well within the state of the art, but semantics are problematic. For example, if one metadata standard includes a field for illustrator but another standard does not, it is impossible to do lossless translation from the latter standard to the former. Note that many different types of metadata are employed. Metadata describing the provenance of a document, the content of a document, and the economic/legal terms and conditions surrounding the document's use are common. Making use of this metadata can improve the sensemaking experience. For example, metadata can be invaluable for information filtering. State-of-the-art systems like the PARC Placeless system find that active metadata properties allow useful computations to be associated with documents.

Resource discovery

Even if full search interoperability is achieved, there are still significant costs associated with searching a large number of sources at once. Some sources may be slow to return results, others may return lesser quality results, and still others may be financially expensive. Technologies for resource discovery work to prune the set of sources to just the most relevant set of sources before a detailed search begins. Sources may be differentiated by statistical measures of the associated content, by metadata describing the source at different levels, by usage statistics, or by the results of probe queries. State-of-the-art technologies for resource discovery have been developed at Stanford, UC Santa Barbara, and other locations. Challenges are to make resource discovery computationally efficient and effective.

Annotation

Annotation is traditionally a human-computer interaction technology, but we consider it in the context of digital libraries because of the attention it has recently received in this area. One of the mantras of digital libraries is that there is more to digital libraries than search. Traditional libraries themselves are more than warehouses of books: not only do they filter, archive, preserve, and make documents accessible, but they also provide people with places to read, photocopy, and interact. The design of digital libraries can benefit from understanding the cyclic model described in Appendix G. Digital libraries can provide people with facilities to obtain and interact with materials. Annotation is one example of a technology that is geared towards interaction. Annotating a document often helps a reader to understand a document better, even when the annotation is merely highlighting. Comments that are left behind can be useful to a reader in a future rereading, as index terms to the document, and as communication to other readers of the document. Current annotation technology addresses issues ranging from how to create annotations to how to display comments to how to share comments appropriately. FXPAL, Stanford, PARC, UC Berkeley, and Georgia Tech are just a few of the institutions to have state-of-the-art annotation technology. In addition, companies like ThirdVoice and the nonprofit Foresign Institute are already bringing state-of-the-art annotation technologies to the market and to the public.

Knowledge-Sharing Technologies from Document Management

The area of document management is particularly relevant for Xerox to consider when it comes to knowledge sharing. Document management systems are successfully in use in many companies, and thus offer a platform upon which knowledge sharing solutions can be built. Document management systems perform the first necessary step in archiving, versioning, and sharing documents. In addition, they can provide infrastructure for workflow. Current document management systems are successful at managing the lifecycle of a document; work on interoperability standards for document management systems continues to progress. The main challenge facing document management systems is their integration into the larger work practices in which they are embedded.

Example: A purchasing department in a large corporation needs to route its paperwork along a prescribed authorization path that involves various members of the department. It also wishes to monitor purchasing patterns. The department is able to use a document management system into which the paper documents can be scanned, archived, and automatically sent along their authorization route. In addition, the system is interfaced to a data mining system that periodically runs over the purchasing archive and detects patterns.

Repositories

Repositories are essential to document management technology. State-of-the-art repositories offer document archiving, document compression, access control, and versioning. Repository interoperability has also become an important feature. Standards such as ODMA and the Web-based WebDAV have emerged as interoperability "glue." The Web has not only brought a new standard to document management systems; it has also introduced new issues for consideration. The creation of XML-based repositories and the introduction of Web-accessible repositories are now areas receiving widespread attention by companies like Documentum and Lotus Notes. The simple Web-based repository DocuShare has been successfully deployed to companies that require straightforward Web publishing.

Document services

Document services provide additional value to users of document management systems. A typical document service is format conversion. For example, document management systems that make their documents accessible in PDF likely require conversion services. PARC's Documents.com system exemplifies the state of the art in document services. It provides an infrastructure and consistent user interface to a multitude of composable services, including conversion, OCR, summarization, printing and delivery, and glossing.

Workflow

Workflow technology, at the intersection of document management and human-computer interaction, strives to facilitate the coordination, communication, and control of organizational processes. This primarily rule-based technology is the subject of standards efforts and is an important tool in organizational management. Workflow issues include security and flexibility. These issues have become increasingly important as companies decentralize and undergo rapid change. Workflow rules often turn out to be too rigid for companies with more freewheeling patterns of interaction. In addition, the physical/virtual boundary can be problematic for workflow. Actions taken and documents produced in the physical world do not always make their way into the document management system, leaving the system with an incomplete record or a subset of the documents it should have. For accounting, data mining, or trend spotting, this incompleteness can be an obstacle.

Knowledge-Sharing Technologies from Human-Computer Interaction

Advances in human-computer interaction (HCI) include not only new technologies but also innovative designs, design methodologies, and evaluation strategies. An interdisciplinary field, it often produces domain-specific contributions that are then classified as belonging to that domain. For example, we will discuss relevance feedback technologies in the section on Information Retrieval. Thus, summarizing the technological state of the art is difficult for this field. One generalization we can make is that the state of the art in HCI is much more advanced for 2-D graphical user interfaces intended for standalone personal computers than it is for other types of platforms and work environments. Topical areas in HCI that are particularly significant to the realm of digital knowledge sharing range from user modeling to visualization to interactive rooms. The following example highlights just one of the many ways in which current HCI technology and methodology could contribute to a knowledge sharing solution.

Example: A large company has just hired a new vice president who comes from another industry. The vice president is interested in understanding the organizational structure of the new company. A standard tree view of the company's online organizational chart is hard to explore, navigate, and manage. State-of-the art HCI technologies could provide the vice president with better tools for exploring this hierarchy. 2-D examples of state-of-the-art hierarchy explorers include the hyperbolic browser (developed at PARC), elastic windows (developed at University of Maryland) and even dynamic outline-style browsers like Microsoft Explorer. 3-D hierarchy viewers like the PARC Cone Tree are likely also be useful for this task once 3-D work environments become more common.

In this section, we survey a potpourri of HCI technologies, including not only visualization technologies and interfaces, but also hypertext, interfaces, user profiling and modeling, groupware, recommendation systems, and non-desktop environments.

Visualization

Good visualizations of data and information leverage human perceptual capabilities in order to improve sensemaking. *Readings in Information Visualization,* a recent collection of classic papers in the field, summarizes the state of the art in this area. Note that some of the classes of technologies that this volume includes in visualization (e.g., information workspaces) are covered elsewhere in this white paper, primarily under interfaces. The classes of visualization technology that we cover here include mapping data onto visual structures, interaction techniques, focus+context techniques, and document visualization. Technologies for mapping data onto visual structures fall into four major categories: techniques that position data on orthogonal axes, techniques that position data on many or complex axes, techniques that employ trees to arrange data, and techniques that use networks more generally to display data. The placement of data into visual structures is only part of the visualization story. Interaction techniques play an important role in visualization. Many techniques allow for the selection of a subset of data in the visualization, often in order to locate data or to reveal patterns. Other techniques give the user control of data transformations. Magic Lenses, developed at PARC, provide the user with control over the current view onto the data. Dynamic queries, developed at the University of Maryland, give the user direct manipulation of value ranges for variables. These technologies are making their way into the commercial realm. Likewise, focus+context technologies, which enable users to see simultaneously both a large-grained contextual view and a fine-grained focus are entering the commercial arena. Classic systems in this area include a number of technologies from PARC, including the Perspective Wall, Table Lens, and the Hyperbolic Browser. Document visualization systems, on the other hand, are still emerging as a technology. One of the important issues in this area is how to map document attributes onto spatial dimensions. Systems like the Galaxy of News from the MIT Media Lab use topic similarity (measured using Information Retrieval metrics) as the basis for a three-dimensional mapping of document collections.

Hypertext

Popularized by the Web, hypertext technologies date back several decades to the early vision of Vannevar Bush. Many advances have been made that are not incorporated into the infrastructure of the WWW. For example, hypertext technologies for bi-directional links, typed links, and versioned links have all been incorporated into research systems that predate the Web. Current focal points for advancing hypertext technologies include scaling to large numbers of nodes, providing better authoring tools for hypertext, addressing the problem of navigation in hyperspace, integrating multimedia into hypertext, and dynamically creating hypertext structure. Centers of excellence include FXPAL and PARC, Texas A & M, and Brown.

Interfaces

Summarizing the state-of-the-art in interface technology is made complex by the close relationship that user interface technology has with design. A user interface that *works* for its users is often the result of understanding the user's work practices and context (often by performing ethnomethodological studies or other anthropologically-informed inquiries), developing new user interface technologies that are appropriate for the user's context, prototyping a design for the user, evaluating how the user interacts with the prototype artifact, and then iterating on this process. In addition, note that some of the technologies we have discussed elsewhere (e.g., focus+context, dynamic queries, Magic Lenses) could be described as interface technologies. Nevertheless, we provide a sampling here of a few other user interface technologies that are particularly relevant to the set of value propositions we have outlined in this white paper: fluid user interface technologies, zooming user interface technologies, interactive workspace technologies, wizard user interface technologies, multiple view technologies, and programming-by-demonstration technologies. Fluid user interface technologies, which have their origins at PARC, employ techniques from the realms of animation and graphic design in order to unfold user interface changes without loss of the visual context. Zooming user interface technologies, exemplified by the Pad++ system (originally conceptualized at NYU and substantially improved later at Bellcore and the University of New Mexico), use a spatial metaphor of a zoomable surface in which different items can be located. Zooming user interfaces are often considered examples of another class of technology, information workspaces. These workspaces move beyond the desktop to provide users with virtual spaces in which to accomplish tasks. The PARC Rooms system, the PARC WebBook and Web Forager, and the Stanford DLITE interface for digital libraries are state-of-the-art examples of this technology class. Wizard user interface technologies, common in many commercial applications, use natural language text instructions and a limited set of user responses in order to guide a user through a task. Interfaces that provide multiple views offer users different perspectives on a single computational entity. Examples of multiple view systems include DataSplash, Visage, and EDV. Finally, programming-by-demonstration technologies allow users to build patterns of computational behavior by providing examples of the desired behavior.

User profiling and modeling

Like hypertext, user modeling and profiling have been under study for years, but have recently received a burst of attention because of the World Wide Web. An important reason to model users is to provide them with user interfaces that are tailored to these needs. Users have different cognitive styles, degrees of experience, and even physical limitations. State-of-the-art user modeling techniques are almost always adaptive, and may be either statistically based or rule based. Many user-modeling technologies have inherited features from more general artificial intelligence technologies. In addition to the actual modeling mechanism, this area of research and development includes technologies for monitoring users and gathering the information upon which the models depend. For example, the commercial Internet company Alexa gathers data about where Web users travel on the Web and then uses the information to provide recommendations about which Web sites to visit. PARC is well-known for its strengths in characterizing user behavior. For example, PARC has developed models for information seeking behavior and for Web "surfing" patterns.

Groupware

Groupware technologies, as well as technologies from the larger field of computer supported collaborative work, are particularly relevant in the context of knowledge sharing. The traditional desktop metaphor (which of course originated at PARC) can break down when a group needs to accomplish a task together. Shared electronic whiteboards that could be controlled by an individual at the board and/or by the desktop computers of a group surrounding the whiteboard were invented at PARC. The Eureka system, also developed at PARC, has been successful in its mission of gathering and distributing tips from technicians. Shared calendars, shared presentation systems, and shared co-presence systems are other forms of groupware that have been successfully developed and evaluated. One of the important challenges in this area is to design systems that respect and extend users' notions of privacy and security into the digital realm. Awareness technologies that give users a sense of who is observing and what is being observed take an important step in resolving this issue.

Recommendation systems

Recommendation systems take as their premise the idea that decisions about what to read, what to buy, where to go, and more can be made easier with input from others who have similar overall preferences. Not surprisingly, recommendation systems borrow technologies for both user modeling and groupware in order to make systems that perform well for people. One key challenge for recommendation systems is the bootstrapping problem. Often, a significant amount of data is required before a recommendation system will perform well. Enticing users to provide that data in advance of receiving quality recommendations can be difficult. Strategies for coping with this problem include seeding the database with data acquired by other means and providing additional value to users beyond the recommendations themselves. State-of-the-art recommendation systems include those provided by Firefly, a startup that has its roots at the MIT Media Lab.

Non-desktop environments

For most people, the term "computer interface" suggests a windowed interface on a desktop computer with a mouse, keyboard, and monitor. Technologies for nondesktop-based environments are becoming more and more practical as computing power becomes cheaper and smaller. Indeed, we may be nearing the age of ubiquitous computing envisioned by Mark Weiser. At one end of the spectrum are interactive rooms. Technologies for these rooms—which may include interactive walls, floors, and furniture—are under development at a number of institutions, including GMD Darmstadt, Stanford, the MIT Media Lab, and PARC. Important components in these rooms may include smart matter, which embeds computation in materials. Virtual reality environments are very different in spirit from interactive rooms because the technology is at the foreground rather than at the background. At the same time, these environments are similar in that they enclose the user in a smart environment. At the other end of the spectrum is the small handheld computer that moves with the mobile worker. The recent popularity of Palm Pilots has demonstrated the market demand for such devices. Research into interfaces and applications for handheld devices has increased accordingly. PARC (the birthplace of ParcTabs), Compaq Research, Stanford, Berkeley, NYU, and Georgia Tech have focused efforts in this area.

Knowledge-Sharing Technologies from Information Retrieval

Reading has long been associated with knowledge acquisition. Ever since the invention of the printing press, the number of documents available cheaply and easily to readers has grown at an enormous rate. In recent years, more and more of these documents have become available online, thanks to the World Wide Web. The field of information retrieval has focused on helping users to find relevant online documents and to browse online collections of documents. The current state-of-the art is fairly good at term-based search and clustering, but accounting for higher-level semantics remains an important challenge. The technologies of information retrieval are intrinsically bound to our value propositions around finding and organizing information. The following example is one of many that illustrate how current information retrieval technology could contribute to a knowledge sharing solution.

Example: A computer scientist wishes to learn more about XML, a topic that has received increasing coverage over the last several months. A wealth of online documents about XML is available, and so the computer scientist has trouble determining where to begin. She turns to the Google search engine first, and performs a keyword-based search for "XML." The Google engine uses the link structure of the Web to order its search results, making it likely that authoritative Web pages for XML will be receive high rankings. In addition, a clustering algorithm is applied to the top 50 results from Google, making it possible for the computer scientist to browse through the top ranked results at a higher level.

In this section, we cover the technologies employed in keyword-based query systems, categorization and filtering, clustering, query interfaces, question answering and information extraction, and summarization. Note that not only is there overlap between information retrieval and natural language processing, but there is also overlap between information retrieval and the domains of multimedia, human-computer interaction, and digital libraries. In fact, many of the recent advances in the field (for example, the development of the Google Web search engine) have taken place at the intersection of multiple domains of inquiry.

Keyword-based search

Systems that allow users to type in keyword-based queries are increasingly prevalent on the Web. Research centers of excellence in this area include AT&T Labs, the University of Massachusetts, various institutions in the United Kingdom (Glasgow, Sheffield, Microsoft Research), and the University of Washington. Typical keyword-based retrieval systems work primarily by performing term matching on the entered keywords (often modeling documents as vectors of words) and are evaluated by the metrics of precision and recall (or other statistical variants of these). These measurements are very dependent on a corpus. For the well-known TREC corpus, typical measures remain at sub-40%. The problem is that term matching is insufficient to determine whether a text is "about" the same topic as the query. Not only may a term represent different concepts in different contexts ("bank" as financial institution vs "bank" as a heaped up mound of earth), but different terms or collections of them may represent the same concept ("forecast" and "prediction"), again depending on context. On the Web, other issues further affect keyword-based search. Users typically issue very short queries, on the order of one to two keywords long. Discerning what is meant by these short queries is problematic (the upcoming section on query interfaces discusses work in the area of better eliciting user needs). The economics of search engines also may play a factor in search engine performance. Given the prevalence of advertising as a revenue source on the Web, many search engines lack incentive to provide a user with a set of results that causes the user to leave the search engine site promptly. Nevertheless, the state of the art is progressing in several directions. Using non-linguistic metadata, authority-based systems like Google and the IBM CLEVER system are proving to be valuable, as are usage-based systems like Direct Hit. Commercially, systems that are use symbolic and statistical techniques are enjoying success when targeted to a particular domain.

Categorization and filtering

Like ad hoc retrieval, most commercial approaches to categorization and filtering tend to be very term oriented, relying on simple, user-specified Boolean queries as categorization or accept/reject criteria. Accordingly, the same semantics-related obstacles and challenges arise for this set of technologies. In contrast, several research approaches use statistical learning to identify texts "like" some previously categorized ones. Specific techniques include the machine learning algorithms k-means, nearest neighbor, support vector machines, and neural networks. Current centers of research excellence include AT&T Labs, PARC, and CMU. Some challenges include large-scale categorization, which typically has the attendant problem of sparse data. Of note also is Murax, an innovative term-based approach for hierarchic fragment categorization developed at PARC and further developed by Inxight. Unlike keyword-based retrieval, additional sources of meta-information (e.g., links, usage data) have not been incorporated into the algorithms for categorization and filtering. This may be a potentially fruitful area for research. Current centers for excellence include AT&T Labs and CMU.

Clustering

Text clustering differs from categorization in that the categories to be found are not predefined. It is potentially of considerable value for text mining and exploratory text analysis. While current clustering algorithms do not always produce "natural categories," these algorithms can create a partitioning of a document collection that is useful when browsing a large corpus. For example, the technology has proven valuable for exploring large collections of legal documents. On the interface side of information retrieval, browsing interfaces like Scatter/Gather from PARC use computationally efficient clustering algorithms to provide users with overview of a collection. Specifically, fractionation is a highly accurate algorithm that is used to do static offline partitioning of the whole corpus, while buckshot is a fast clustering algorithm for reclustering smaller subsets of the corpus. Centers of excellence for clustering work include PARC, the University of Washington, Brown, and the University of Tennessee (for clustering based on latent semantic indexing). Of recent import is the development, by Thomas Hofmann, of an EM (expectation maximization) based algorithm for clustering that can deal with texts involving multiple topics. Evaluation of clustering algorithms is even more problematic than evaluation of retrieval algorithms. Determining what a good set of clusters would be is difficult for people, making evaluation difficult to perform. Finding better evaluation metrics is a key challenge for this whole domain. In addition, the question of how to incorporate additional metadata is again important here. The work at PARC on multimodal clustering focuses on this question.

Query interfaces

In general, users find it challenging to formulate queries that match their information needs. On the Web, user queries average 1-2 words long, which makes it difficult to determine exactly what the underlying need might be. Boolean queries have been a staple of online library catalogs, but many users have trouble formulating them. Mix-ups between "and" and "or" are common occurrences. Accordingly, new technologies for querying and new query interfaces are under study in information retrieval, digital libraries, and human-computer interaction. Relevance feedback, a classic information-retrieval technique for finding more documents, locates documents that are judged similar to a user-selected set of documents. Measurements of document similarity are term-based; again, semantics becomes an issue. Other forms of query-free retrieval and interfaces for information exploration provide some alternatives to relevance feedback, though they also encounter the thorny problem of semantics. Free-text query systems allow users to enter full sentences, which are analyzed using natural language technologies. For example, the DR-Link system from TextWise (a spin-off from Syracuse University and Manning Napier Information Services) has successfully marketed a free-text system to organizations with large, highly valued databases (e.g., the patent office, Edgar Online). In contrast to free-text query systems, which generally elicit queries that are longer and more expressive than keyword-based queries, query-free retrieval technologies remove the need to formulate explicit queries entirely. Query-free retrieval involves monitoring a user's activity and finding documents that appear relevant (for example, one system watches what a user types during document creation, and looks for documents based on recognizable keywords). The SenseMaker interface (from Stanford University) uses document meta-information as an organizing criterion for grouping search results, and then allows users to treat groups of documents as pseudo-queries. The state of the art for visualizing query results also continues to advance. Systems like TileBars (from Berkeley and PARC) give users finer-grained information about how their results match their queries; MDS visualization algorithms are valuable for displaying clusters of results.

Question answering and information extraction

The technologies under this heading are concerned not with document retrieval and/or grouping, but with extracting specific pieces of information from document collections. Domain independent systems capable of fully achieving this goal are far beyond the state of the art. Nevertheless, interesting and useful approaches have been developed. The PARC Murax question-answering system represents the state of the art in systems employing smart pattern matching to identify phrases likely to represent all or part of the answer to a question. The commercial system AskJeeves also employs pattern matching, but rather than matching a user's query to text, the query is matched to a hand-constructed database of predetermined questions. Natural language technologies that play a key role in systems of this kind are phrasal analysis and partial parsing. One of the key challenges for these systems is the question of how to determine when a good answer is not available and what the fallback strategy should be in this case. In 1999, the well-known Text REtrieval Conference (TREC) workshop series launched a track specifically devoted to question answering. This track is likely to drive work in this area.

More demanding tasks, requiring the instantiation of templates describing events or objects of interest, are generally considered to fall in the category of "information extraction." For example, an event template for business acquisition event might require the determination of the two firms involved, and specifics of the acquisition terms. Evaluation for such systems is made possible by the Message Understanding Conference, which provides participants with templates and a corpus. Systems can then be compared by applying precision and recall metrics to fields within a template. In general, these systems can work fairly reliably if the domain is very restricted and the genre remains constant. SRI and other research institutions have successfully used finite state automata in this application domain. In fact, a recent spinout from SRI is looking to commercialize this technology. Centers of excellence include not only SRI, but also the University of Massachusetts, BBN, and

SRA. Challenges for the set of technologies echo those mentioned for other information-retrieval domains: semantics, scale, speed, and domain sensitivity. While restricting systems to narrow domains alleviates the semantics problem to some extent, because it allows the development and use of purpose specific knowledge bases, it has been estimated that moving even the best of current systems to a new domain requires approximately six months.

Summarization

Summarization is an active area of research in information retrieval oriented natural language processing. Centers of excellence include PARC, University of Massachusetts, GE, Columbia University, Cornell University, the University of Southern California, and the University of Edinburgh. Summarization techniques differ along many dimensions. First, the sizes of the source documents imply the level of summaries that are appropriate, and, indirectly, the kinds of techniques that are suitable. For small documents, summarization typically consists of extracting (or generating) some key phrases or sentences, while for large documents, finding key paragraphs is a significant approach. The techniques used in summarization are generally synthetic, combining methods of identifying the most salient words, phrases or concepts, methods of text segmentation (such as the PARC TextTiling method) to identify subtopics, and superficial discourse processing (e.g. cue words) to assist in the detection of candidates for extraction. Deeper methods, similar to those of information extraction, are also the subject of experimentation, to find conceptual representations of primary content, from which natural-sounding summaries can be generated. While the current state of the art focuses on indicative summaries (descriptive of the document's content), more research is underway on informative summaries that use more linguistic information in order to provide the main findings of a full document in a summary.

One of the immediate challenges facing the area is determining how to evaluate the existing summarizers. Precision and recall can be used if summarizers work by sentence extraction, but not if they work by keyword extraction. One current approach to evaluation is to evaluate summaries in a larger task context. For example, one can compare relevancy judgments that are made with respect to the full document versus judgments that are made with respect to the summary. Another criterion for evaluation is fluency – a measure of how natural sounding the resulting summary is. Another challenge for this set of technologies is the development of genre-specific summaries. For example, e-mail summaries are likely to have different characteristics from book summaries. Summaries for non-text genres are also beginning to receive attention, both at Xerox (work on audio summaries) and elsewhere (for example, recent Intel work on video summaries).

Knowledge-Sharing Technologies from Multimedia

Although many of the examples in this white paper focus on text, sensemaking and knowledge sharing of images, audio, and videos are becoming increasingly important to corporations as these genres become accessible and ubiquitous on both the intranet and the internet. Questions of how to acquire, share, and act upon documents in these genres can be informed by ongoing research in the field of multimedia. This field has seen rapid progress in the last ten years. Speech recognition has advanced to the point where it can be incorporated into commercial products. Image browsing and search, which once received little attention in comparison to text browsing and search, is blossoming. Multimodal systems that require a combination of speech recognition, natural language, and image processing techniques look increasingly promising for the study of genres like video. The continuing challenge for all of the recent work is to make multimedia systems less reliant on training and on circumscribed domains. Limited domains are critical because they provide a way around the difficult issue of semantics that arises for so many of the knowledge sharing technologies. The following example illustrates how current technologies could be of value in a knowledge sharing solution.

Example: Trademarks and logos are important elements in the virtual collage of words, sounds, and images that corporate identity experts employ when representing a company to the public. If a logo is used inappropriately or a derivative logo is used by another company, the integrity of the company's identity image can suffer. How can a company know that its logo is in use elsewhere? Image matching technology that can identify images that are on the Web and are similar to the ones owned by the company can be used in keeping the company informed about the uniqueness of its logos.

The field of multimedia covers a diverse array of areas, including such well-studied areas as video streaming, multimedia synchronization, and video compression. In this section, however, we consider only those technologies that are Level 3 technologies. Specifically, we discuss image retrieval and browsing (under which we subsume video retrieval and browsing) and speech recognition.

Image retrieval and browsing

Internet-focused companies like Virage have drawn attention to the value of image retrieval. Text-based search engines like Excite, AltaVista, and Google play important roles in widespread knowledge sharing. In comparison, image retrieval is still in its infancy. State-of-the-art retrieval systems depend on using low-level image features—such as color histograms, shape, and texture—to match images. Just as in text-based information retrieval, performance is typically measured in terms of precision and recall. Image retrieval of labeled images remains highly preferable to retrieval of non-labeled images. Current issues and challenges include determining what features are most valuable for matching, what distance metrics to use, and what query paradigms are appropriate for end users. One apparent trend in new query paradigms is to combine image retrieval with image browsing in order to make systems more intuitive for end users. For example, a common technique is to ask users to compare two images in order to assess which one is closer to the desired type of image. A recent variant on this approach is to use this information not merely to converge on a set of features, but also occasionally to diverge to prevent premature feature choice and to encourage exploration of the full space of images. Overall, progress on image retrieval and browsing techniques is happening not only in academia and industrial research labs, but also in commercial companies in this area. Centers of excellence include Columbia, CMU, IBM Almaden, NEC, and various startup companies (e.g., Virage). Even though many challenges remain in image retrieval and browsing, many of these institutions have simultaneous research and technology development programs for video in addition to static images. One reason for this trend can be found in looking at how video retrieval and browsing are often performed. Combining data from multiple modalities appears to be a valuable strategy for making progress in this domain. In a sense, this trend is similar to the trend in information retrieval around using metadata to improve text search results. While no one modality may be sufficient for identifying useful video segments for a user, the combination of speech, video, and sometimes transcripts may make the problem easier. Of course, work is also underway on technologies that are specific to video. For example, video summarization technologies have been developed at the MIT Media Lab, Interval Research, and Intel.

Speech synthesis and speech recognition

Both speech synthesis and speech recognition have great potential as knowledge sharing technologies. Speech synthesis becomes increasingly useful as work moves off of the desktop and into the fabric of our environments. For example, speech synthesis could be used to read documents (or summaries of documents) while the listener is driving. In terms of the state of the art, speech synthesis technology is fairly good at making automatic speech intelligible, but challenges remain in making it more natural. Work on intonational phrasing seeks to address this issue.

Speech recognition promises to make captured speech as accessible (or almost as accessible) as written documents. This has implications not only for acquiring information, but also for sharing information. On the acquiring side of knowledge sharing, audio transcripts of meetings, recorded speeches, and video may become searchable by end users. On the sharing side, end users may find that speech is a valuable input modality to a variety of knowledge sharing applications. Already, microphones are beginning to appear on personal digital assistants. Accurate speech recognition would make this feature more ubiquitous. Already, high accuracy is possible in limited domains. For example, the Nuance speech recognizer has 97% accuracy as of this writing for the specific domain of obtaining stock quotes. Situations that are more demanding for speech recognition are those in which there is continuous speech, a noisy audio channel, a variety of possible speakers, and an arbitrary domain in which disfluencies are common. In addition to improving the accuracy of recognition in these contexts, a technology trend in this area centers on topic identification and detection. FXPAL has looked recently at speech summarization in the context of video. Centers of excellence in speech recognition and speech synthesis overall include Microsoft Research, AT&T Labs, Nuance (an SRI spin-off), MIT, Cambridge University, Carnegie Mellon, and BBN.

Knowledge-Sharing Technologies from Natural Language Processing

Technologies that automatically analyze and process the content of text can contribute to several of our articulated value propositions. Ideally, the automatic methods applied would provide results that humans judge to be appropriate and correct. Today, state-of-the-art natural language processing technologies are advanced with respect to the formalization of linguistic knowledge, but still primitive with respect to the formalization of world knowledge. Nevertheless, knowledge sharing solutions can benefit from what the field of natural language processing already has to offer, ranging from rule-based and statistically-oriented methods for representing and processing text to large linguistic knowledge bases that are accessible from higher level systems. The following example illustrates how natural language processing could contribute to a knowledge sharing solution:

Example: A pediatrician wishes to read the latest literature on lead poisoning, as he has just made this diagnosis for a young patient. He turns to a search engine that covers the medical literature. A search for the word "lead" might uncover not only articles on "*lead* poisoning" but also articles that include phrases like "*lead*ing causes of amyotrophic lateral sclerosis." Natural language processing techniques could resolve this problem by using part-of-speech tagging to disambiguate the use of the word "lead" in a given document.

We consider here those technologies from natural language processing that appear particularly relevant to building knowledge sharing solutions in the relatively near future. For example, we consider tokenization but not discourse processing, though the latter is the subject of active research. On the analytic side, we discuss language identification, tokenization, morphological analysis and stemming, part-of-speech tagging, phrasal analysis, shallow parsing, full parsing, and generation. On the applied side, we discuss database query interfaces and machine translation. Several other technologies with roots in natural language are discussed in the section on information retrieval.

Language identification

When combined with user profiling and modeling, language identification can serve as a filter to a search engine by only returning documents in the information seeker's language. It can also act as a pre-step to the automatic translation of documents that match a user's needs. The technology itself, which is based primarily on statistical techniques (notably letter trigram frequencies), is fairly mature; most languages can be identified with 99% accuracy on the basis of a single line to a paragraph of representative text. The technology has made its way to the marketplace (e.g., AltaVista) and so competitors abound. Nevertheless, Inxight and XRCE remain centers of excellence for this technology.

Tokenization

Tokenization technology is necessary for determining the word boundaries in a stream of text. This technology is used independently in keyword-based search, and is a first step in more inclusive analyses. When used independently, good tokenization can contribute to retrieval accuracy. Among other examples, tokenizers can isolate fixed multi-words, time and date phrases, and lexical representations of numbers ("twenty five"). A keyword-based search for the key phrase "sensemaking" could rely on tokenization to determine that documents with the word "sense-making" match the entered query. Like language identification, this technology is very accurate. Inxight has the state-of-the-art technology in that its tokenizer (based on finite state automata research from PARC and XRCE) is not only accurate, but it is also three to four orders of magnitude faster than the nearest competing tokenizer on the market. In fact, tokenization is considered a solved problem for many languages. Languages that are difficult to tokenize include some Asian languages, which are notoriously difficult for even native speakers to tokenize. Fuji Xerox has a state-of-the-art tokenizer for difficult-to-tokenize Japanese, also based on PARC research and now included as part of the Inxight offering.

Morphological analysis and stemming

Like tokenization, the analysis of words into their components (e.g. walked → walk + Past) can be used either directly in representations of document content, or as a step in further analysis. For morphological analysis, the finite state transducer (FST) approach developed at PARC and XRCE is state of the art. The technology is highly accurate, fast, and compact. In addition, it is completely bi-directional, making it useful for both the analysis and generation of natural language. Furthermore, Inxight (based again on research from XRCE and PARC) makes available a large number of language-specific transducers.

Part-of-speech tagging

Part-of-speech tagging is critical for realizing the scenario described at the beginning of this section. Using the results of morphological analysis, which generally produces a number of candidate interpretations, part-of-speech tagging is a form of partial parsing that disambiguates those assignments. In the opening pediatrician example of this section, recognizing when "lead" is used as a noun rather than as a verb is necessary for determining if a document matches the query. Part-of-speech tagging is accomplished either with statistical techniques (usually, Hidden Markov Models) or rule-based techniques (usually, finite state automata). This technology has also made its way to the marketplace, but Inxight's offerings (based on research from PARC and XRCE) again put it at the top. Most of the other companies are small and focus on a few languages.

Phrasal analysis and shallow parsing

The goal of phrasal analysis technologies is to find useful elements of syntactic structure that are valuable for specific applications. For example, the Inxight "thing finder" focuses on robustly locating mentions of named entities from categories such as people, places, and dates. This function is particularly useful for information retrieval. Another potential application of "thing finding" is document sanitizing. For example, consulting companies often need to systematically replace the names of real corporations and people with fictional names. "Thing finding" can also be useful for text mining, as IBM's text mining software shows: it incorporates a variant of "thing finding" as well as morphological analysis. Terminology extraction, another application of phrase analysis, builds up an index of the key phrases in a corpus. This index is useful for ensuring the consistency of terminology and terminology translation, as well as building tokenizers and lexicons for processing documents in the domain. In general, the underlying technology is not as highly accurate as the lower level technologies we have discussed. Some ambiguity problems can be overcome by combining the analysis with word sense disambiguation techniques. Again in this area there are both rule-based and statistical techniques used. PARC has expertise in the statistical methods and XRCE in a combination of rule-based methods.

Whereas phrasal analysis extracts syntactic elements, shallow parsing aims to identify some of the predicate-argument relations in a sentence. The predicate-argument relations of a sentence provide an indication of "who did what to whom," arranged in a standard format independent of the many paraphrase-variations in which a given event may be described. The standard format makes it easier to use this information for a variety of further knowledge-processing tasks such as higher-precision retrieval, summarization, and converting textual information into conventional database representations. In general, the problem of identifying predicate-argument relations accurately, consistently, and completely requires fairly sophisticated processing techniques and substantial grammars that characterize how predicate-argument relations are encoded in particular languages (see Full Parsing below). Shallow parsers operate with reduced ambition in the hope that simpler algorithms and simpler grammatical specifications will provide results that are good enough for many practical tasks while circumventing the problems of robustness, efficiency, and unneeded specificity that full parsers often encounter. Variants of the shallow approach, for example, are used for those question-answering and information-extraction tasks in which full parsing is still judged too slow to be practicable. XRCE Grenoble is among the pioneers in using finite-state technology for this purpose.

Full parsing

Full parsers also aim to identify the predicate-argument relations in a sentence, but they are more ambitious than shallow parsers in that they try to produce a complete collection of relations that is accurate and consistent. The more precise results of a full parser are required for tasks such as translation and fact retrieval that depend on a deeper understanding of a text. Contemporary full parsers generally take as input the results of prior tokenization and morphological analysis, and they rest on formal grammars that describe how predicate-argument relations are encoded by the words in a sentence and the order in which they appear. The formalism of Lexical Functional Grammar, developed at PARC and Stanford, is well-known for its ability to describe language structure succinctly and expressively. Full parsers usually produce at least two different kinds of analytic results, one a "parse tree" representing a syntactic analysis of the sentence structure, and the other a deeper interpretation that represents the predicate-argument relations, expressed either as annotation to the parse tree or as a separate structure (as in the case of LFG). The Xerox Linguistic Environment, developed at PARC with contributions from XRCE, is regarded as world-leading technology in the efficiency with which it carries out full parsing. LFG grammars that have been developed for English, French, and German are also recognized as covering broad ranges of grammatical constructions.

Often, parsers deliver multiple alternative structural results to represent the different interpretations of ambiguous sentences. Statistically-based methods for resolving the ambiguities are becoming increasingly common, focusing either on estimating the relative probabilities of specific kinds of alternatives, or, more ambitiously, developing full probabilistic grammars. The former approach is currently being investigated at XRCE and at various universities. The latter approach, while still exploratory, is becoming increasingly sophisticated; an example is the work at PARC (LFG-DOP) in developing a probabilistic grammar utilizing both parse-tree and associated feature-structure probabilities.

For some applications, parsing must be followed by explicit semantic analysis, to identify deeper "logical-forms" giving quantification relationships, and to isolate word senses (whose identification can be assisted by statistical processing of the larger text). In turn, this can be followed by pragmatic/discourse analysis to resolve inter-sentential references and obtain an understanding of the rhetorical structure of the document.

Full parsing is still a subject of ongoing research, but full parsers have found their way into the commercial world as the base for grammar checkers, database query interfaces, machine translation systems, information extraction vehicles, and for the creation of annotated corpora used in statistical learning for other natural language processing tasks.

Generation

Whereas much of this section has focused on analyzing text, natural language generation seeks to produce text that is understandable and fluent to human readers. The inputs and outputs for natural language generation systems vary depending on the intended application. State-of-the-art generators can produce good results when the input material is largely factual and when the structure and vocabulary of the output are pre-determined. Some advanced generators (e.g., the PARC LFG generator in the Xerox Linguistic Environment) work by performing a kind of "inverse parsing," using the same grammar and lexicon as a parser would employ. Other systems have generation-specific architectures.

Even when generator input is quite detailed, generator output can specify alternative orderings. Selection of the most appropriate ordering should be dependent on the structure of the surrounding text, for example, what is "given" and what is "new" information in the sentence. For less detailed input, questions of word selection and combination also arise; their resolution can be assisted by statistical information on text genre and collocation. Integrating these considerations into working generators is an ongoing challenge, as is the fluent use of conjunction, ellipsis, and reference (not just pronouns, but alternative forms of identification, such as "Bill Clinton" vs. "this president").

While these challenges remain, generators have already made their way into the market place. Generators are used commercially in machine translation and in database query interfaces. They are also used in situations where nonlinguistic data must be converted to linguistic form for general comprehension. Cogentex, for example, has a generator that automatically transforms meteorological data into weather reports for the Eastern Canada maritime provinces, and it also produces linguistic summaries of statistical reports. More experimental work in generation has focused on expressing summarization results (see the corresponding information-retrieval section earlier in this paper).

Database query interfaces

Like machine translation, natural language interfaces to databases have a long history; the first was IBM's Transformational Query Analyzer (TQA), dating back to the early 1970s. The technology typically associates one or more expressions (usually a verb and its arguments) with a database relationship. The expressions can then be employed in natural language queries by end users. Representative queries for a purchasing database might be "What department buys staplers?" and "Who supplies staplers?". More complex queries might imply relational joins and unions and quantitative processing. Users learn the controlled language needed by looking at query examples. Avenues of further development for this technology include loosening the permitted language, and allowing queries to range over both database information and text sources (e.g. XML-tagged text with database-oriented content).

The most obvious value proposition is that natural language interfaces allow people to query a database without needing to learn more formal languages like SQL (or to find someone who has). Another important potential value appears in high-level data base integration; the natural language forms can smooth over different ways of expressing the same information in the different databases.

To date, database query processors have not been a commercial success. One probable reason for early failures was simply the lack of a target audience; only recently has the phenomenon of a computer on every desktop become widespread. Microsoft has recently acquired an established firm in the areas, NLI, and has released the NLI query technology as part of the MS DB2 server.

Machine translation

Machine translation, an active area of research for the past several decades, aims ultimately to produce document translations that humans find fluent and accurate. Many challenges, including ambiguity and the use of idioms, are prevalent. Yet machine translation does not have to be perfected before knowledge sharing solutions can make use of the state-of-the-art techniques that exist today. For example, intelligence agencies during the last three decades have found machine translation useful for screening foreign language documents. PARC and XRCE together have a small French-English prototype tuned to the translation of technical manuals. One new center of excellence that has emerged in the translation arena is a commercial company, Lernout and Hauspie. It is a Belgium-based company that has been acquiring both natural language companies in general and translation companies in particular. It has recently acquired INSO, which was InXight's former biggest competitor.

Recently, cross-language information access (at the intersection of machine translation and information retrieval) has emerged as a magnet area for technology innovations. "Foreign language understanding aids" fall into this class. They are somewhere in between translation and bilingual dictionary look-up in that they do not translate the whole sentence but go beyond simple word by word translation by finding multi-word expressions of various types. For example, the French sentence "elle cherche toujours la petite bête" will not be translated as "she always looks for the little animal" but as "she always nitpicks." There is currently one product on the Web that takes this approach (http://www.babylon.com). XRCE has a more sophisticated prototype that also incorporates some syntax normalization (some word order differences between languages are neutralized). A second example of a cross-language information access technology is cross-linguistic information retrieval. The goal is to allow for queries to be expressed in a language other than the language of the target corpus. For example, this application would allow one to search for "leveraged buyout" in a French corpus (the French equivalents are "rachat d'entreprise financé par emprunt" or "rachat d'entreprise financé par l'endettement"). A recent DARPA initiative has served as the springboard for this research; a number of institutions are performing research in this area, including XRCE and the University of Maryland.

Knowledge-Sharing Technologies from Security

At first glance, one might think that security would hinder knowledge sharing rather than contribute to it. However, if we think about security technologies that allow people to put more trust into systems, to articulate and protect rights in intellectual property, to provide for secure e-commerce, and to allow for privacy (and sometimes even anonymity) then we can indeed see that security technologies are critical for knowledge sharing to take place. Centers of excellence for this domain as a whole include PARC, IBM Research, AT&T Labs, Bell Labs, Microsoft Research, Compaq SRC, SRI, MIT, Princeton, Stanford, UC Davis, UCSD, and Cambridge University.

Example: A company that issues highly coveted financial reports wishes to make these reports available on a pay-per-view basis to users of secure portable document readers (PDRs). The company uses a rights management language to describe how many printouts the customer will be able to make, how many times the customer can view the document, etc. A rights management system protects the rights assigned by the company. The system also provides information about usage of the document, which the company finds valuable for engaging in targeted marketing.

Trust management

Trust management is the area of security that considers access control in its broadest sense (note that some researchers continue to term this larger research enterprise as "access control"). It emphasizes reliance on digitally-signed statements and a more flexible mechanism for resolving statements into access decisions (e.g., various languages and/or logics as opposed to simple access control lists). Controlling information flow remains a challenging topic within this area, as it is difficult to determine and block all of the covert channels that might exist. Within the realm of trust management, the DPRL language and ContentGuard (its commercialized counterpart) exemplify the state of the art. They allow for the articulation of rights in terms of predicates and actions, and also provide protection of those rights.

Authentication

Like so many other technologies, authentication technologies received additional attention with the advent of the Web. Most of these technologies are based on cryptographic systems. State-of-the-art technologies typically employ hybrid cryptography, in which both asymmetric cryptography and symmetric cryptography methods are combined together. For example, public key cryptography might be used to gain a session key that is then used in a symmetric protocol. The symmetric protocol that is widely in use today is DES, but that is likely to change when NIST replaces it with a new choice. That choice will most likely remain the standard of choice for the next decade or two. In the realm of asymmetric cryptography (and public key infrastructure more specifically), many variants abound. Key management remains a critical issue for all of these schemes. In an Intranet environment, this problem is less severe because there is a natural root to the certificate hierarchy. In an Internet environment, the lack of a widely accepted hierarchy root has resulted in scalability becoming an important issue, as the schemes must devolve to n^2 schemes. Well-known hybrid systems that combine symmetric and asymmetric cryptography include PGP and SSL. Of these, SSL is the current pragmatic choice for Internet transactions.

Distributed systems security

Due to the complexity of the environment and the number of systems involved, distributed systems security has emerged as one of the new focal points for security research. It requires advances in both authentication and authorization. Two well-known technologies are SPKI and SDSI, which are delegation and attribute-based technologies. PARC has developed SPKI/SDSI schemes for repository security, as well as other schemes for anonymity in Web interactions.

Knowledge-Sharing Technologies from the Web

The emergence of the Web as a popular communications and publishing medium has spurred research and development in a number of areas, both in research institutions and in the legion of startups that have surfaced in recent years. Many of the new technologies that have arisen in the context of the Web can be grouped into one of the more traditional domains of inquiry we've presented already. However, others seem uniquely focused on the Web and we cover these here.

Example: A SOHO owner needs to determine what hardware configuration she should buy when she upgrades her current computer system in order to be Y2K compliant. She first uses a Web search engine in order to locate trusted reviews of the latest computer peripherals. She finds a report that promises to tell her what she needs to know and agrees to pay $2 in order to read this report once. The report includes not only text, but also a dynamically updating graph showing reader satisfaction with the various items covered in the review.

Web crawling

Web crawling is necessary in order to build the index that a Web search engine uses to answer user queries. Over the last few years, the technologies for crawling have become much more efficient and scalable as the developers of these technologies have raced to keep pace with the growing size of the Web. Even now, most Web search engines cover only a fraction of the pages available on the Web. In the early days of the Web, one of the main questions about crawlers was whether they performed a depth-first or breadth-first search. Many other issues have since become important, including how to distribute crawler code in order to search the Web more efficiently, how to recognize cycles and site identity on the Web, how to handle dynamically generated pages on the Web, how to handle intentional misrepresentation by Web sites, and how to decide what meta-information to keep for a Web page. Most of the state-of-the-art technologies for crawling come from the well-known search engines on the Web, including AltaVista, Google, Excite, Lycos, and Infoseek.

Dynamically generated documents

Documentation systems and other systems that predate the Web have long made use of dynamically generated documents. The number of dynamically generated documents has exploded with the advent of the Web and so we cover them here. Many documents are generated by the use of standard scripts known as CGI scripts, but new technologies such as Java servlets and database publishing technologies also contribute to the number of dynamically generated documents on the Web. State-of-the-art technologies allow for pre-designed graphics and pre-authored text to be mixed together with programmatically generated results. For example, the WebWriter system from PARC and Microsoft's ASP system allow computation to be embedded easily inside of Web pages.

XML

XML, the Extensible Markup Language, is intended to be a universal format for structured documents and data on the Web. A descendant of both SGML (developed in the 1980's) and HTML, it is a text format that uses tags to delimit structural elements of a document. XML is an important standard for information exchange and application interoperability because it provides a common syntax in which information can be expressed. XML does not address the issue of semantics. For two independent systems to exchange data using XML, the designers of the systems must agree upon conventions for the set of tags that will be used and what the expected semantics for those tags might be. One of the advantages of XML technology is that it promotes discussion about data exchange and prompts system designers to plan for interoperability. In addition to the many companies that are using XML (for example, Microsoft is a key proponent of XML), several technology providers are implementing and marketing XML tools, including parsers and viewers. XML is interrelated with other standards, such as RDF for metadata. At the level of document collections, topic maps represent the structure of information in document collections. The topic map syntax is encoded in XML and is an emerging standard for externalizing collections of references, representing relationships among topics, and for representing multiple views onto collections.

Economic models

The Web is emerging as an important new marketplace for the global exchange of goods and information. It is also serving as an experimental arena for the conceptualization, development, and testing of new economic models. Companies like Cisco systems have successfully turned to e-commerce to generate revenue. New companies like amazon.com have proven that even home computer users are willing to buy goods online. Web sites have discovered that subscription models, pay-per-view models, and advertising models are viable sources of revenue. In addition, companies like eBay have served large numbers of buyers and sellers through the new technology of online auction models. PARC is one of the key players in advancing and suggesting new economic models for the Web. Nonlinear pricing, novel auction models, and information marketplaces may fundamentally change not only the market, but also the field of knowledge sharing.

Other Knowledge-Sharing Technologies

To attain complete coverage of all the technologies that might contribute to knowledge-sharing solutions is a hopeless task. We have included here the technologies with which we are familiar and which seem particularly relevant to our enterprise. Without doubt, we have surveyed these technologies with a computer-science bias. The state-of-the-art in knowledge sharing today is not yet advanced enough that we can be confident that we have even covered the relevant technologies that we have invented and advanced within Xerox. We remain open to rediscovering and recognizing the value of omitted technologies within the domains of inquiry we have covered, as well as the possibilities of discovering and employing technologies from other domains of inquiry.

Contributions from Social Sciences

Although the focus of this Appendix has been on knowledge-sharing technologies, this seems a good place to acknowledge that technologies form only one set of building blocks in knowledge-sharing solutions. In this white paper, we have adopted a conversation-oriented perspective on knowledge. This perspective reminds us that scientific approaches to understanding how people interact, how people think, and how people work are critical for developing knowledge-sharing solutions that really work. The fields of psychology, anthropology, ethnomethodology, cognitive science, design sciences, and social psychology have much to offer us here. In addition, given our primary interest in supporting knowledge sharing within corporations, we can learn much from organizational theory and the larger areas of business and economics.

Appendix D. How do Corporations Acquire, Share, and Act with Knowledge?

How do corporations acquire and share knowledge, and take action with it? What are the barriers to acquiring, sharing, and acting with knowledge? What would constitute a "solution" for alleviating or overcoming the barriers? How can technology help? What are its limitations? What else is needed besides technology?

To address such questions requires having a clear idea of what constitutes knowledge and a model of how people and corporations create, use and share it. It is perhaps no surprise that this is both a research question and a practical question. Our investigation of how Xerox might address such questions took us quickly to Ikujiro Nonaka, who holds the Xerox Distinguished Professor in Knowledge at the Haas School of Business at the University of California at Berkeley. Together with Hirotaka Takeuchi, Nonaka wrote a best-selling book *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation.* This book has been influential in discussions around Xerox and is cited by many industry reports about the knowledge management industry.

The world has now shifted considerably since Nonaka's and Takeuchi's best-selling book was written. Published in 1995, the book predates the rapid rise of the Internet, the widespread use of corporate intranets, the emergence of corporate information portals, the new knowledge management companies, and many of the now-widespread trends in corporate reorganization. In trying to understand the knowledge-sharing opportunity for Xerox, we found it necessary to go beyond the conclusions and time-frame of this book and to combine it with more recent thinking.

In the following we develop a perspective that considers the following topics: differentiating between explicit and tacit knowledge, processes of knowledge creation, knowledge and social networks, communities of practice, knowledge and conversation, how the written word differs from the spoken word effecting knowledge and sensemaking, communities of practice and social capital, barriers to sensemaking, and value propositions for knowledge technology. We present our basic assumptions and a model for how corporations create, use, and share knowledge. In the end, we recognize that our collective understanding of knowledge is still growing. Although this writing represents our current understanding, it will not be the "last word." The informed study of how knowledge is acquired, shared, and acted on will continue to co-evolve with technology over the next few decades. It is worthwhile for Xerox to carry out research in this area in order to inform and accelerate the development of products that better serve us and our customers.

Dimensions and Definitions of Knowledge

We begin with basic definitions and ideas from Nonaka, Takeuchi, and others.

> First, knowledge, unlike information, is about *beliefs* and *commitment.* Knowledge is a function of a particular stance, perspective, or intention. Second, knowledge, unlike information, is about *action.* It is always knowledge "to some end." And third, knowledge, like information, is about *meaning.* It is context-specific and relational. (Nonaka & Takeuchi, page 58)

Nonaka's definition of knowledge fits well with the knowledge management theme that the core value proposition of knowledge management technology is about speeding or improving business processes. Knowledge is about action. Business processes are by their nature about doing things. An objective test of whether someone "has" particular knowledge is whether they can do something that requires that knowledge. This action-orientation for defining knowledge corresponds to the intuition that mere "book learning" is different from real experience. A sail boating enthusiast talks about people who have obtained a captain's license, but to whom he would not trust his life on the open sea to because they lack experience.

Nonaka and Takeuchi take care to distinguish between knowledge and information. In Appendix A, we characterize data as being about physical representations ("bits") that can be stored or transmitted; information is about messages and writings that can be interpreted by humans. People often ascribe different meanings to the "k-word." In his textbook *Introduction to Knowledge Systems,* Stefik considers seven relevant meanings of the word "knowledge" from Webster's dictionary and other sources. Nonaka and Takeuchi say that knowledge is about action; it is to some end; and it is about meaning. Whenever we talk about meaning or action or experience, there is necessarily some agent who constructs the meaning, takes the action, or has the experience.

When we refer to an "experience," some agent must interact with the world to have the experience. Usually this agent is assumed to be a person. ... Agents are also involved when we consider written representations of knowledge. One of the insights in the past few years about the nature of knowledge and knowledge systems is that when meaning is attributed to systems and symbols there is necessarily an agent. … symbols do not have meanings inherently; they are assigned meanings by an observer/agent and the assignment of meaning is in the mind of the observer. (Stefik, page 295)

Introducing the notion of "agent" to the definition of knowledge is useful for illuminating many of the salient phenomena and difficulties involved in creating knowledge and in understanding what's going on in processes involving learning and communication. In an information economy, much of the relevant knowledge for businesses is about creating, finding, and using information.

Following Michael Polanyi, Nonaka and Takeuchi characterize a dimension of knowledge as the epistemological dimension, which distinguishes between *tacit knowledge* and *explicit knowledge*.

Tacit knowledge is personal, context-specific, and therefore hard to formalize and communicate. Explicit or "codified" knowledge, on the other hand refers to knowledge that is transmittable in formal, systematic language. (Nonaka & Takeuchi, page 59)

The transition between explicit and tacit knowledge often results in shifts in precision. An agent using tacit knowledge to perform a task may be perfectly adept in the performance of the task, but may be unable to explain exactly what they know. When an expert is asked to explain something, the process of articulation is often experimental, while the expert reaches for precision and clarity.

Experts sometimes can give only vague accounts of their thinking. But they do know how to make their thinking more precise through experiments. In this experimentation process, knowledge is often created or discovered. (Stefik, page 295, 296)

Nonaka and Takeuchi give numerous examples of tacit knowledge. One example is the knowledge of a baseball player who consistently bats well but can't "explain" how he does it. Another example is a baker who made the best bread in a city. The secret to his bread making – a way of twisting the dough as well as kneading and stretching it -- was not understood by others until they apprenticed under him and rediscovered the technique through careful observation and imitation. The baker and baseball player carried the tacit knowledge in their hands rather than just in their heads.

According to Nonaka and Takeuchi's theory of organizational knowledge creation, knowledge is created by a process in which it is converted between tacit and explicit forms. The four kinds of conversions are summarized in Figure D-1.

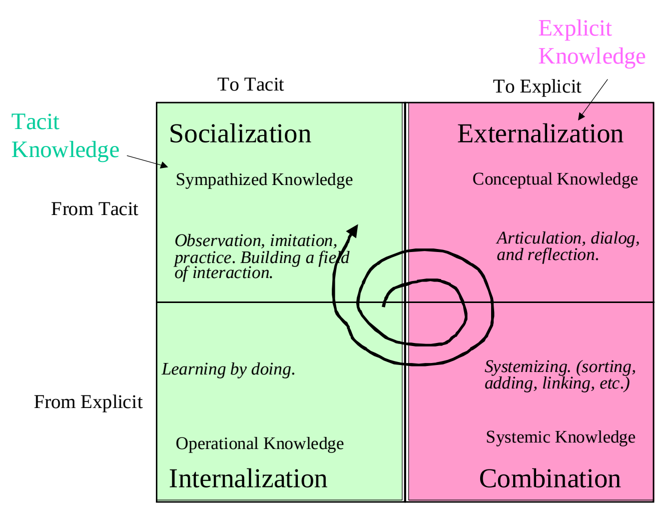

Figure D-1. The "seci" diagram **-** four modes of knowledge conversion. (Adapted from Figure 3-2 (page 62), Figure 3-3 (page 71), and Figure 3-4 (page 72) in Nonaka and Takeuchi, *The Knowledge-Creating Company*, 1995.)

Apprenticeship programs exemplify tacit-to-tacit knowledge conversion where learning takes place through observation, imitation, and practice. Externalization creates an external representation of knowledge and is carried out through dialogue and collective reflection. In case studies, Nonaka and Takeuchi point to the use of metaphor in conversation and to cross-functional teams that engage in problem-solving sessions.

Combination is a process of systemizing concepts and reconfiguring explicit knowledge. This is the area most often associated with computer automation of business processes. Explicit knowledge is embedded in work flow systems, online business processes, and expert systems, as described in Stefik's 1995 *Introduction to Knowledge Systems* textbook.

> Knowledge, as the word is used for knowledge systems, refers to the **codified experience** of agents. By *codified,* we mean that the knowledge has been formulated, recorded, and made ready for use. … Experience must be articulated to become explicit knowledge. … Codified experience must be organized and generalized to guide future action. (Stefik, page 294)

Internalization is closely related to learning by doing.

> For explicit knowledge to become tacit, it helps if the knowledge is verbalized or diagrammed into documents, manuals, or oral stories. Documentation helps individuals internalize what they experienced. … In addition, documents or manuals facilitate the transfer of explicit knowledge to other people, helping them to experience the experiences of others indirectly. (Nonaka and Takeuchi, page 69)

Nonaka and Takeuchi characterize tacit knowledge as context-specific, but they are often ambiguous about the role of context-specificity for explicit knowledge. Our analysis of this point follows Stefik, who characterizes knowledge as always having a *degree* of generality whether it is tacit or explicit.

> Knowledge cannot be isolated from its creation or use. Experience involves agents in particular situations working on particular tasks with their own background assumptions. For example, a doctor seeing a patient with particular symptoms may quickly propose an explanation of why the patient is sick, knowing that there has recently been well-water contamination in the area. Such knowledge is particular to experiences of the time and place. (Stefik, page 295)

The implication of this point is that articulating or writing an expression of knowledge does not immediately make the knowledge portable to other places, times, or contexts. This point takes on particular relevance given computer networks, which can transport documents instantly from one part of a corporation (or the world) to another. Having a document does not necessarily imply that some particular knowledge can or will be used.

> The metaphor of knowledge as a "transportable substance" has long been used in informal English conversation. For example, books and teaching are said to "convey knowledge"; elders "pass along" knowledge to the next generation; knowledge and secret information sometimes "leak out." These metaphors reflect a truism. During education, knowledge becomes known by additional people. … However the process of spreading and knowing involves more than transportation. (Stefik, page 295, 296)

The sending agent must find a way to represent his meaning in a message. The receiving agent must find a way to construct a meaning from a message. Skill and context are needed at both ends. Figure D-3 illustrates this point, showing that although the bits of a message may be transmitted with perfect fidelity, the meaning behind the message may be missed entirely. Understanding does not come for free. We have to work to "make sense" of a message.

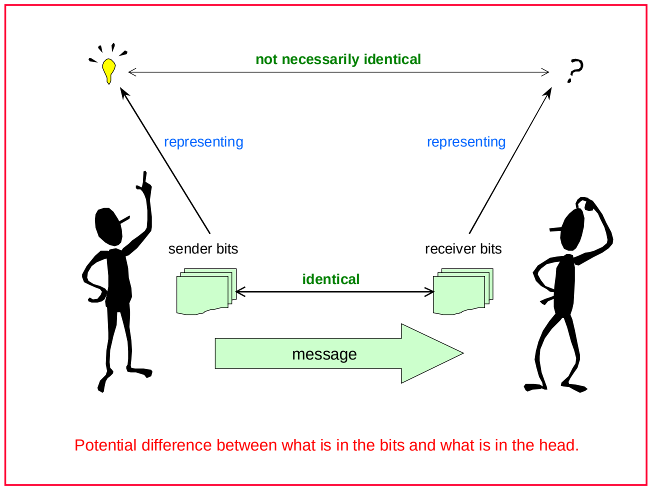

Figure D-3. Understanding and using knowledge requires more than receiving a document.

Michael Reddy calls the model in Figure D-3 the conduit metaphor. He uses it to explain the role of language in communication and to debunk overly simplistic models of knowledge sharing:

> Language seems rather to help one person to construct out of his own stock of mental stuff something like a replica, or copy, or someone else's thoughts – a replica which can be more or less accurate, depending on many factors. (Reddy, page 167)

People cope with the ambiguities and uncertainties of messages through the art of conversation. Many of the messages exchanged in a conversation are not so much about the content directly as they are about guiding the joint construction of meaning: "Is this what you meant?" "There's more than one way to understand that. Did you mean …?" These questions, called "speech acts" in discourse theory, enable participants to check whether they have constructed the intended meaning. Figure D-4 portrays parts of such a conversation.

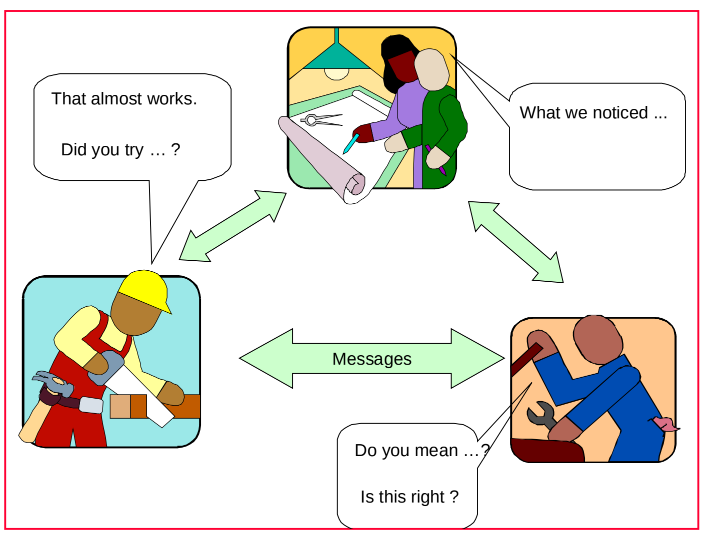

Figure D-4. Creating and sharing knowledge takes place in a working conversation, where people make sense of what others say in the context of shared work.

Another aspect of Figure D-4 is that participants in a conversation share a context. In Figure D-4, the context is the coordinated work. Doing the work involves sawing the wood, adjusting the bolt, consulting the plan, or whatever. The work could be about spreadsheets or customer accounts or anything else. The work constitutes an external reality in which the participants try things, see what happens, and share knowledge.

When people talk to each other, they often try to keep their speech simple. We often omit saying things that we believe the other person already knows or will infer from the context. For example as in Figure D-5, a speaker may say "turn the handle" without indicating whether to turn it to the left (counterclockwise) or the right (clockwise). In trying to follow the instructions, a listener may wonder which way to turn the handle. He could ask the speaker to clarify the instructions. If the listener asks for clarification, he mixes listening and speaking to make sense of the instructions.

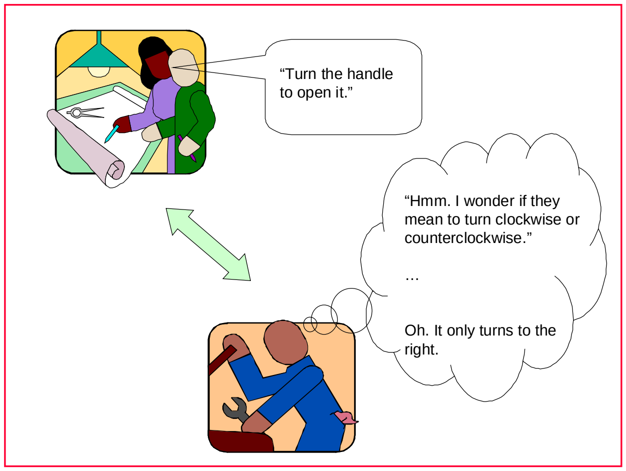

Figure D-5. Spoken instructions are sometimes incomplete or ambiguous. People use the environment as a resource for making sense of instructions.

Alternatively, if the listener experiments and tries to turn the handle, he may notice that it turns only one way. He may also look around and notice that there is an arrow on a sign above the handle that indicates which way to turn it for opening. In such ways the listener can use clues in the environment to resolve ambiguities in the instructions. If he experiments and tries turning the handle, he mixes listening and taking action to make sense. Conversations are always situated in an environment. Sensemaking in conversation often mixes the steps of listening, speaking, and acting.

Creating, Using and Sharing Knowledge

For Nonaka and Takeuchi, creating, sharing, and using knowledge in an organization invokes a spiral process involving knowledge conversion.

> … socialization … starts with building a "field" of interaction. This field facilitates the sharing of members' experiences and mental models. Second, the externalization mode is triggered by meaningful "dialogue or collective reflection," in which using appropriate metaphor or analogy helps team members to articulate hidden tacit knowledge that is otherwise hard to communicate. Third, the combination mode is triggered by "networking" newly created knowledge thereby crystallizing them into a new product, service, or managerial system. Finally, "learning by doing" triggers internalization. (Nonaka and Takeuchi, page 71)

This leads to the spiral in Figure D-1 which leads from socialization, to externalization, to combination, to internalization, and then back to socialization. Figure D-6 adds group size -- the what Nonaka and Takeuchi call the ontological dimension – to visualize the expansion of knowledge through an organization in a spiral of knowledge creation.

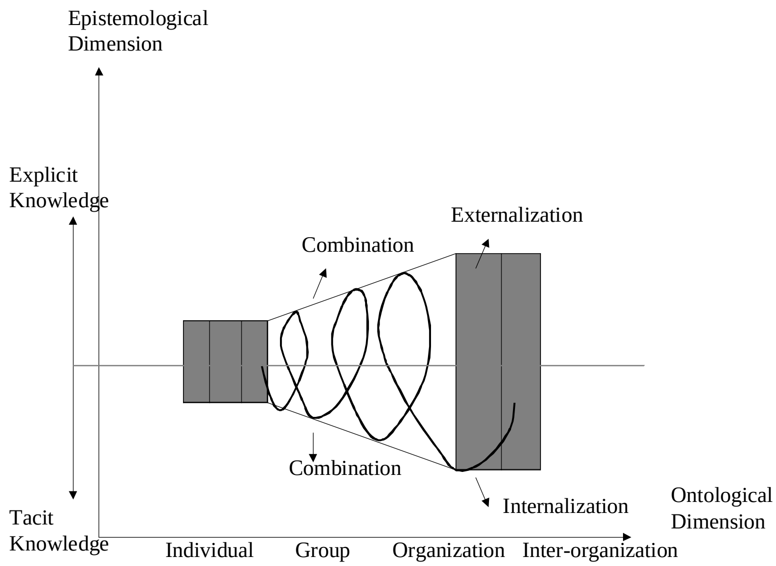

Figure D-6. Spiral of organizational knowledge creation. (From Figure 3-5 (page 73) in Nonaka and Takeuchi, *The Knowledge-Creating Company*, 1995.

The knowledge creation spiral provides a metaphor for a process of knowledge creation and use in an organization, but it also sets the stage for several questions. Does knowledge creation really have a starting place, characterized by the first quadrant in Nonaka and Takeuchi's diagram? How much knowledge should be shared in a company, or more specifically, what are the important patterns or profiles for sharing knowledge? Going back to the example of the baker and bread-kneading, is it really useful for a driver of a bread-delivery truck or the bread salesperson to know how bread dough should be twisted? More likely is the case that key knowledge of a special way of twisting and kneading dough should be kept as a company trade secret, held only by a few people.

Knowledge is not a physical substance or artifact that we can see directly and point to. Unable to see it directly, people often feel a little puzzled when we ask "how" knowledge is created or "where." In the context of the "seci" diagram, we can ask: Is tacit knowledge created in the socialization phase, when the apprentice learns from the teacher? Do we "create" explicit knowledge when we write something down in the externalization phase? Is that writing just the conversion of knowledge to an explicit written representation, or is that also an opportunity for creative combination of ideas and genuine new knowledge creation – like a scientist finally getting the equation right as he rearranges symbols on his scratch pad? Does knowledge creation take place in the conversations of small groups, as people listen to each other and sometimes even "creatively misunderstand" each other – leading to new ideas? In the end, we believe that knowledge can be created in all of these ways.

In the middle ages in the west, it was widely held that all ideas and knowledge came from God. This was a legal justification for the rule of kings and their role in regulating such things as the book printing industry. Kings claimed that power as deriving from their role as representing the power of God. This example shows that notions of knowledge and its creation can be bound up deeply in social and philosophical beliefs that differ from culture to culture. Our purpose here is not so much to explore or debunk any of those ideas. It is to understand knowledge in the context of modern enterprises and modern life how we might foster the creation and productive use of knowledge.

There are many factors that bear on knowledge creation. It is said that enough "quiet" time is needed for a creative person thinking. It is said that a certain amount of chaos is useful to destabilize habitual thought. It is said that play and informality are useful to encourage creative offerings by members of small groups. It is said that mixing is good to bring together perspectives of people with different backgrounds. It is said that voting and polling is good for combining the ideas and preferences of individual members in large groups. The question of how to foster the creation and use of knowledge is taking on practical importance as well as philosophical curiosity because the traditional methods are breaking down. Work is no longer just done by small groups working in one place and living in a small town. The world is changing in ways that change how we work and know each other – with globally-distributed corporations, virtual teams and emergent communities.

How many people work together in a company, or more specifically, what are the dynamics of the formation and dispersal of collaborative work groups in a company? How do work groups or larger organizations actually collaborate, create joint intentions, and share knowledge? These questions are about scale and change. In addressing these questions the next sections go beyond and extend the framework of Nonaka and Takeuchi.

Scale, Change, and Social Networks

The depiction of the ontological dimension of Figure D-6, from individual to group to organization to inter-organization, portrays a simple linearity. This portrayal is misleading in that it does not suggest that the process is affected by the size of an organization. Figure D-7 redraws the connections between individuals, work groups, and larger organizations to show how a typical tree3-like branching of hierarchical organizational structure leads to an exponential growth in the number of connected individuals.

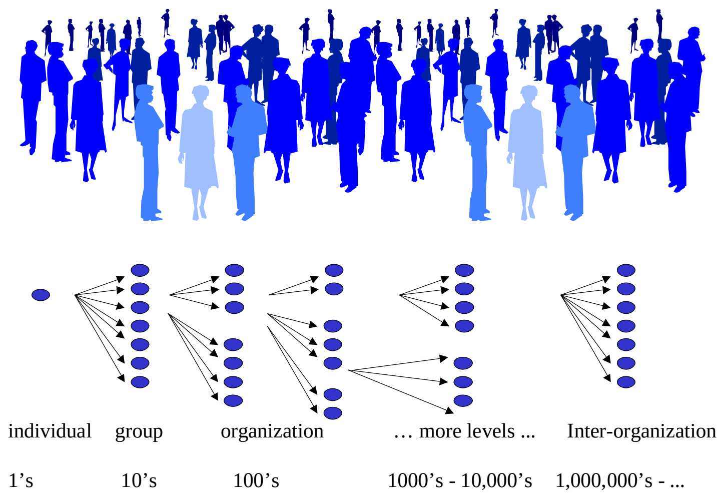

Figure D-7. Fan-out of relationships in organizations.

Organizations are not usually organized as strict trees – they often have shared groups and fuzzy boundaries. Indeed, in any large organization the institutional hierarchy is interlaced with connections among special interest groups and functional connections: strategy and planning groups in one part of an organization link up with their counterparts elsewhere; marketing and manufacturing groups do the same; special interest groups even form for groups that use the same software. Although people in an organization usually have one manager, they are also typically members of several informal groups. For this reason, the connectivity of groups and individuals in an organization is better represented as a network or graph.

It is not practical for everyone to know everyone else in a large organization. Political and business influence tends to be transmitted over chains of acquaintances. People tend to associate in social networks or social circles. These networks have certain mathematical characteristics and are called "small world" graphs. The name is derived from the comment "it's a small world, isn't it?" made by two new acquaintances who discover that they have an acquaintance in common. A small world network is characterized by a high degree of clustering as suggested by Figure D-8. For example, within a social network, no more than 10 or 12 links may be required to go from any person to any other person via the relationship "knows," where "knows" could be defined as "can recognize and be recognized by face and name."

How can a modern and evolving organization make sense of itself? The idea that all of the knowledge created in every small work group is shared with every one else is just silly. In a large organization undergoing change and turnover, even knowing how the cast of players is changing can be a challenge. People specialize. Some people become experts in different fields. Some people become experts about particular regions of the world, or about particular customers. An organization as a whole operates effectively when the threads of personal connection, conversation and shared knowledge enable it to act as a sufficiently coherent team.

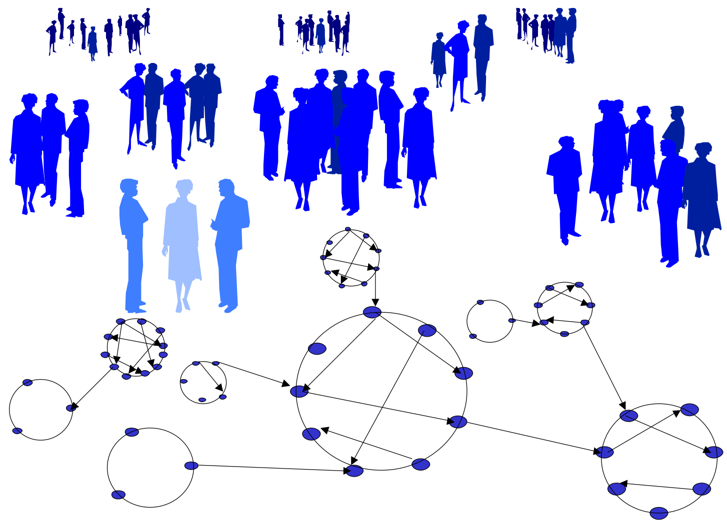

Figure D-8. Relationships in a small world – with small path lengths and a high degree of clustering.

This requirement relates to the old parable of several blind men around an elephant. One man touching the tail says that an elephant is like a snake. Another man touching a leg says an elephant is like a tree. Another touching the trunk says an elephant is like a hose. A corporation is like this – where some people are looking at technology, some are looking at markets, and different people are looking at different kinds of customers. Only instead of there being just seven men around an elephant, there may be tens of thousands of people around the world. No wonder that knowledge sharing is a serious challenge.

Communities of Practice

So far our discussion characterizes knowledge as residing either in individuals (if it is tacit) or in writings (if it is explicit). This view is consistent with psychology – perhaps best exemplified by the information processing psychology of Newell and Simon. Over the past few years, however, another perspective on just where knowledge may reside has emerged from the study of communities of practice. A basic notion is that much knowledge can be held by communities – in their activities, work places and practices.

> While knowledge is often thought to be the property of individuals, a great deal of knowledge is both produced and held collectively. Such knowledge is readily generated when people work together in the tightly knit groups known as “communities of practice.” As such work and such communities are a common feature of organizations, organizational knowledge is inevitably heavily social in character. (from Brown and Duguid, *Organizing Knowledge*, page 91)

As Brown and Duguid put it, the idea of knowledge embedded in practice augments but does not replace the idea of personal knowledge.

> What people have by virtue of membership in a community of practice, however, is not so much personal, modular knowledge as shared, partial knowledge. Individual and collective knowledge in this context bear on one another much like the parts of individual performers to a complete musical score, the lines of each actor to a movie script, or the roles of team members to the overall performance of a team and a game. Each player may know his or her part. But on its own, that part doesn’t make much sense. Alone it is significantly incomplete: it requires the ensemble to make sense of it. (from Brown and Duguid, *Organizing Knowledge*. page 96)

The notion of a "community of practice" comes to the fore in the writings of Jean Lave and Etienne Wenger. They noticed that work is a collective and cooperative venture. To understand knowledge in a community of practice it helps to differentiate between know-what and know-how. Know-what is about facts and to a significant degree this is something that people can carry around in their heads. Know-how is the ability to put know-what to work. Know-how is the product of experience. Since most work is a cooperative venture, most know-how is collective. Restated – much of know-how is knowledge that a group of people possess about how they accomplish something together. A group that shares know-how and sensemaking is called a "community of practice."

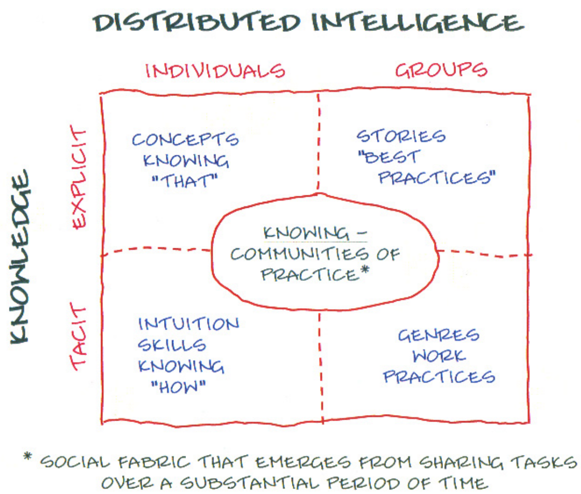

Figure D-9. Extending the notions of tacit and explit knowledge from individuals to groups.. (From Brown, J.S., *Learning, Working & Playing in the Digital Age: Creating Learning Ecologies*.)

Figure D-9 shows a way to think of the extension of the notions of tacit and explicit knowledge from individuals to groups. Explicit knowledge for an individual is knowing what or knowing that, but for a group explicit knowledge can include documentation for best practices. Tacit knowledge for an individual includes knowing how as well as developing skill and intuition. For a group, tacit knowledge is the collective know-how of work practices.

The Sensemaking Cycle

When two or three people work together, they share knowledge in conversation. Figure D-10 illustrates this process by which each party creates and shares knowledge using a sensemaking cycle. The steps of the sensemaking cycle are the familiar parts of conversation in which each person listens to the others, speaks to them, and creates an understanding of the knowledge and information while doing the work. Sense is made, not just conveyed. As in the Xerox marketing slogan "keep the conversation going," the sensemaking cycle represents a basic part of how people create and share knowledge.

One problem with the simple form of the "conversation metaphor" is coping with scale, that is, understanding how the nature of conversations changes when they involve larger numbers of people. Whenever we talk to many people at once, the simple notion of conversation inherently breaks down. If one person says something to a group in a small meeting, the same issues of incompleteness and ambiguity arise, but not everyone can ask questions at once. It gets confusing if everyone tries to talk at once. Conversation in a small group requires rules for turn-taking. If we have a question of the speaker, it often happens that someone else will ask our question before we do. With multiple people listening, speaking, and doing – sensemaking becomes collaborative.

.

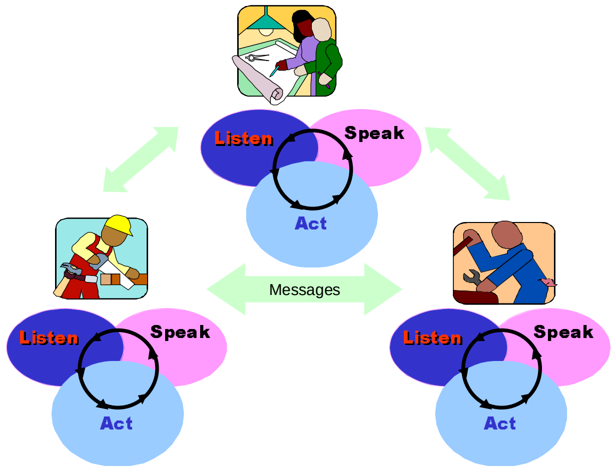

Figure D-10. People creating and sharing knowledge in a sensemaking cycle of conversation. Each person listens for knowledge and information from others, speaks to share knowledge and information, and tests understanding of the knowledge and information through action in a work context.

The term "fan-out" refers to the number of people that one person reaches if they speak or send a message. If one person speaks and a hundred people hear him, then the fan-out is one hundred. When we use technology for fan-out beyond the size of a small working group – such as a loudspeaker for a town meeting or a radio for an even larger audience – then the notion of simple conversation breaks down further. Consider the reaction of an audience after a speech. When a speech is over, there is often a buzz in the audience as people discuss among their neighbors what they heard and what it means. It is clearly impossible in a town meeting for everyone to address the speaker. People talk to their neighbors, comparing insights, trying to construct and repair their understandings of the meaning of the speech. In this example, the collaborative sensemaking spreads over the entire town meeting. Sometimes there is a vote or some other kind of collaborative activity. In a large group, there is a very good chance that people will leave the meeting with different and perhaps contrary understandings of what was said.

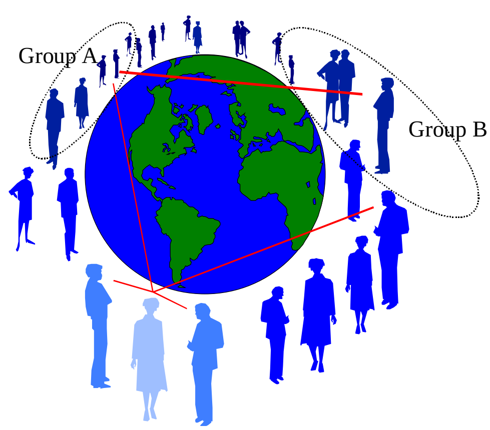

Figure D-11. Effective global corporations organize to facilitate an appropriate level of focused communications and information gathering among distributed individuals and groups.

How can a modern and evolving organization make sense of itself and its competitive world? Figure D-11 redraws the communicating groups of people, distributing them around the world. The people of an organization are both its sensors and its actors. Collectively people act as the "eyes" on different markets, different developments, different regions of the world, and different fields or specializations. They form a networked ecology of work groups and interests. More crucially, they form an ecology of communities of practice.

> … most formal organizations are not single communities of practice, but rather, hybrid groups of overlapping and interdependent communities. Such hybrid collectives represent another level in the complex process of knowledge creation. Intercommunal relations allow the organization to develop collective, coherent, synergist organizational knowledge out of the potentially separate, independent contributions of the individual communities. The outcome is what we think of as organizational knowledge, embracing not just organizational know-what but also organizational know-how. (Brown & Duguid, *Organizing Knowledge*, page 97)

Keeping organizational objectives in mind, how is it possible to have an organizational conversation? How is it possible to be effective, think globally and act locally?

Document-Oriented Sensemaking

Another problem with the simple version of the "conversation metaphor" is that it is limited to audio. In audio conversation we listen and we speak. For many centuries people have used other media to get beyond the limitations of audio conversations and oral tradition. Perhaps the oldest and most powerful non-audio medium is writing.

How does "conversation" work with writing? Anyone who has corresponded with a friend by writing letters knows that letter writing has some similarities to conversation. The pace of exchange is slower than with face-to-face audio conversation, even with e-mail. Nonetheless, operations for "making sense" must be present in conversation by letters. We may ask questions to clarify what is meant when a statement is ambiguous or incomplete. For written instructions, we may try things out in the environment to make sense of what is to be done. Figure D-12 extends the "listen-speak-act" cycle of conversational sensemaking to encompass variations for writing such as "read-write-act." Seeking media-independent terminology, we propose acquire, share, and act as the three general steps. As with conversational sensemaking, knowledge and meaning are created – not simply transported. The making of sense involves combinations of all three steps of the cycle.

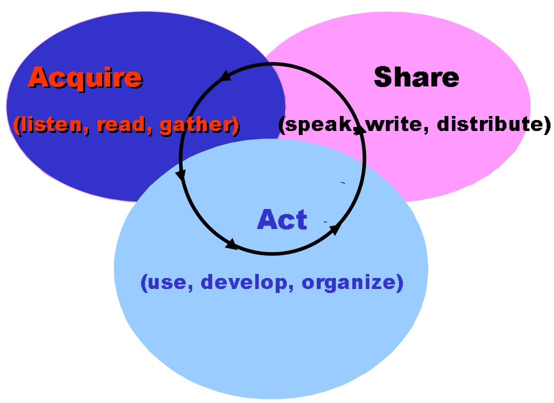

Figure D-12. The document-oriented analogy to the conversational sensemaking steps "listen-speak-act" are "read-write-act". In order to generalize beyond any particular medium, we propose the more generic steps "acquire-share-act".

As with audio conversations, written conversations take different forms when they involve large groups. Consider scientific journal publications. Even for scientists specializing in the related fields, making sense of journal articles can be a challenge. How do the steps of sensemaking – analogous to conversational speech acts – take place in the context of journal publication? One way to look at the process of peer review, journal editing, and copy editing is that these processes have been invented to reduce the misunderstandings that might otherwise arise from the ambiguities and incompleteness of casual writing. In some journals, the writing takes on a formal methodology where specific sections in every paper cover “materials,” “methods,” “earlier work,” “hypotheses,” “experimental design,” and conclusions.

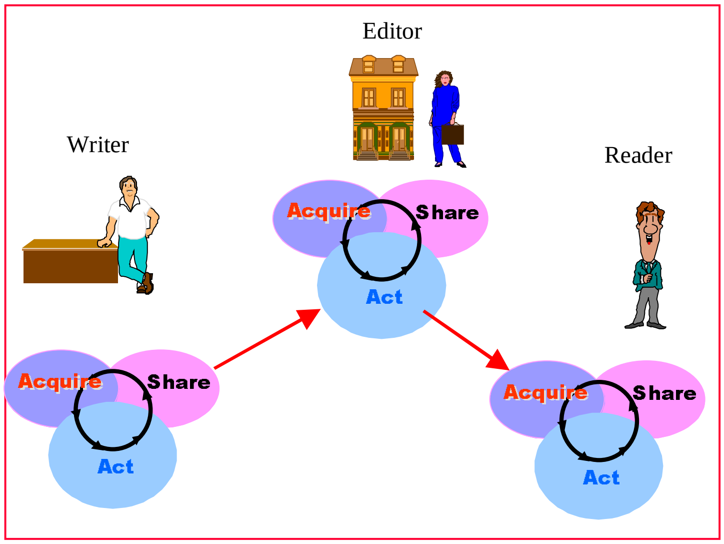

Figure D-13. When there is an intermediary in communication, it is typical for each party to operate primarily in one of the three steps of the sensemaking cycle. In this example, the writer primarily writes or shares, the editor primarily acts, and the reader primarily reads or acquires. When we classify the value proposition for situations like this, we can take a different point of view (and corresponding sensemaking step) according to which party we primarily have in mind.

Much of the literature of the scientific community reflects the expectation that each publication adds one more “brick” to the “knowledge edifice.” In that context, if a paper needs corrections or clarifications, this amounts to “taking apart” some of the knowledge structure of science. Although practical science is not really this idealized, the process of editorial judgement and peer review works to assure the enduring value of journal articles by formalizing conversation among informed experts prior to publication. In this way, the vetting process for journal publication uses a “quality control” process to reduce the need for conversational repair. This partial substitution – from conversational repair to quality control – is beneficial whenever the fan-out of document distribution is large.

In using the “acquire-share-act” conversational model in situations like publishing or vetting where there is an intermediary party, the point of view matters in characterizing the value proposition. As in the editing situation in Figure D-13, the writer is seen as primarily in the “share” role in the situation, the editor is primarily in the “act” role, and the reader is primarily in the “acquire” role. The value proposition for the writer is about how it improves his activity to share information; the value proposition for the reader, it is about how it improves his ability to acquire information. For another example, translation from one language to another has become increasingly important for global enterprises. As in the vetting example, the classification of the “translate” value proposition can depend on point of view. If someone translates a document in order to understand it, then translation relates to acquiring knowledge; if someone translates a document to make it more readable for others, then it relates to sharing knowledge; if someone routinely provides a translation service as an intermediary, then this relates to taking action as a regular job.

Thus, although we list the generic sensemaking steps as if they were distinct categories, there is often more than one reasonable way to place particular value propositions under them. Sometimes multiple categories are reasonable even when there is no intermediary. For example, “learning” can take place in acquiring (“What you say makes sense.”); it can take place in sharing (“I never understood it until I had to explain it.”); it can take place in action (“I finally figured it out when I had to do it.”). For another example, consider the placement for “organizing” information. Organizing can take place in acquiring (conceptual organization and note taking to aid comprehension); it can take place in sharing (organizing material for effective presentation); it can take place in action (organizing for efficient use). This overlapping and interacting nature of the sensemaking steps is fundamental to the interwoven threads of how knowledge is found, invented, shared and assimilated.

From the Spoken Word to the Written Word

Writing shifts the workings of conversations and sensemaking in other salient ways, especially the durability, accumulation, and complexity of messages.

The written medium stretches the time-period of conversations, sometimes dramatically and in ways that make conversations become essentially one way. Readers inevitably read documents from writers of the past, and writers inevitably write documents for readers of the future. In extreme cases, centuries may intervene between the acts of writing and reading. In such cases, the construction of meaning or sensemaking offers additional hurdles and challenges. Michael Reddy characterizes these challenges in a fairly extreme case, although a case not unfamiliar to those who have tried to find meaning in ancient writings.

> There are no ideas whatsoever in any libraries. All that is stored in any of these places are odd little patterns of marks … Now, if a human being comes along who is capable of using these marks or sounds as instructions, then this human being may assemble within his head some patterns of thought or feeling or perception … and if he does not have a rather full and flexible repertoire of thoughts and feelings to draw from, then it is not likely that he will reconstruct in his head anything that deserves to be called “his cultural heritage.” … All that is preserved in libraries is the mere opportunity to perform this reconstruction. “ (Reddy, page 17)

Because documents can be duplicated easily and persist, they can be read by large numbers of people. In contrast with an oral tradition in which the tale changes slightly each time it is told, documents are relatively immutable. Historians of the history of printing have noticed that these properties have contributed to the perception that relative to the spoken word, written documents often convey more authority – from “authoring” to “authority” as it were. Contracts are written. Famous treatises and reference works are written.

Another way in which the durability of documents changes the role of time in conversation when large numbers of readers are involved is that different readers can read a given document at different times. This property makes it easier for knowledge and information to propagate in writing because of convenience for the reader. Documents can be important for creating a sense of community. People can reasonably assume that others have read the same newspaper, and thereby, have a shared sense of the news and events taking place in a large community.

Another salient consequence of the durability of writing is that people and organizations build up of collections of writings. Reddy’s example refers to libraries. In a corporate context, writings accumulate as corporate records in many forms. As collections become large, this raises issues in acquiring knowledge that are not so apparent when we look at simple listening in conversation. Specifically, finding specific information in a large collection can become difficult. This has led historically to indexes and catalogs in libraries, and with online systems, to computer indexing and computer-mediated search. We will explore further aspects of this in the next section. For now we observe that in the generalized “acquire-share-act” sensemaking cycle, acquire involves searching and browsing in addition to listening and reading.

One insight about how people and organizations consume information is that they have an “information diet”. In rationing their reading or listening or viewing time, they consume a certain amount of international news, a certain amount of local news, a certain amount of gossip, a certain amount of technical information, a certain amount of economic information, and so on. Different people have different information diets. Some of the information comes to people unbidden – such as letters, e-mail, and advertisements in magazines, in newspapers on billboards, or on radio or television. When there is too much information people need strategies to browse and to filter out what is of low value. When information is hard to find people need strategies to search -– by asking other people or looking things up. Researchers who study the different strategies that people use to get information refer to “information foraging.” Knowledge is a key part of information foraging: knowing where to find information, how to evaluate it, and what to keep and what to discard.

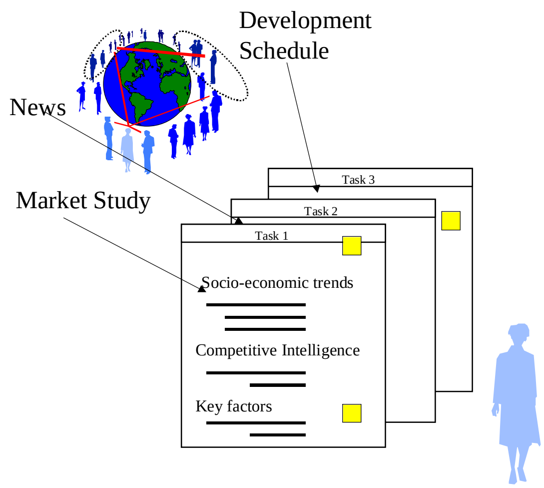

Figure D-14. The "acquire-share-act" sensemaking cycle involves finding information and putting it to work.

Even when we have information in hand, using it effectively can be a challenge. Much of the success in the competitive drive to efficiency in enterprises has been enabled by automating information handling. Examples range from automatic logging and analysis of data from point-of-sales terminals to the use of online forms and tracking in the re-engineering of corporate processes. Figure D-14 gives a visualization of this general theme for individual work. In this example, an individual has a variety of different tasks that they perform using information of many different kinds from many different sources and groups. Getting efficient means more than just getting all the information. It means putting information to work. The figure suggests that automation could be involved not only in gathering information about development schedules, or markets, or international news – but also putting that information to work in the right place. We call this just-in-time and just-in-place information.

How Writing Changes Sensemaking

Psychologists who study how people do knowledge or information work pay attention to the use of external representations, that is, to their writings and drawings on computer screens, paper tablets, or blackboards. The term external cognition refers to human information processing that combines internal cognition with perception and manipulation of external representations. We create, use, and share external representations in ways that augment our mental work.

The use of such external writing for understanding is another way in which writing changes the nature of sensemaking. Written words are tangible artifacts. We can move words and pictures around when we are working on something that is to big to "fit" in our heads. When we write on a blackboard or on the index card as part of the process of figuring something out, writing takes on a role different from conversation with another person. In effect, we have a conversation with ourselves in a process of invention and sense making.

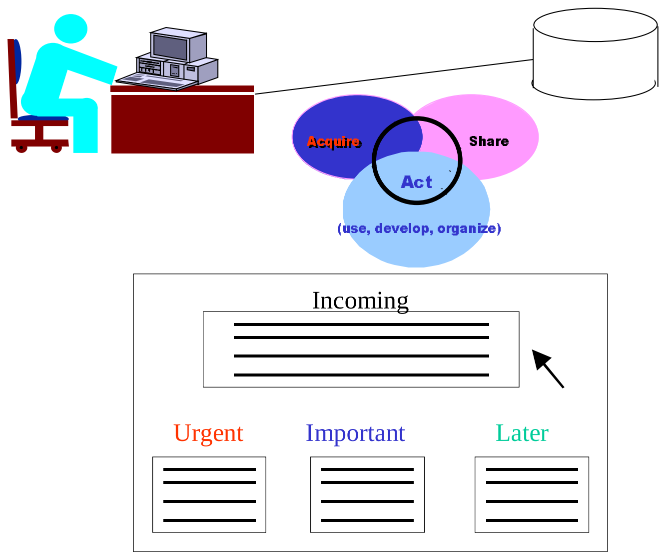

Figure D-15. External cognition. When the information task is too big to fit in their heads, sensemakers work with external representations. In this figure, the sensemaker categorizes incoming tasks into three categories. In this way, the "act" step in the sensemaking cycle takes on physicality as a sensemaker organizes the representations, augmenting his mental activity with physical activity to "make the sense."

Different kinds of external representations serve different kinds of tasks. For example, certain kinds of charts help us to see certain kinds of patterns. Pie charts make proportions of something easy to understand. Calendars make allocation of temporal events easier to manage. Many people use their e-mail systems to organize their activities – where the arrival of messages often signals something to do and filing a message marks the completion of the task.

An effective example of a designed external representation was made clear to some of us as members of a business planning team preparing a proposal to start a new business. The proposal was presented for review to a corporate oversight group. The financial parts of the plan were presented in spreadsheets according to a specified format. During our meeting with the oversight group, we were struck by the efficiency of the oversight group as they flipped quickly to particular pages and stepped through the columns of numbers. Their probing of our plan followed well practiced lines. Why do you believe you can hire people so rapidly during this quarter? Do these expense figures account for the staggering of employment start dates during the period? Why do you expect the income to climb so rapidly during this period? The formatted spreadsheet was a familiar external representation for them. It was designed to make certain kinds of information easy to see. The efficiency of their use depended on its grouping together exactly the information that they needed.

Sensemakers craft external representations for a task. In a typical example, a sensemaker gathers information from many places. The making of sense involves selecting and organizing information so that it can be used for making a decision. The external representations are used both for joint figuring and for communication. Initial scribbles and informal notes evolve into formal reports. This invites a useful twist in terminology -- referring to the external representation as the "sense." Viewed this way, sense is not just an internal understanding. In writing a report or crafting a representation, sensemakers literally *make the sense*.

The theme writing as part of knowledge work has another interesting consequence: work leaves a document trail. Many corporate processes are regulated in part by the "paperwork" that surrounds their activities. The "paper" work need not literally be paper, as the trend towards electronic and network-based processes has continued. One effect of such electronic record keeping is that the records become an accessible source of information about corporate activities. Harvesting such information -- gathering and understanding the records – still requires sensemaking. The opportunity for harvesting valuable knowledge and information from such records is the force driving the new interest in data mining and knowledge mining.

Communities of Practice and Social Capital

How can a modern and evolving organization make sense of itself and its competitive world? The groups and organizations in Figure D-16 – communities of practice – can be visualized as an organized federation for collective action. Organizing knowledge across enterprise subcommunities is the essential activity of management.

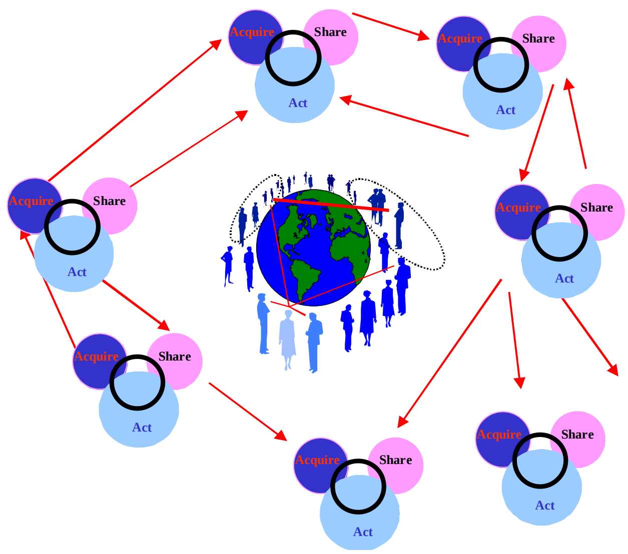

Figure D-16. Productive conversations. An organization at work involves groups of different sizes – communities of practice – collaborating in their work.

With explicit knowledge, we can easily imagine how the "publications" of one group could become the "readings" of another. But how is tacit knowledge shared, or more elusively, the know-how that is produced by communities of practice? The social embedding of know-how influences how it can move. Again, as Brown and Duguid put it:

> … knowledge moves differently *within* communities than it does *between* them. Within communities, knowledge is continuously embedded in practice and thus circulates easily. Members of a community implicitly share a sense of what practice is and what the standards for judgment are, and this supports the spread of knowledge. Without this sharing, the community disintegrates.
>
> Between communities, however, where by definition practice is no longer shared, the know-how, know-what, and warrantees embedded in practice must separate out for knowledge to circulate. These divisions become prominent and problematic. Different communities of practice have different standards, different ideas of what is significant, different priorities, and different evaluating criteria. What looks like a best practice in California may not turn out to be the best practice in Singapore (as HP found out). (Brown and Duguid, *Organizing Knowledge*, page 100-101)

When an organization is geographically dispersed, collaboration among groups has additional challenges. The values and methods of the groups will be different – because of their different environments and roles. Furthermore, geographical distance creates a greater sense of psychological distance.

The term social capital refers to features of social organization such as networks, norms, and social trust that facilitate coordination and cooperation for mutual benefit. The term social capital surfaced in the writings of sociologists in recent years, notably the writings of Robert Putnam who studied the relationship between civic participation and government effectiveness in Italy. As Putnam put it in an interview, "You tell me how many choral societies there are in an Italian region, and I will tell you plus or minus three days how long it will take you to get your health bills reimbursed by its regional government."

The same concept has been applied to the study of organizational effectiveness in companies. Data collected in large multinational companies consistently show that social capital – as measured in terms of social interactions and trust – had significant correlations with resource sharing and innovation. In effect, the quality of the human networks and level of trust has a huge effect on how well people work together. This has led to studies of what has been called "the organizational advantage." The overall thesis is that social capital facilitates the creation of new intellectual capital. Organizations are conducive to the development of high levels of social capital, and potentially have an advantage over markets in creating and sharing intellectual capital. Put more simply, a firm is a social community that specializes in the speed and efficiency by which it creates and transfers knowledge.

The workings of social capital and the organizational advantage depend on groups of people thinking of themselves as a networked community. When a company is small and located in a single location, it is easy for people to get to know each other and to share information. Knowing each other also motivates sharing. People recognize that the organization works better when people cooperate and there is a sense of shared fate. When a corporation becomes large or geographically distributed, it becomes much harder for everyone to know each other and to build up social capital. For example, many consulting companies have become interested in making it possible for consultants to share proposals, project reports, contracts, deals, and other documents. Such sharing makes it possible in some cases to reduce duplication of effort and to speed up project completion times.

However, an issue arises in a distributed organization. What will motivate people to share information in the first place? Presumably there is an advantage to the consultants who receive the reports or draft deals – because it speeds their work. But what is the pay-off for the individual consultants to contribute information? One much studied variant of this problem is the "tragedy of the commons" – where a common resource gets overused and depleted because everyone has an incentive to consume it but no one has a proportional incentive to sustain it. The situation also relates to a much-studied problem known as the "Prisoner's Dilemma," where there is a pay-off to two parties to cooperate, but an even bigger but asymmetric pay-off if one prisoner cooperates and the other cheats. The resolution to both dilemmas lies in social capital – which arises when all parties recognize that they will have many engagements – not just one. It is to their mutual benefit to cooperate over and over again because that maximizes their benefits over the long haul.

In a large organization where there is no habit of cooperation, there is a chicken-and-egg problem about how to begin to accumulate social capital. Different companies have approached this in different ways. There is a perceived benefit for the organization as a whole – so one approach is to just pay their people to share information. There can be a bootstrapping problem of achieving both a critical mass and a cultural shift which encourages people to share. Some companies deliberately recruit leaders to start the process going and work to make sure that people understand that they will be rewarded.

There is a big leap from recognizing the importance of social capital and the potential advantage of firms to assuming that corporations always succeed at exploiting that potential advantage. The forces of globalization that lead to geographic distribution in an organization also increase the psychological distance between individuals in virtual teams and the challenges of knowledge sharing between distant work groups. Cecere and Rosenberg identify several increased problems for distributed groups, including loss of trust, loss of urgency, loss of context, inappropriateness of electronic media for the subject matter communicated, and loss of motivation.

Moving from individuals to groups, the problems compound. How can individuals join a group, or sample its work? How can individuals establish reputations for themselves for exchanges with distant groups? How can groups build up reciprocal relationships, so that they proactively support each other in their work? The urgency of understanding communities of practice and social capital today is that the world is changing in ways that uproot the traditional processes for accruing social capital. If we ignore and fail to adequately support groups and communities in the evolving enterprise, we risk damaging the learning, innovation, and coherence that make enterprises effective.

Barriers to Acquiring Knowledge

What are the fundamental operations and barriers that the groups must face? We begin with barriers related to acquiring knowledge and information. What makes acquisition difficult? The right information can be hard to find. It can be buried in a pile of other information or lost in a flood of information. Somebody in the organization may have the information but we may have difficulty in identifying that person. The information may be in the system but scattered across many documents. To get it we may need to combine a little information from one document with information from several others. Each of these barriers also represents a potential opportunity.

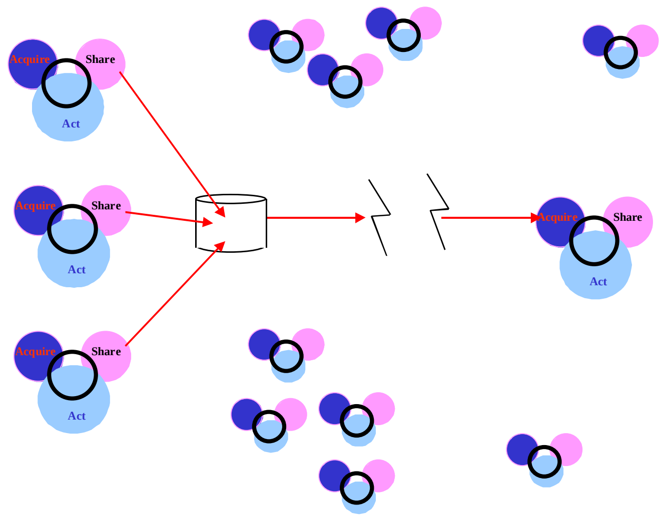

Figure D-17. For every organization, relevant information can come from many places: news sources, technology reports, competitive analyses, customer data, and so on. Attention management becomes an issue. Organizations need strategies for intelligently monitoring information from various places.

Figure D-18 shows the relationship between several barriers that arise in acquisition and corresponding value propositions that can be offered by knowledge-sharing information systems. In addition to the barriers for knowledge being hard to find, there are barriers for knowledge being hard to notice or hard to assimilate. There is an enormous untapped potential for knowledge-sharing systems to help us to get the gist of complex collections of information or to get oriented in a complex new area.

Acquiring (listening, reading, gathering knowledge and information)	
Barrier	**Value Proposition**
Hard to Find Lost in pile or flood. Lost in organization. Scattered across files and differ media	**Access rapidly** - Rapid access to information - Finding expertise - Fusing sources of information
Hard to Assimilate Too many details. Missing a framework. Importance not known	**Learn quickly** - Getting the gist - Getting oriented - Managing the flood - online apprenticeships - Distance learning
Hard to Notice Don't see what's happening Don't see what's changing. Importance not known	**Monitor effectively** - Finding what's new - Showing changes - Catching the buzz - Extracting intelligence - Profiling - Collaborative filtering

Figure D-18. Barriers to acquiring knowledge and information and corresponding value propositions.

To paraphrase Herb Simon, with a wealth of information there is a poverty of attention. A person cannot focus on everything at once. When there is enough information, it is not possible for a person to read it all – even working around the clock. This observation is grounded in the realization that people have limits in terms of how quickly they can find or digest information. When the amount of information is overwhelming, it helps to introduce it a little at a time – with processes for progressive disclosure and affordances for browsing. When we are pulling together information from different sources and trying to establish the relations between different elements of information, we create structures that facilitate understanding. The representations of the information progress from unstructured or chaotic forms to highly-structured forms. Knowledge-sharing technology can assist in this process by applying knowledge about information to the process.

Barriers to Sharing Knowledge

Our next general category of barriers relates to sharing knowledge and information. It is sometimes said that the net provides opportunities for using the big work of a few and the small work of many. However, orchestrating this work so that it is effective and effectively shared in an organization has many challenges.

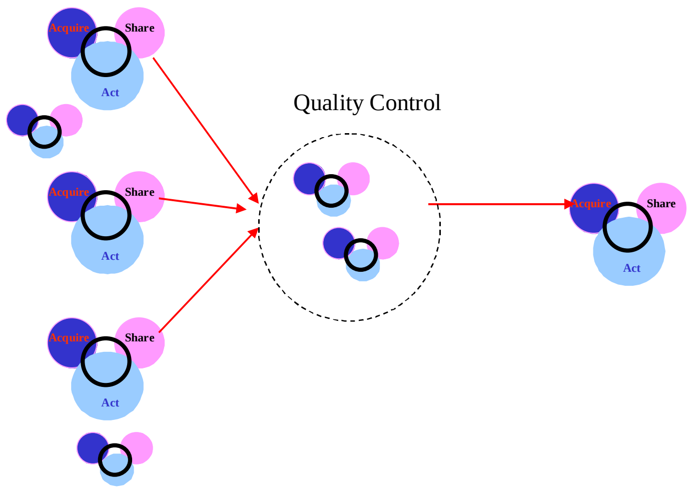

Figure D-19. One source of power in sensemaking is for organizations to solicit input from many different groups and individuals. Before such information is shared widely, it is reviewed for quality. This vetting approach is analogous to peer review in journal publication.

Figure D-20 shows the relationships between several barriers to knowledge and information sharing and corresponding value propositions for knowledge-sharing information systems. Information can be hard to share. People scattered across many areas and specializing in different concerns can often have different and isolated viewpoints. How can these viewpoints be collected or combined?

Communities left on their own can have blind spots caused by limitations of their own point of view. They can become stuck in ruts. Organizations as a whole can act to yoke together diverse communities of practice, challenging the limits of each community's beliefs. Means for connecting and supporting cross-functional teams are crucial to the success of such challenges and essential for effective joint work. From a pragmatic perspective, it is essential in an organization – and complicated in a geographically-distributed one – to find ways to build trust and relationships across groups that have a natural tendency to evolve distinct points of view. Diversity can be a strength – but communities of practice need means and motivation to evolve their practices to mutual advantage.

Another challenge to sharing knowledge and information by means of writings is that information can be of uneven or context-sensitive quality so that it is of high quality in one context but low quality in another. There can be mismatches between readers and writers. Warrants or reasons that carry weight for one community of practice may seem irrelevant to another.

Even when information is on the net, it can sometimes be in many places on the net creating headaches for keeping different versions of it coordinated and up-to-date. Finally, it can be difficult establishing connections and building community. How can we build relationships, social capital, and trust when many interactions in virtual teams and electronic communities are electronically mediated?

Sharing (speaking, writing, distributing knowledge and information)	
Barrier	**Value Proposition**
Hard to Share Isolated viewpoints Uneven quality of information	**Share easily** - Collecting thoughts - Vetting for quality

Poor writing or organization Mismatch between reader writer Redundant and stale	- Author expressively - Enter once
Hard to Coordinate Cross-department confusion Uncoordinated planning	**Mix creatively** - Wide-spectrum searches - Cross-functional teams - Sharing workspaces
Hard to Connect Unknown in the organization Communities not formed Knowledge not portable Hard to socialize and join Hard to develop trust	**Building community** - Knowledge maps - Latent communities - Diversity maps - Belonging by degrees - Establishing trust

Figure D-20. Barriers to sharing knowledge and information and corresponding value propositions.

A central theme in knowledge creation has to do with how organizations cross-fertilize their teams and manage the development of their organizational goals and plans. With globalization, the time-honored approach of just walking around to talk to people and see what's happening becomes less practical. Companies are turning more to technologies to support collaboration at a distance and the discovery of opportunities and needs through knowledge maps and other visualizations of distributed enterprise.

Barriers to Acting on Knowledge

Our last general category of barriers relates to the acting step in the sensemaking cycle, which includes using, developing, and organizing knowledge and information. Information sometimes comes in a form that is difficult or at least time-consuming to use. For example, e-mail is a major conduit of information for many people. Many business people receive many more e-mail messages than phone messages – and these message include things to do and information about things that a person is already working on. However, information in e-mail is all mixed together. Part of what makes working with e-mail so time-consuming is the work of collecting together information that needs to be used together, extracting information from e-mail, and moving it to other documents or programs. Information is harder to use if we have to organize it. Information can be hard to use if it is made excessively complicated by being repeated or when stale and outdated information is mixed up with newer and more accurate information.

Acting (developing, organizing, using knowledge and information)	
Barrier	**Value Proposition**
Hard to Use	**Use readily**
Mixed up	- Putting information in context
Not known at right time	- Just-in-time & just-in-place
Repeated and outdated	- Catching up
	- Coordinating action
Hard to Understand	**Understand deeply**
Rigid Filing.	- Fluid categories
Stale information	- What-now documents
Narrow viewpoint	- Multiple perspectives
Important and trivial matters mi together	- Mapping collections
	- translation
	- Documents for thought
Hard to Develop	**Innovate creatively**
Missing opportunities.	- Digital sandbox
Solving wrong problem.	- What's needed invention
Parochial viewpoint	- Collaborative invention

Figure D-21. Barriers to developing, organizing, and using knowledge and information and corresponding value propositions

Figure D-21 shows the relationships between several barriers to developing, organizing, and using knowledge and information and the corresponding value propositions for knowledge-sharing information systems. When information and documents are used for multiple purposes, then the organization and categories used for filing for one purpose are often not optimal for a second purpose. When the most complete document analyzing a situation or proposing a strategy ages, it can be difficult to determine what parts of it are still in alignment with how the situation has changed. When information is presented from one narrow viewpoint, it can be difficult to see how to apply parts of it to another. The value propositions on the right correspond to proposed kinds of solutions that incorporate knowledge that addresses the barriers. These include, for example, technologies for fluid approaches to categorization that admit multiple classifications for a given document, documents akin to spreadsheets that maintain links from their conclusions to underlying assumptions, and systems for organizing collections dynamically according to the purpose of the user.

Appendix E explores concrete examples of knowledge sharing value propositions that address the barriers in the sensemaking cycle.

Appendix E. Concrete Examples of Knowledge-Sharing Value Propositions

Concrete examples stimulate understanding. When we know clearly what a system or solution does, we can develop ideas about what it is good for, who would want it, and what technology is needed. In the following we introduce concrete examples of knowledge-sharing processes using the intermediate language of knowledge-sharing terms and value propositions.

	Content	**Context**	**Community**
Acquiring (listening, reading, gathering knowledge and information)			
Access rapidly	Providing an Information Portal Filter for relevance	Fusing sources	Finding expertise
Learn quickly	Getting the gist Managing the flood	Getting oriented	Distance learning online apprenticeships
Monitor effectively	Extracting intelligence	Finding what's new Web change detection Showing changes Topic maps	Catching the buzz Profiling Collaborative filtering Knowledge immersion
Sharing (speaking, writing, distributing knowledge and information)			
Share easily	Author expressively Translation in authoring	Enter once	Collecting thoughts Vetting for quality
Mix creatively	Wide-spectrum searches Rough documents	Cross-functional teams	Sharing workspaces
Build community	Knowledge maps Latent communities Web notes	Diversity maps	Belonging by degrees Establishing trust
Acting (developing, organizing, using knowledge and information)			
Use readily	Catching up Identify Document Redundancies	Putting information in context Just-in-time and just-in-place information	Coordinating action Reusing work
Understand deeply	Fluid categories Mapping collections Content analysis Information visualization Documents for understanding	What-now documents Translation for understanding	Multiple perspectives
Innovate creatively	Digital sandbox Knowledge discovery	What's-needed invention	Collaborative inventing

Figure E-1. A Taxonomy of knowledge-sharing value propositions

Figure E-1 presents a set of concrete examples for the knowledge-sharing value propositions, organized according to their leverage points in terms of the three C's – content, context, and community.

The value propositions vary in their state of readiness. Sometimes the discussion there are known solutions or products that realize the value proposition. Sometimes technology is emerging that may be employed in realizing the value proposition, but it has not been tried yet. Sometimes there are current solutions with shortcomings and ideas for addressing those. Sometimes the value proposition raises many questions and alternatives to be considered. We begin with value propositions offered by the current generation of information portals.

Accessing Rapidly

The phrase "access information rapidly" refers to a generic knowledge-sharing step. There are many ways to access information rapidly using indexing, search, and so on. Many regular business processes -- in manufacturing, finance, marketing, customer service and other industries -- are information and knowledge intensive. Speeding up access to the right information can accelerate the whole process.

The Gartner Group sees the trend exemplified by the rise of information portals in 1999 as spreading well beyond the early adopters by 2003.

> By 2003, the primary competitive differentiator of 70 percent of enterprises will be that they are enabled by IT (0.7 probability).
>
> The successful enterprise of 2003 will develop and implement an IT architecture that acts as a baseline for decision making and reflects the dynamic business and technology environment of 2003. (0.7 probability) (from *Enterprise 2003: The Technology-Enabled Enterprise*, January 29, 1998)

We now consider several concrete variations of this knowledge-sharing value proposition.

Providing an Information Portal

The rise of information portals in 1999 is based primarily on the value proposition of giving people rapid access to the information that they need. For example, here are the main value propositions offered by Autonomy, one of the new companies currently in the spotlight for information portals:

> *How much time do your employees waste each day on the following:*
>
> > *Surfing the web to stay up-to-speed on industry developments?*
> >
> > *Wading through the results of less-than-accurate search engines in search of a specific document or the answer to a particular question?*
> >
> > *Reinventing the wheel because they don't know that a colleague across the hall – or across the ocean – recently researched the same issue or developed a similar proposal?*
> >
> > *Tagging, hyperlinking or otherwise manually preparing documents or articles for publication online?*

Information portals are online interfaces that provide rapid access to online information. A portal provides tools for displaying selecting information in selected places in a window. If it enables searching over a corpus or web, it includes a web walker or a directory walker to traverse the set of documents, an indexing system to enable rapid search for terms, a matching algorithm to rank hits, and a browser to control the search and display results. Well-known accessible versions of this on the Internet include Altavista, Northern Light and Google. Personalized versions include myYahoo, myExcite, and others. Alternatively, the underpinnings can involve a system for classifying and possibly ranking the quality of documents. A well-known free public version of this approach is Yahoo, which employs people to search through sites on the web, cull out the most worthwhile ones, categorize them, and present their findings in an organized way.

As companies have gained experience with portals, the shortcomings of simple approaches have become more obvious and an appetite has grown for more effective approaches. The following examples explore variations on this value proposition.

Fusing Sources of Information

The vital information of the knowledge worker is typically fragmented into a variety of separate, non-interoperable (or awkwardly interoperable) sources. It is fragmented into separate buckets for e-mail, calendar events, presentations, reports, web documents, file system directories, and so on, and often multiple non-interoperable versions of each. In the absence of fusing, information must be found, viewed and accessed using different applications. Thus, web pages are accessed via a browser; e-mail is accessed via a mail reader; presentations, reports, and data are accessed via other programs. When searching for information, a user must typically perform separate searches with different programs that access different kinds of buckets. Look in here for e-mail; look here for presentations on topic X but there for presentations on topic Y; use the separate phone database for contact information.

Some current search engines and portals fuse separate buckets of fragmented information together so that a user can search through all sources of information at once. This addresses the fragmentation issue when there are multiple databases, multiple repositories, or multiple networks that would otherwise need to be searched separately. This saves search time and leads to greater success in finding information because the user does not need to remember where particular information is stored. The underlying technologies that support fusing of information include glue or interface software for accessing separate repositories and disparate representations of information, so-called "filters" for indexing documents that include format controls, and various online document services for such functions as format conversion.

Simply fusing information sources is not the end of the story. If all the information that a user was interested in was mixed in one big place, focussing would be difficult. Boundaries and persistent categories are not just obstacles – they can also be resources. An open question is how information systems can support the creation and use of boundaries and divisions in the "information soup" that are useful and supportive for a user at work.

At a lower level the fragmentation problem crosses multiple media. A digital "road warrior" receives messages on the phone, voice mail, faxes, sticky notes from someone who drops by, e-mail, and so on. Addressing fragmentation at this level involves a step of media fusion that precedes information fusion. Fusion at this level is being addressed by various companies, including the recently announced start-up, Onebox.

Finding Expertise

Sometimes finding the right information is best done by finding the document, but other times it is better to find the right person. Orbital Software is an example of a new company that is working this angle. Here is how they describe their offerings:

> Organik PersonaServer is a … personalization and profiling engine, designed to provide the basis of a new breed of people-aware applications. … PersonaServer builds dynamic, digital personas for the users and groups within your organization.

Whenever people use information systems, their activities leave a trail in the computer. Systems like the PersonServer help to capture that trail and to mine it for information about people and groups. For example, consider e-mail discussion lists. If someone subscribes to an e-mail discussion list about Lotus Notes, this suggests that they are interested in Notes. If they contribute to the list answering questions that others ask about Notes, that suggests that they have expertise. If they mainly access introductory documents about Lotus Notes, that suggests that they are beginners. If they spend much time with arcane and sophisticated documents about advanced Notes features, that suggests that they have a higher level of expertise. An integrated approach to monitoring use can combine bits of information about a person's activities and build up profiles of the expertise in an organization.

Sometimes the issue is not so much finding the right person as it is finding the right group. There may be no individual who has the needed knowledge. Sometimes what is needed is finding a community of people with whom to interact. Sometimes what is needed is to find out about the competencies or activities of a group of people. In this way, the "finding expertise" example blends into the community-building value propositions discussed later.

When an organization incorporates a "people finder" system in its business processes, it gets more than a phone directory and an organization chart. It gets a knowledge-sharing tool that the organization can use to understand itself. It gets a way for people to network with each other more effectively to find expertise and to carry out the company business.

Learning Quickly

The value proposition for speeding access to information focuses on accelerating routine work. However, in times of rapid change, much work is not routine. This group of examples of value propositions emphasize cases where the business situation is not stable and routine because it is undergoing change. When conditions change, organizations need to learn. Business forecasting and consulting organizations such as the Gartner Group see knowledge management technology as crucial for accelerating learning in the enterprise, especially for globally-distributed organizations:

> The successful enterprise of 2003 will learn from its successes and failures to better compete in the future (0.7 probability). Learning will combine the enterprise's human and systemic knowledge assets with information scanned from operations and the external environment. Learning will be a CSF (critical success factor) and a core competence of successful enterprises of 2003 (0,7 probability). (from *Enterprise 2003: The Technology-Enabled Enterprise*, January 29, 1998)

The nature of organizational work in a responsive and dynamic organization can involve less rote and more learning. How can technology support an organization in doing better at the time-intensive early part of the learning curve?

Getting the Gist

One strategy for dealing with the increasing pace of work and the growing volume of information is to allocate time differentially to various materials and to reduce the amount of time spent on individual items. This speed-up strategy is familiar in its pre-digital form. For example, when people shop for books at a bookstore, they browse and sample the pictures and writings in a book to get a feel for it. They may turn to the summary, jacket material, and the table of contents for an overview.

The same value propositions apply when people want to get the gist of digital information. Underlying technologies supporting this include summarization technology, which creates short abstracts of longer documents. Variations of this technology can be used for condensing videos of slide presentations – speeding up the presentation, showing only the introduction and conclusions, or skipping rapidly through key slides. Current techniques for gisting rely on statistical analysis of word frequency and cues about placement of sentences in documents.

There is room for creativity in addressing these value propositions. Current summarization techniques typically do not reflect the specific information needs or context of the user in the selection of sentences that summarize a document. (Inxight claims to support query-driven summaries, which is at least a start.) There is an opportunity for combining summarization technology with technology for capturing and using context to focus the summary.

Getting Oriented

Consider the example of helping a newcomer to an organization to become oriented. Many organizations have a complex collection of information resources that they use to organize their work activities. For concreteness in this example, assume that the information is organized in a portal or web site. Newcomers first encountering a complex web can find it difficult to get oriented. The idea is to give a newcomer a quick orientation to the site. On first encountering the web site, the newcomer (or "tourist") gets an experience analogous to an establishing shot in a movie. In an establishing shot a camera quickly pans over a scene to give a movie viewer a sense of orientation, and then focuses the camera on the point of action. An establishing shot in a movie can range from a simple a zoom to a lingering three or four minute pan of a scene. Similarly, a guided tour of information resources could be short or long, and would come to a point of focus.

Where does the "knowledge about knowledge" come from that defines a tour? Some of it could come from the people who create the information resources. For example, they could provide annotations about the purpose of different information resources. Another source of knowledge would be the behavioral patterns of people using the site. Part of the tour could be about what parts of the information are new and rapidly changing, and what parts seem to be more stable. Part of the tour could be a tour of the people in the organization and an overview of how they actually use the information. Part of a tour could be an introduction to different communities and work groups in an organization and how they work together. Part of the tour could be a summary of some of the important documents. Part of the tour could be an overview – akin to a table of contents. An important quality of a tour may be that it suggests a coherent narrative or story to help the tourist to understand and remember the material.

What is the underlying technology for creating a tour? The technology would support collecting, indexing, and summarizing information. It would support gathering information about how the information resources are used. It would support displaying the information cogently – providing short, long, and focused tours depending on the particular interests of the tourist. For example, user interfaces like the hyperbolic browser can give an overview. Animated graphics could zoom the view into particular information resources and then back again. If a user can select optional parts of a tour, details about the parts of a tour could appear – pop-up summaries or other data – animated when the user passes a mouse over them.

Managing the Flood

Metaphors abound about the experience of being overwhelmed by large quantities of information. The experience is called a flood of information, feast versus famine, information overload, and the information explosion. Herb Simon sums it up with his observation that a "wealth of information creates a poverty of attention."

Regular users of online search services are keenly aware of the very variable results of search. Sometimes a search leads to a flood of irrelevant information, dominated by irrelevant items based on a wrong word sense. For example, the word "stock" could refer either to warehoused materials or to news about recent changes in a share price. Time is wasted when a user has to filter through a collection dominated by documents that match a query term but which are irrelevant to the user's purpose.

Some portals provide a means for automatically focusing a search – such as what Scient Corporation calls "task-focused search" on their internal Scient Zone portal. In its simplest version, the portal asks the user to select a purpose for the query. For example, if the user's purpose is to find someone who can answer a question about Java, then the search focuses on employment records or biographic files for terms like "Java" near terms like "experience." If the purpose is to locate a Java-based system, a different corpus would be used. The Scient Zone also uses context information to guide the search of e-mail. Some e-mail discussion lists are about technical matters, some are about general business matters, and others are more specialized. The relevance of retrieved e-mail documents is enhanced by ordering the search to focus on e-mail distribution lists on the closest topics.

Potential resources for managing a flood of information are not just limited to the expertise and context of an individual. How could the experiences, priorities, evaluation criteria, or values of a work group be harnessed to assist in sorting through a flood of information?

In general, the idea is to provide technologies that reduce the flood of information by more accurately focusing what is delivered to correspond to the user's task, interests, needs, and preferences. Technologies relevant to managing large quantities of information include technologies for building up online ontologies, technologies for vetting documents to a panel of quality judges, and technologies for harvesting knowledge about organizational expertise. Also relevant – but more experimental – are technologies that extend syntactic matching to semantic matching.

Distance Learning and online Apprenticeships

The pace of change in technology is making it more difficult to stay current – even for professionals. How can the net be used to support new approaches to learning?

Needs for learning come in different shapes and sizes. At one end, universities are exploring new ways to support long-distance and lifelong learning activities around college-level courses. In this context, the network provides opportunities for digital delivery of lectures and instructional materials, as well as a means for field feedback and questions from students. Demonstration systems have shown how the network can support remote and time-shifting viewing of lectures, including capabilities for pausing, "rewinding," repeating portions, and asking questions. Students watching a lecture can type in questions that arise at particular points, even when the lecture is time-shifted. They can also see answers from the professors. The context of a question can be gauged by the point at which it arises in the lecture. To help find answers automatically, a question-matching facility can match the words and context of a question against ones that came up before. Students can also network together, sharing comments and questions about the material. These approaches leverage a community of students and faculty to enrich the lecture and learning experience over time.

In other situations, the most effective means for a motivated learner to acquire new knowledge is not a lecture, but rather engagement with a tutor. A tutor understands the domain in question and can structure a learning situation that presents the key concepts (content) in the appropriate context. In contrast, learning from text is more difficult, as it requires the learner to first grasp the organizational structure of the media, understand the terminology employed (which may be different from the terminology already familiar to the learner), find the particular topic of interest within the text, and assimilate the information presented. This information is rarely couched in the context of the user's situation, although examples may help bridge this gap.

A competent human tutor presents the material on an as-needed basis and in a useful format; for example, the meaning of any unfamiliar terminology may be clarified on the spot. The interaction need not require much time to be effective. For instance, coaching somebody on using a particular feature of PowerPoint may only require ten minutes. Unfortunately, requests for tutoring arise randomly, and can pose a significant burden on an experienced tutor.

Can the network and a community of practice be leveraged to support online tutoring? Network-based tools could in principle support online tutoring in several ways such as checking for cached answers, locating potential tutors, supporting a sharable work space, and accumulating educational snippets. Currently available collaboration tools already provide some of the functionality for shared workspaces and asynchronous communication. In principle, the network medium could usefully combine pre-stored materials with real-time tutoring. For example, a tutor could identify the student's confusion, give a quick answer, and then offer some pre-canned materials for further reference. A community of practice could use such tools to build up a "teaching corpus" to facilitate knowledge sharing with other communities.

Beyond improved communications, opportunities for research include tools for supporting a pool of experts who could give occasional online help. There is a considerable amount of work on the development of automated tutoring systems in the artificial intelligence community. While completely automated systems are unlikely to achieve human-level competence any time soon, techniques from this area may help in textual navigation and re-framing content for a particular context.

Monitoring Effectively

When the world is changing rapidly, it is not always possible to know from where important information may come. Routine work tends to have a center area for information – established news sources, certain fixed technologies, and certain fixed topics of relevance. Some changes, however, seem to come out of the blue. Rather than arising from the center of our focus, they come in from the periphery. For example, in state-of-the-art manufacturing businesses a new and potentially destabilizing technology could come from many different sciences and from almost any part of the world.

According to *Enterprise 2003: The Technology-Enable Enterprise* from the Gartner Group, an ability to monitor for changes and react appropriately will be increasingly characteristic of the competitive enterprise.

> The competitive landscape of 2003 will be littered with the rotting carcasses and parched bones of competitors too slow to seize opportunity and too slow to recognize challenge. A new form of competitor will emerge that continuously scans the environment and its own internal processes, sensing changes that signal opportunities or challenges. This enterprise uses its knowledge assets to learn from that which has been sensed, and innovates new products, services, channels and processes. It then mutates rapidly to bring its innovation to market-seizing opportunities and repelling challenges. …
>
> The successful enterprise of 2003 will be aware of its environment, continuously scanning for changes that signal opportunities and challenges (0.7 probability). It will constantly explore its relationships to detect changes in the condition and activities of customers, competitors, complementors and suppliers. Equally important, the successful enterprise of 2003 will scan to sense information about consumers who are not customers and alternative suppliers to those being used (0.7 probability). The successful enterprise of 2003 will scan its own operations and those of its value-chain partners to be aware of changes (e.g., quantities, quality and times) in processes and stocks (0.7 probability). (from *Enterprise 2003: The Technology-Enabled Enterprise*, January 29, 1998)

Knowledge workers create strategies for gathering information – where they focus much of their attention on the center but also find some way to monitor for important changes that come in from the periphery. This group of value propositions is about monitoring the periphery.

Presenting What's New

It is no accident that the word "news" refers to the latest, *new* information. A big part of staying informed is keeping track of recent events. News services specialize – so that there is international news, sports news, finance news, start-up news, and so on. The value proposition of any news service is that it provides the latest information in a topic area. In an enterprise setting, the relevant news is not just the general public news. What's new in a competitor's products? What's happening with our customers? What's happening in other parts of the enterprise? What technological advances may create new opportunity or new competition? Still the value proposition is the same.

A prerequisite for tracking what's new is logging new information. When information comes from news-gathering organizations, the responsibility for collecting the news is delegated to that organization. For information accruing in the files and information systems of an enterprise or on the web pages of competitors and technology providers on the net, it is quite possible that new information is entered but that it is not noticed by or distributed to people who need it. Level three technologies for processing news need to be integrated with lower-level technologies for noticing and collecting new information.

Given that new information is noticed, collected, and indexed by time, there are different ways to access it. Library search systems enable specifying dates in searching for information. This illustrates a "pull" approach to finding information where a user requests information and limits the search to items within a given time period. However, the efficiency-minded user does not want to search through files of information to pull out what's new. A variation on this is a display of items over time, such as a time-line display of e-mail from different people on a given topic. News can also be distributed by "push" technology. One approach using push technology allows clients to subscribe to topics so that they receive e-mail on a topic whenever news becomes available. Subscription technology is available within many information portals.

In general there is much opportunity to augment the systems used by knowledge workers to provide more automatic and effective ways of gathering and summarizing potent and relevant new information. The relevant technology includes agents and polling techniques that monitor the web or file systems for changes, technologies for detecting which changes are relevant and important, technologies for mining and summarizing information about the change, technologies for classifying the urgency of the news, technologies for fusing data from multiple sources and for vetting or interpreting the changes, and technologies and for routing it to relevant parties. A graceful blending of push and pull technologies could be employed so that a user would get a single note that there have been changes – rather than a slew of notes about each change – and a simple way to get a summary of the changes that is specialized to what has changed. The summary could omit changes that have been superceded by subsequent changes, could summarize who is making the changes, could cluster the changes to show their relationship, and could prioritize the report to reflect the specific interests of the user.

To be concrete, consider how web sites are often designed to present a particular organization of information. However, a user wanting to see "what's new" on the site may have to search through many pages – and even then, it may not be easy to determine what's new and what's old or even stale.

Finding out what's new is relevant to understanding organizations as well as people. It may be just as interesting to ask "who's new" as "what's new," or to get a perspective about how an organization or its activities are changing. These are areas where the community-oriented value propositions come into play.

Environmental Knowledge Immersion

A complementary notion for presenting what's new involves using environmental channels to present peripheral information. When people become fully immersed in something, they often know things unconsciously that they don't consciously realize. In effect, they develop subtle or intuitive beliefs.

Although the most powerful learning comes from sustained practice, immersion itself can arise even from peripheral cues. Advertisers understand this very well, in their use of billboards and other spots to build a sense of awareness and associations for a brand. Messages to promote awareness are important in non-advertising contexts too. Modern life is busy and many things compete to be noticed. As individuals, we routinely manage our attention. To promote awareness of various issues, organizations distribute literature, make posters, and send e-mail about such things as events, news, and policies. Such messages are akin to "public postings." Not all announcements and postings in an organization are about corporate or management initiatives. In principle, people with different interests throughout an organization use email, flyers, and posters to bring attention to certain things.

It is possible to create a working or living environment in which the environment itself is a vehicle for creating context and a sense of immersion. For example, displays (such as plasma panels and large LCDs) could be mounted in strategic places in the hallways and public spaces of a building – playing out elements of news and information. For another example, the header sheets of documents sent to printers can carry messages. Screen savers on desktops through a building could carry messages.

The use of such environmental channels raises both social and technology issues. Where should displays be placed? When do people just tune out messages? What kinds of messages are effective? From a technology perspective, there are opportunities to personalize messages. For example, displays in a public setting could be designed to sense the identities of people near by and to key the messages to the people. They could show different messages when visitors are near by. They could be made to be interactive, so that a person could walk up to an interactive message board and request more detailed information about something that just went by or request information on a topic.

How can an organization set policies about what gets aired? Policies could take into account the genre of the message, the source of the message, the timeliness of the message, the set of people near by, whether a fire drill or other emergency is in progress, and so on. Where are best places to put message displays? Should there be message-free zones?

Showing Changes in Context

The previous example focused on the use of push technology to focus attention on new information. It can be helpful to show new information in a context that highlights what has changed. Conventions for such highlighting have been developed over time to indicate what portions of a complicated document have changed – with cross-outs and underlines and so on. These notations indicate what words were deleted, what sections stayed the same, and what has been added or modified.

This same basic idea – using automatic formatting to indicate change -- could be used in other settings. Suppose that someone subscribes for alerts whenever the “accounting rules” of a company change or whenever new personnel are hired. Instead of just getting an e-mail message that someone was just hired into a group, or even just getting a link to a web page of the current members of a group, a system could show indicate the changes in assignments. A further variation – useful for understanding trends -- is to give an overview of recent changes. In the personnel example, one could create an animated movie to rapidly review the recent hiring or transition events in the group as a “movie” of changes or as a storyboard that shows how and why an organization is evolving over time.

Catching the Buzz

Looking at *recent* documents (“news”) is just one strategy for selecting what to read. Another is finding out what documents or topics are interesting to other people. Consider the online analog to joining a conversation down the hall when it sounds like people are excited about something. On the net there are many people and several ways to find the “buzz.” One idea is to keep track of the most popular sites (or documents) that people are reading. Another idea is to keep track of the sites that important sites link to. This is more of an “info-mediary” idea, where some sites stand out as opinion leaders. A third variation on the theme is used in recommendation systems to focus the search for buzz to topics of interest to delineated subcommunities. A user answers a number of questions, allowing the recommendation system to create and classify or cluster the user’s profile. The user’s interests (favorite movies, favorite books, etc.) are then tabulated so that when someone asks for a recommendation, the system can base its recommendations on the choices made by other users with a similar profile.

One way to look at the differences between these approaches is as variations on measures for popularity and authority. These variations – from info-mediary approaches to profile-based recommendation systems – illustrate the kinds of information that can be built up from usage patterns, reference patterns, and other sources. The underlying technology includes statistical modeling techniques as well as usage monitoring technology. In an enterprise setting, the inherent utility of this information is enhanced when it is combined and specialized by business context.

Another kind of “buzz” to catch is not just knowledge about what information people are using, but also knowledge about what people are doing. What are the hot activities in an organization? Where are groups re-aligning their activities? What new opportunities are getting people excited? This is an area where this value proposition relates to the community-oriented value propositions.

Extracting Intelligence

The basic value proposition of monitoring effectively is looking at the periphery. Each solution for monitoring effectively is an antidote for when a company is so focused on some immediate business that it is blind to and surprised by events coming from the side. A periphery is much bigger than a center, and therefore much more difficult to focus. One difficulty with monitoring raw data from the periphery is that the effort to digest the information can be overwhelming. This creates a premium for automatic processing to extract and organize the information for easier assimilation.

To be concrete, suppose that we want to benchmark a corporate research activity. For concreteness in the example, suppose that we want to benchmark a group doing research on a specific topic like "image retrieval" or "digital libraries," or "blue lasers". A human investigator might proceed by looking for conference proceedings in this area. She might gather up the various papers in the field, note the different topic areas that are currently hot, and note which ones the benchmarked group is publishing in. She could check through the lists of authors on papers and make aggregated lists of researchers in the area as well as centers of research activity. She might consult a citation index to get a sense of which papers are most cited, and use this as a quality measure for different researchers and their groups. She might then go to the web and search through the web pages of the identified sites, gathering information about the number of current and recent personnel, university connections, and funding. Finally, she would organize her findings in a short summary and tables that enable ready comparison of the benchmarked research group to competitive research groups at other companies or in academic settings.

To what extent could this benchmarking process be automated, running on computer cycles rather than investigator "cycles"? To what extent could it be generalized to cover product comparisons or scanning for new or destabilizing technologies? In general, a data mining program capable of a specialized search and report is highly knowledge intensive. There is a power-versus-generality trade-off in the creation of such systems. Computer programs as powerful as the investigator example above stretch the state of the art.

The commercial state of the art for intelligent data mining is exemplified by comparison shopper programs. For example, the e-compare shopping network (www.ecompare.com) and the ACSES shopper (http://www.acses.com/) both compare price and shipping information for books, music, office supplies, toys, golf equipment and other goods from multiple vendors on the web. These services and others use the net to connect to online book sellers, extract price and shipping cost information from the web pages, and then present the aggregated results from multiple online stores in a table for ready comparison and purchase. Another example of a comparison shopper is the Microsoft Network shopper for cars (http://carpoint.msn.com/Compare/Dual).

These examples suggest that there is a value proposition for intelligent data mining. Intelligent data mining systems will have a growing role in helping organizations to maintain a peripheral awareness of events that leave a trail in online sources. However, there are tradeoffs between generality and power in these developing systems, and the opportunity points in commercially-viable systems have not been deeply explored yet.

Sharing Easily

The previous value propositions for knowledge sharing have focused on finding and retrieving information. The flip side of knowledge sharing is in putting information into information systems in the first place. This group of value propositions is about solutions that simplify that part of the work.

In a Gartner Group survey about drivers for knowledge management initiatives, the highest ranked purpose (76 percent) was to improve knowledge and information sharing across operating units. Gartner Group focuses on information systems that integrate multiple information sources with each other and with the work of the organization. They characterize these systems as "polycentric" – meaning that they provide support for many different constituencies. Finally they speak of the value of making knowledge maps available to an organization.

> Through 2001, construction of a knowledge map will be a CSF (critical success factor) for an enterprise-level KM program. (0.8 probability)
>
> Integration will no longer be the orphan project assigned to AD (application development) and infrastructure staff as they become free from other projects. By 2003,

> application integration will become an IS business unit in 70 percent of enterprises (0.7 probability); and applications archeology will be a skill in high demand.
>
> Because of the emergence of knowledge architectures and business intelligence, by 1998 70 percent of enterprises will require polycentric application designs (0.7 probability) (from *Enterprise 2003: The Technology-Enabled Enterprise*, January 29, 1998)

Collecting thoughts

The web can be used to tap the hard work of the few as well as the light work of the many. In a networked knowledge-sharing context, a challenge is to find easy ways to collect and present small nuggets of information from many people in an organization. This value proposition is about collecting thoughts and opinions from the many.

One variant of this idea developed within Xerox research is "shared bookmarks" or GroupFire. This value proposition is that the useful bookmarks of discovered and valued web sites can be published and shared by a community. The underlying technologies include level-two technologies for databases, privacy, and sharing, combined with level-three technologies for clustering and visualizing sets of bookmarks. In an e-commerce setting, GroupFIre has access both to the bookmarks that people use and how often they use them. This makes the underlying data valuable for discovering latent communities (by measuring shared interest) and buzz (by measuring popularity). In the context of an intranet, such data could play a role in discovering and supporting interactions between communities of practice.

Online recommendation systems are another variant of the "work of the many" idea. Recommendation systems create profiles of related interests as well as ranking systems for judging the quality of online sources.

During the writing of this white paper, several companies for web notes either went public or announced products. There are several variations of this idea. An underlying theme is that web pages provide a sense of place in cyberspace, in effect, a place to leave notes or write graffiti on the wall. Companies like Third Voice, hypernix, uTOK, and other offer facilities for people to leave notes that appear on web pages when they are viewed using the appropriate software. It is also possible to search for web pages that have certain kinds of comments on them. Such web notes provide a vehicle for capturing the "small work of the many" in the context of the web pages that they are accessing.

A key element in systems designed to harvest the work of the many is that the overhead for anyone contributing information must be very small.

Vetting for quality

An inherent issue in harvesting the work of the many is that people tend to value some people's opinions more than others because some people bring more thought or expertise to their opinions than others. This basic observation underlies professional publishing in journals in all scientific and medical fields. Journal publishing is organized in a process of editorial and peer review. This institutional process separates the wheat from the chaff, and – through a process of professional review and dialog – incrementally improves the quality of what gets published.

The Eureka system within Xerox offers this sort of value proposition in an enterprise setting, Eureka collects so-called "tips" about fixing Xerox copiers and printers. Tips are intended to be an authoritative and timely source of service knowledge that augments product release information. Tips can be submitted by any technician. Before they are posted in the general database available to all technicians, they are vetted for quality control by a board of experts. From a perspective of knowledge sharing, this arrangement combines the strength of the many for discovering, identifying, and articulating tips with the strength of the expert few for assuring the quality of what is made available.

There is room for more experiments in processes for assuring the quality of online information. When someone joins a group, how do they gain stature and credibility in the community? How can peer review be done quickly? How do groups control or manage jargon? How do groups manage the discarding of stale data as well as the publishing of new data?

Authoring expressively

Word processing tools in the past few years have incorporated features such as spelling checkers and grammar checkers intended to help people to write more clearly. These tools have focused on syntax and barely scratch the surface of what is possible. Furthermore, they are mainly oriented to producing passive, non-interactive "documents" that evolved before the computational capabilities of the net became so widely available.

One direction for increased support in creating and organizing information is a shift from syntax to semantics. Information does not spring from nowhere into the mind of a writer; nor does it immediately emerge from chaos and take on a fully-structured shape. More realistically, writing derives at least in part from other writings or sources. Structure emerges incrementally, as writers grapple with gaining clarity and mastery over the ideas, and as they grapple with issues of teaching and rhetoric such as presentation, order, dependencies, and disclosure.

If a writing system included an information retrieval system for identifying and ranking sources, then it could help a writer both to find information of relevance and to keep track of how it was used. For example, it could maintain a network relating arguments and conclusions in a document to underlying source. It could remind a writer if certain critical sources were found and not used. It could help a writer keep a log of unfinished changes and rhetorical threads. Furthermore, this net of "support links" need not be discarded when a document is "finished." It could be kept available to the writer for subsequent rewriting and updating; it could also be made available as part of an extended "what-now" document for readers who want to easily delve into the sources and reasoning that underlie the outer presentation of the document.

The new interactive technologies of computing and networking make it possible to create documents with presentation capabilities beyond what is possible for conventional paper reports. Even casual users of the net are familiar with the range of multi-media mixing of sounds, animation, and video on web pages and with the basic hypertext linking structure of the net. Even for documents that are mainly textual, new ways of enhancing the reading experience are emerging. For example, fluid user interfaces add an animated capability for mediating the progressive disclosure of information. A reader can pass a cursor over a portion of the text causing footnotes, marginal notes, or inter-line notations to gracefully and temporarily appear with the text. Such affordances change the reading experience – enabling a reader to see notes in context – understanding them more easily and less disruptively.

Another property of conventional documents is that they assume that one size fits all. When document contents can be computed, it becomes possible to alter the selection and presentation to fit the needs of the reader. Such flexibility becomes more important when information is being shared by a wide range of people in an organization. This is an area which relates to the community-oriented value propositions.

There is much room for research in this area. How can choices about content and presentation be influenced by the knowledge of the reader and the writer? How can the creation of dynamic documents be made easier for the writer? How can sensemaking processes be incorporated to support both a reader and a writer? How can information systems support collaboration and latent collaboration in an organization that tries to make sense of itself and its environment?

Translation in Authoring

As mentioned above, word-processing tools have barely scratched the surface of what is possible. While spell-checkers and general-purpose grammar checkers can be helpful, translation-related authoring tools can further increase the possibility that what is written is accessible and comprehensible.

In a later section we discuss translation mechanisms per se. But a major barrier to fully automated, idiomatic, translation is the resolution of word-sense and syntactic ambiguities remaining after source analysis. One way to deal with that problem is to introduce translation technology at the authoring stage, to identify aspects of a document that a translation system would find ambiguous, and to interact with the author to resolve those ambiguities. The results can be reflected as rephrasing and/or recorded as annotation.

While this might be a very tedious process for a long document, the difficulty can be alleviated by use of user-driven methods. For example, a system can use carefully designed displays to express its "best guess" resolution of difficulties; synonym glosses can show word-sense assumptions, and parentheses can express assumed phrase attachments. Also, an author can decide to devote such attention only to abstracts.

An example of this kind of approach which is less specialized to a particular translation system is interactive reduction of a document to a controlled language, that is, a subset of the source language that is carefully designed to exclude many possibilities for ambiguity.

Such rewriting can help writers expand the audience for a document beyond readers of the source language. Other tools can help authors writing in a second language to produce more readable documents. This situation is increasingly common; the "lingua-franca" of the web and international communication is English, and the web itself is probably increasing use of that language. Easily accessible on line translation dictionaries, with lots of examples, are the most important tools here. Beyond that, "writing checkers" for second language speakers can look for anticipated kinds of grammar errors (for example, problems of prepositions, missing articles) and give useful advice. Also, checking of lexical co-occurrence is important, for example, "high mountain" vs. "tall building", as well as checking for the presence of "false friends", words which appear cognate but are not.

Finally, we can interpret the notion of language broadly to include usages local to a technical or professional community, and consider facilities to help an author to make a text comprehensible beyond their home community. For example, a technical author may want help in writing for a non-technical audience. A tool could help by flagging terminology which would not be comprehensible to that audience, based on a corpus of terms drawn from a periodical intended for a general readership. And even if the intended audience is a very specific one, similar checks could be made over abstracts intended to provide the gist of the material for other readers.

Enter once

In an organizational setting, groups are often required to report on their output and productivity. For example, in a research setting, research groups report on invention proposals, patent filings, projects and deliverables to business groups, customer engagements, papers published, talks given, and so on. This information tends to be reported in periodic reports, on web pages, and in other ways.

For concreteness in this example, consider the task of reporting on current publications via web pages. It is very common for individuals to create web pages that list their publications, and often, to provide a means for downloading electronic versions of the documents in various formats. At the same time, it is typical for workgroups – areas in a research center – to maintain aggregated lists of the publications for all of their members, to give an indication of the scope of their research interests and the productivity and professionalism of their members. The same notion of aggregating lists of publications takes place one level up – in laboratories. At this level, however, the work of maintaining an accurate list of publications for forty to fifty people is already a substantial chore. At the level of the research center with several laboratories, the task of maintaining an accurate web page of publications is enough of a chore that it is not currently done in the absence of dedicated staff.

The task of maintaining a list of publications is more than just creating the web pages. Suppose that the data about a publication changes or an error is made on the listings of publications at any level. Part of the job is in maintaining consistency across all of the different aggregations of pages.

From a perspective of knowledge sharing, there are opportunities to radically simplify this entire process. Since essentially all publications are written in digital form anyway, suppose that the publication clearance process is an online web-based service rather than a paper-based process. An author or other party fills out the clearance information for a paper and is prompted for the location of the digital document. The document can then be copied to a central repository. If the document is not in a convenient format, it can be converted automatically and at once to an Acrobat or the archival DigiPaper format using a conversion service such as Documents.com. An online workflow process can then be started which passes information about the document to all of the stakeholders in the clearance process.

A side-effect of this process can also be the entry of the document contents and descriptive data into a publication database. This is where the “enter once” value proposition comes forth. Once the document is entered and automatically converted to readable formats, it is a small matter to arrange that it appear on all of the web pages that display documents. In effect, each web page computes its contents by performing a query on the publication database – dynamically creating a date and author sorted list of published documents, complete with downloadable versions. This is done for the author’s web page, the area’s web page, the laboratory web page, and the center web page. The arrangement could automatically and correctly account for minor anomalies that can plague the process – such as when a publication has multiple authors who are members of multiple laboratories. Such an arrangement eliminates the problem of inconsistency in reporting, missed documents, unintended duplication, and the headaches of maintaining consistency across levels of the organization.

This same "enter once" value proposition applies not just for publications, but also for phone directory information, invention proposals, customer visits, and other items. It applies not just for research organizations, but for any kind of hierarchical or layered organization. Once institutionalized, it becomes a powerful adjunct for programs that reason about communities. For example, the earlier value proposition on finding expertise, could be powered in part by relying on up-to-date information gathered more easily by this sort of approach. A key factor in promoting the adoption of this approach is having simple tools for bootstrapping the data base from existing information sources so that up-to-date and relevant data seems to come available all-at-once and ready-to-go.

Mixing Creatively

Ideas come from many places, and sometimes the best ideas are combinations that come from mixing things from different places. Following the dimensions of content, context, and community – mixing can involve mixing documents from different sources and media, mixing ideas from different contexts, or mixing ideas from people in different communities. This generic value proposition addresses the barriers to different forms of mixing.

Rough Documents

A key practice in knowledge work is the sifting and recombination of fragments of material from existing documents. Through a process of ``document bricolage'' people intentionally extract printed text, outlines, scribbled notes, diagrams, charts, references, URLs, and so forth to generate new documents reflecting the evaluation and creative synthesis of ideas. Development and combination of a bricolage of printed document fragments is as easy as building a scrap book. Indeed, the "cut-and-paste" metaphor in computer user interfaces is in many ways a pale imitation of the paper-based operations that inspired it.

Although most documents are now authored online as formal structured entities, there are still advantages for treating documents and document parts as rough materials which can be annotated freely, sketched informally, and mixed loosely. Unfortunately, there are stubborn barriers to combining elements from diverse document sources. Scanning paper documents by-and-large delivers either raw images or pure ASCII; online, narrow and incompatible electronic genres make it impossible to mix images plainly visible next to one another on the computer console. Significant opportunity lies in breaking down these barriers. This is the yellow area of the Figure E-2.

The Rough Documents Opportunity

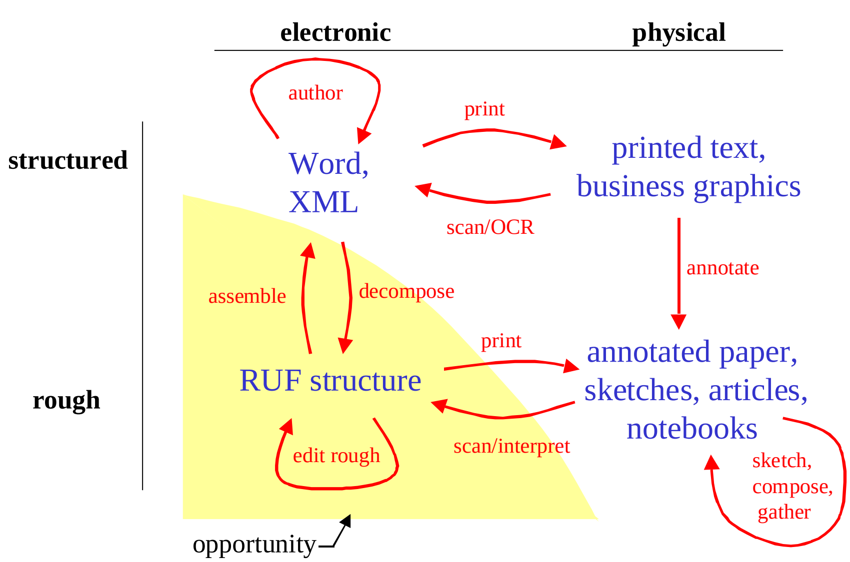

Figure E-2. Rough Documents

Rough Documents refers to technologies for dealing with document images in terms of meaningful aspects of visual appearance, as opposed to either textual content exclusively, or else unstructured bitmaps. Sources of rough material arrive in varied media, both online (email, web pages, Word documents, spreadsheets) and physical (stapled papers, bound notebooks, loose scraps, whiteboards). Core image analysis processes interpret document images in terms of perceptually significant objects such as text blocks, outline structures, collections of marks comprising annotations, diagrammatic elements, and so forth. Data structures reflecting perceptually salient entities may be called *RUF structure*. As independent objects, these elements can now be manipulated, combined, and printed in new arrangements. They can also be exported to electronic composition tools to participate in the composition of new formal documents. The value proposition is about mixing document materials creatively: a natural cycle becomes fulfilled whereby people can create and share knowledge around documents in ways that embrace informal and physically-embodied artifacts and practices as well as purely online applications.

Cross-functional Teams

Sometimes the most effective way to develop the knowledge for addressing a problem is to mix knowledge from different parts of an organization – mixing people in order to mix ideas. One clear place where ideas come from different areas is when the customers of enterprises span cultures. It is crucial to understand differences between customers in different places, and these understandings are most likely to come from a diversity of people. The Gartner Group addresses this in terms of diversity and the distribution of expertise. Thus, diversity outside is a challenge for meeting the needs of customers, and diversity inside is a resource for innovation.

> The ability to address cultural differences that may exist across geography, political boundaries, ethnicity, gender and other differentiators of needs will be critical to the success of product and service offerings in 2003. This issue is a major driver of the shift from international business to transnational business and the focus on customers in the marketplace. …
>
> KM speeds the recognition, capture and distribution of innovations throughout the enterprise. This is especially important in a geographically dispersed enterprise. Without KM, there is no infrastructure or support for systematic capture and leverage of solutions and process improvements. (from *Enterprise 2003: The Technology-Enabled Enterprise*, January 29, 1998)

In any large organization people need to find a balance in managing their attention between local and global concerns. Typically most attention goes to local problems. Local issues are generally more tractable than global ones; local expertise is generally more applicable to local problems than global problems; and locally-controlled resources are generally scaled more for local problems than global ones. However, it is frequently the case that choices based strictly on considerations for local optimization are in conflict with global optimization. Too much attention to global issues leads to organizations that cannot function; too little attention to global issues leads to organizations that cannot compete effectively.

In their studies of organizational knowledge creation, Nonaka and Takeuchi give several examples of how corporations work to combine the thinking of different groups. They review strategies ranging from focussed and transient cross-functional teams, to institutionalized cross-functional divisions, to highly fragmented divisions that are deeply engaged in focus groups with lead customers. Dialogue is a common element in all of these approaches. Dialogue amounts to focused listening and focused exchange between parties with different backgrounds and responsibilities. Dialogue is a key by which individuals and groups transcend their individual perspectives and synthesize new knowledge.

In the case of transient cross-functional teams, the teams are typically charged with a focused objective and given a substantial level of resources and support from upper management. When larger departments forge links, having a certain level of deliberate redundancy can be a crucial element that enhances an organization's ability to understand itself and to provide adequate communication and diversity. All of these approaches support a degree of mixing. Knowledge is often born in the fluctuation and creative chaos.

Market analyses of the offerings of knowledge management companies point to technologies that support collaboration and discovery. Beyond telephone systems, e-mail, and fax machines, basic collaboration technologies include teleconferencing facilities, digital support for networked meetings, real-time messaging, and so on. Highly networked individuals in corporations become familiar with the challenges of scheduling real-time meetings across international time zones.

How do the trends in enterprise organizations – such as globalization – affect these issues? As companies have become international, the idea of casually "walking around" to find out what's going on becomes expensive and impractical. How do the changes in computer networking and other technologies create opportunities for supporting these processes?

An open research question is what human services and technologies can be brought to bear on the new enterprise situations. One of the effects of outsourcing is that this introduces discontinuities in the processes that forge informal relationships in organizations. Outsourced organizations negotiate specific responsibilities and tend to be governed by explicit rules. In the absence of informal connections that can be used to create context and adjust procedures as needed, these organizational structures sometimes lead to rigidity and inability to adapt.

Inside Xerox, the outsourcing of human resources functions sometimes leads to situations where members of the organization have been known to spend hours on phone trees and in frustrating long-distance conversations to resolve non-routine issues. For another example, the Rheem company is a national supplier of water heaters. *Consumer Reports* recently reported a product defect in a batch of their water heaters and gave a phone number to be called to determine whether a water heater of a given serial number was part of the batch. It turns out that Rheem outsourced their fielding of customer calls and had an inadequate map of whom to call for actual service in the company. All calls in the Bay Area to get help servicing the defect were directed inappropriately to a plumbing company in Santa Cruz. That plumbing company found itself unable to correct the information used by the outsourcing company in forwarding calls. Such rigidity and low quality of service can arise in any large organization, but additional barriers arise when members of the larger organization do not maintain informal links.

In summary, the development and nurturing of cross-functional teams appears to be one of the most important vehicles for knowledge creation. Globalization and other trends toward distributed and evolving organizations tend to break down informal connections and increase the need for ways to create and support these teams.

Sharing Workspaces

One of the basic results of studies of effective collaborating groups is that they develop shared external representations and organize their work around them. An example of a shared external resource is a map of the physical layout of a computer chip, used by groups which share responsibilities for designing different portions of the chip. This practice of sharing external representations is common across a wide range of collaborative design tasks, include automobile design, airplane design, copier design, software design, and so on. External representations are resources. For information work, they make it possible to refer tangibly to work under construction. They surface a need for clarity in the negotiated meanings of terms and other symbols. They offer tangible expressions for places – interfaces -- where work of one subgroup must mesh with the work of another. Nonaka and Konno introduce the concept of "Ba" for shared spaces.

> Ba (equivalent to "place" in English) is a shared space for emerging relationships. It can be a physical, virtual, or mental space. Knowledge, in contrast to information, cannot be separated from the context - it is embedded in ba. To support the process of knowledge creation, a foundation in ba is required. (Nonaka & Konno, page 40)

For persistent organizational processes, it is increasingly common for groups to use online representations and databases to drive their processes. Where different parts of an organization come together, a crucial question is how their information systems and representations come together. Do they mutually design and then share a database? How do they trade off the expense of developing new and improved designs with the expense of modifying and maintaining legacy systems? How do they keep the interface between groups flexible enough to accommodate change, so that interfaces developed at one point in an enterprises lifecycles do not become too expensive as hindrances to further adaptation and change? How do enterprises discover their "latent communities" and retrofit connections that emerge as the situation changes?

Wide-Spectrum Searches

The basic search facilities commonly used on the web check for the appearance of specific terms. There is a great deal of variability in how people refer to things in writing. For example, an article referring to the United States may use words such as “U.S.” or “America” or even “Uncle Sam.” In the retrieval of whole documents this is less of a problem than it might seem because authors often use a variety of equivalent phrasings in writing, in part to keep the text from becoming monotonous. This natural advantage does not apply so much when searching for short documents or searching for short sections -- “snippets” -- in long documents..

Semantic matching can improve recall, that is, the probability of retrieving documents or portions of documents. Some techniques for reducing the requirements for exact word matching are routinely applied in information retrieval. For example, most retrieval systems use word stemming to remove suffixes and prefixes in order to convert terms to a standard form for the purposes of matching. For example, the words dream, dreamer, dreaming, and dreams would all be converted to the root word dream for the purposes of matching.

Beyond stemming, there are other ways of broadening the basis of a match. One way to broaden a search is to look for synonyms to use in the matching process. Thus, a search for documents including the term “city” could be broadened to gather documents using words like town, metropolis, suburb, and so on.

Synonomy is not the only relationship that can be used to broaden information retrieval. Suppose that someone was writing an article about mammals eating fish and there was an article about a cat eating a fish. It would be useful to include that article in the retrieval, even though the term cat is not a synonym of mammal. We can say that a cat is a kind of mammal, so there is a class specialization relationship (hypernym) between the terms. For another example, suppose that someone was writing an article about governments and taxation and there is a snippet that mentions court rulings on taxes. That snippet may be relevant. In this example, the term court is not a synonym of government, nor is a court a special kind of government. A court is a part of a government. Thus, there are a variety of kinds of relationships between terms that can be used to loosen the requirements of exact matching of terms in information retrieval. Extensive online dictionaries for matching and retrieval have been developed by the National Library of Medicine for medical domains.

Techniques for broadening retrieval have the potential benefit of finding useful information that would otherwise be less accessible because of variations in phrasing. In a broader sense, this sort of technology might be useful in looking for ideas or examples in documents. Suppose, for example, that a business analyst wanted to find examples of foreign companies buying telecommunications companies. One can imagine that the analyst would like to retrieve snippets containing phrases like “MCI to merge with British Telecom,” or “NTT considers buyout of Motorola,” or “Siemens increases holdings in Deutsche Telecom.” Notice that none of these examples contains the terms “foreign company,” “buy,” or “telecommunications company.” Identifying these phrases as matches requires that substantially more world knowledge be brought to the searching process. The value proposition of such techniques in a business context is that they could make it possible and automatic to find and combine many more kinds of relevant information than is the case today. Searches similar to this example have been used as one of the test domains for MUC.

As technology for broadening retrieval with wide-spectrum searches becomes available, different communities of practice are likely to build up their particular dictionaries of word relationships for synonym, class, and part relations. They may also build up rules for disambiguating word sense. There may be opportunities for different communities of practice in an organization to share these dictionaries both among themselves and across groups. This is an area where there is overlap with the community-oriented value propositions.

Building Community

In globally-distributed, dynamic organizations, the one-on-one approaches for building trust, relationships, and community break down. Walk around management is not feasible. Chance meetings at the coffee pot or in the lunchroom don’t happen. Community meetings at the town hall and auditorium don’t work. The people of such an enterprise need a way to function in electronic communities and to “network on the network.” Solutions must incorporate technologies with affordances for degrees of openness, group membership and anonymity.

Studies of social capital have suggested that enterprises may be able to exploit an “organizational advantage” for yolking together different communities of practice into coherent organizations. As Brown and Duguid put it:

> When [firms] get to the point they are so loosely connected that there is no synthesis or synergy of what is produced in their various communities – when as Teece and colleagues argue, there is no "coherence" – then a firm has indeed lost its edge over the market. … firms are valuable exactly to the extent … that they make communities of practice that expand their vision and achieve collective coherence. (Brown & Duguid, *Organizing Knowledge*, page 101)

The knowledge-sharing value propositions in this group echo other themes of knowledge sharing, but focus on means for discovering, building, supporting, and evolving communities in the enterprise.

Knowledge Maps

How does an organization make sense of itself? How does it keep track of its evolving structure? IBM's knowledge management teams construct what they call "knowledge maps," that model the explicit and tacit knowledge resources of an enterprise. Variants of knowledge maps include diagrams of knowledge value chains – which show how different organizations in a value chain apply their knowledge to an end. Knowledge value chains can include internal and external organizations. Knowledge maps and value chains are no substitute for dialogue and no panacea for knowledge creation. However, they can be useful for providing an overview of an organization and in discovering opportunities.

In the current state of the art, the information for creating knowledge maps comes from interviews. Such processes create a snapshot of an organization and relationships among groups. Because the process of creating a knowledge map is expensive and outside of routine work practice, knowledge maps tend to get out of date.

There is room for experimentation and creativity in the technology for automatically creating knowledge maps. For example, tools that "catch the buzz" in an organization – the patterns of activities and the identification of hot spots – can also be used to create information about an organization undergoing change. Tools that notice "what's new" can be an alert both to external forces on an organization and to internal activities. Tools that notice "who's new" in an organization also inform directly about how groups are changing. Tools that tap into the automated processes of an organization, or that summarize or aggregate information about activities over time, can help to create a map of an organization at work – in contrast to a static organization chart of how an organization at one point in time was intended to work.

Knowledge maps – and analysis of the underlying data used to generate them – may also provide an interesting force in the evolution of communities. Left to themselves, communities sometimes have a tendency to become rigid. The difficulty of creating consensus and methods of "best practice" can also serve as inertia that limits changing community knowledge and methods. Furthermore, communities inevitably have a localized point of view – according to their geographic location, their work context, and other factors. As such, the knowledge held by a community will not always be what the larger organization needs for a related problem.

Diversity Maps

A force powerful enough to drive evolution of a community comes from a source more powerful than the average individual in the community – other outside communities. The creation of explicit knowledge by communities – in the form of knowledge maps and other shared data – can be a kind of "boundary object" relating the community to others. For example, consider what happens when a community has developed and is willing to share dictionaries and other knowledge sources that are used in its analysis of online information. Suppose that another community seeks to use the same knowledge in a somewhat different context. Because the context is different, there will be certain breakdowns. Some searches and online information activities will fail. An analysis of the underlying rules may expose the context-dependent assumptions.

For example, in the technology of printing and copying, it is well known that there are problems in paper handling that arise in hot and humid climates that do not arise in cold and dry climates. This illustrates how experience and rules that work in one region do not necessarily work in another. Hence the knowledge produced in one region does not automatically have value in another.

A diversity map is a variant of a knowledge map that shows changes or variations in the knowledge in a different context. Restated, a diversity map is a knowledge map that is specialized by context. A different context could be a different work context, a different geographic region, a different market segment, and so on.

There is much room for experimentation in the development of diversity maps. How can the rankings, evaluation criteria, term equivalence dictionaries, and other data of a community be systematically tested for suitability in a different context? How could the evaluation criteria of a group be viewed with an eye toward understanding underlying assumptions and perhaps shifting the assumptions to avoid a parochial or regional bias? Tools for identifying context-dependence can be used either for evolving the knowledge of one community for use by another, or for use by a community when it seeks to extend its reach to another context.

Belonging by degrees

When people join an organization, they need to figure out how they fit in and how they can contribute. Most large organizations have a nearly-hierarchical structure so that individuals may join a small work group which is a part of a sequence of nested larger organizations. Part of what newcomers do is figure out how their work fits in with the work of their work group. They also need to understand how the work of their work group relates to the work of the larger organization.

In many organizations, the work of a work group relates strongly but not exclusively to the work of its parent organization. There can be an extensive web of relationships – an ecology – connecting workgroups to other parts of the organization. In a dynamic situation, these relationships are in a state of flux as new opportunities arise and as the organization itself evolves. Sometimes the rationale for locating a work group in an organization changes so much that the overall structure of the organization is changed.

Knowledge maps can be a useful tool for a newcomer in figuring out what people do and where to go for help. However, what a newcomer needs besides help is a way to fit in. In the context of a small group, local social conventions play a big role in introducing someone. However, walk-around introductions, group meetings, social occasions, and face-to-face collaborations that are the mainstays of becoming integrated into a workgroup do not scale well for covering an introduction to a large enterprise.

Within a large organization, an individual is likely to belong to several communities of practice. These might correspond to different roles a person performs in carrying out their job, or it may correspond to bridging activities to dispersed groups in an organization. For many years on the web there have been special interest discussion groups. These groups are evolving online communities. Different groups have different cultures and ways of relating to newcomers or “newbies.” Studies of these groups have identified different strata of membership – ranging from old timers, active participants, and “lurkers.” Lurkers are people who join the list and often don’t participate much at first in the discussions. They are sampling the message traffic to figure out how they fit in.

E-mail discussion groups are just one facet of the technologies for supporting geographically-distributed communities. The notion of “lurker” reflects an important part of socialization in a large organization. One reason for exploring distant groups is to learn what they are doing – and this is a crucial step not only for picking up a new capability, but also for becoming competent in engaging with a community in joint work of all kinds. People learn by positioning themselves on the periphery of a community. It is crucial that they be allowed to move slowly from the periphery into the community or into engagements with the community.

This socialization process involves not only learning about but also becoming known. How can online technologies assist a person in becoming known, establishing credentials, and establishing authority? How can individuals and groups accumulate social capital?

There is much room for experimentation here. Systems like eZmark make it possible for individuals and groups to reciprocally share bookmarks. This is similar to “belonging-by-degrees” in that sharing can be incremental. It can be set up so that people do not reveal more about themselves than they learn about the other party.

Establishing trust

According to O’Harra-Devereaux and Johansen as cited in a Giga report, ”Trust is the glue of the global workspace – and technology doesn’t do much to create relationships.“ A key aspect of trust is reciprocity – the sense that when people contribute in a community they can expect others to contribute as well. Indeed, one way of identifying the members of a community is that they have relationships involving substantial reciprocity. Social capital accumulates when people contribute to a community.

Trust is not an all-or-nothing proposition. In ordinary social occasions, trust is something that is built up incrementally. When we work with a new person, we may start by trusting them in small ways. If things work out well, we may extend the trust to larger matters. When we deal with institutions, trust is sometimes backed by the warrantee of the large group. For example, we trust banks with our money in part because they are backed by the Federal Reserve.

People sometimes need to experiment when they are trying something new. An issue of trust arises on the network when someone wants to ask a question and is worried about whether the question will seem stupid. They may not be sure who is watching them on the net, and may want to try learn something or something without risking their reputation. In cases like this, technologies that offer limited information about identity and some degree of anonymity may be useful in negotiating the early steps of contact with a community. This is roughly on the same order as when a computer facility offers a "guest" account for people on the periphery of the community. Security technology in support of this makes it possible to have a certificate that indicates "this person is an authorized member of a particular group" without identifying which person it is.

Another area where there is room for experimentation is in supporting reciprocity. How can reciprocity be supported between groups? An example of support for a reciprocal relation is in the identification of interests between two users of eZmark. Two users of eZmark can determine whether they have interests in common by sharing their bookmarks. But consider a situation akin to a computer dating service, where two people may want to discover what interests they have in common, but they don't want to tell the other person more than they learn themselves. In short, they want a "reciprocal" exchange of information. Recent security results can support such an exchange in a way that precludes cheating by either party.

Technologies that can recognize and identify how relations within communities (where reciprocity is established) differ from those between communities (where reciprocity must be cultivated) may help to extend reach between communities.

Latent Communities

When enterprises are large and geographically distributed, finding the right expert can be a challenge – as considered in the "finding expertise" value proposition. That scenario presupposes that one is looking for someone who has expertise on a specific topic – and that the purpose is to make contact so that the information can be found.

Community life, however, is not such a one-shot event. Sometimes what we need is not so much finding the wise owl as finding "birds of a feather." How can technology support the creation of special interest groups, or even the discovery that there are groups of people of a common interest who would benefit from creating a persistent community of people with a similar interest?

online shopping systems have been experimenting with systems that identify and support latent communities. A simple example of this, used to extend sales, is illustrated by the information that is returned when a person inquires about a book title from an online bookstore. A typical addendum to a search request is a summary to the effect that "Customers who bought this book also bought …" In short, it is assumed that people who bought the same book have similar interests. Another datum from the online bookstore is a more controversial service which lists the books in order of popularity purchased by people from a given institution.

This is an area with connection to many of the other value propositions. People who are interested in the same topics might form a latent community. People who use the same resources might form a latent community. People whose information is being fused together by a third party might be a latent community. Recommendation systems draw on personal profiles in making recommendations. People with similar profiles essentially form a latent community.

Latent communities are more tenuous than communities of practice. The fact that people have some common interests does not necessarily mean that they want to form a community or discussion group. There are potential privacy issues involved – in that people may not want their activities monitored that closely – for example, if they are sampling topics out of curiosity rather than deep interest. One capability might be to send a message asking a question to everyone who searched for a given topic – and asking whether they want to join a discussion group. In principle, such a message could be sent without revealing to the sender who the receivers were.

Using Readily

Knowledge work is more than finding information. It also requires making effective use of it. Use involves culling out the right stuff from the potential flood of online information, extracting out the essential parts, and putting it in the right places for use. Here are some examples of concrete knowledge-sharing value propositions that involve the effective use of ready information.

> Strategic Planning Assumption: By 2001, enterprises that lack ongoing KM programs and infrastructure will lag KM-enabled competitors by 30 percent to 40 percent in speed of deployment for new competitive programs (0.8 probability). …

Putting Information in Context

E-mail, by itself, is far from the ultimate technology for staying informed. A familiar lament in the workplace is "I spent the whole morning reading e-mail." If reading e-mail is crucial for staying informed – saving us from searching through a myriad of separate places for news from other people – why does "spending the morning reading e-mail" feel like an inefficient use of time?

For many people, e-mail arrives as a flat list of messages. Processing the e-mail is more than reading. It is about organizing messages for later work. I'll answer this message right now; this message goes to my to-do-later list; this message goes to the design problems section of project X; this message means that I should put something on my calendar. Reading and organizing e-mail is pure information work. It feels bad, however, when much of the work is mindless.

Many knowledge workers receive more than six times as many e-mail messages a day as phone messages. The current list of unanswered e-mail messages acts as the *de facto* "to do" list for many people. E-mail has become the main conduit for document exchange before meetings. For many people, e-mail is now a major conduit for information that requires timely processing.

Current e-mail systems and integrated office systems provide assistance for some automatic processing of e-mail. So-called mail filters recognize patterned classes of messages and handle them automatically according to user-specified options. This is useful for organizing e-mail into categories and for discarding junk mail.

Experimental e-mail systems attempt more sophisticated processing, such as recognizing categories of messages extracting information, and using it in automatic processing:

> "Here is a message from your boss requesting a meeting tomorrow at 10:00 at his office. Your calendar is open at that hour. Shall I put it on your calendar and confirm the meeting?"

Such so-called agent behavior, at least for many routine cases, relies on parsing and extraction technology for processing the e-mail, and on glue technology for interoperability with the calendar and other system elements. Such information extraction and automated behavior appears in experimental and non-robust forms in integrated information systems.

The example of automatically processing the boss's message about a meeting illustrates that e-mail information work involves putting information into context. In the example, the information is moved from an e-mail message into an entry in a calendar program and also to a response message. For another example of putting information into context, suppose that you have circulated a draft report and are collecting comments from various readers. Some comments arrive in e-mail messages. In a manual information-processing approach, this might involve printing out the message and then putting the page with the report or moving the message into a "comments on the report" mail folder. A later step would transfer parts of the message into pieces of the report. In each of these cases, the information is moved to "where" it can be used, that is, to a context which is ready for using it.

This is an area where current tools fall far short of what is technologically possible. In the paper-revising example, imagine that when the writer begins working on the report again, the comments from different readers are organized on the right hand side of the display. The comments could be seen in the margin" next to the point in the report where they refer. In the language of the Notable and Notepins projects in HDI, the comments are extracted from a message and pinned to the relevant portion of the document under revision.

Going beyond currently available systems, there is an opportunity for substantially more automation here. The underlying technology for supporting comment-driven editing would include linguistic matchers for segmenting messages, user interface technology for showing comments in context, and various glue technology to connect the programs together.

Just-in-Time and Just-in-Place Information

In preceding examples we considered how some information processes move information into a context for using it. The reverse notion is also useful, where a work context is used to pull in, categorize, and select information. The intranet portal, Scient Zone, uses context to inform search. It creates a task environment that supports a Scient employee's activities on one or more projects or tasks. In order to provide task-specific help, the environment keeps track of what project or task the user is engaged in. When the user initiates a search for information, the environment supplies additional information to the search engine so that it can carry out a more focused search. In short, the task context provides information to the search engine for disambiguating word sense and for indicating the kinds of documents that are most relevant. This example uses information about task as a kind of context.

Taking this idea of using context to inform search further, query-free retrieval retrieves information without an explicit search request by a user. This concept is different from either push or pull technology. It's about information that is just-in-time and just-in-place.

For example, suppose that someone is sketching out a version of a paper – either as an outline or as a sequence of semi-formed sentences and phrases. An "outlining sensemaker" tool would use the phrases appearing in the draft outline as grist for a search for relevant documents in a corpus. It would assist the report writer or "sensemaker" by locating relevant paragraphs in the selected corpus of documents and offering them up for consideration and information extraction within the work context. In effect, a sensemaking tool is a companion researcher working along side the sensemaker. The sensemaker could decide to include them as notes or background material, incorporate portions of them, or discard them. As the sensemaker changes the paper or shifts focus, the sensemaking tool would shift its focus and display. The information arrives both just-in-time – when the sensemaker needs it – and just-in-place – in the relevant work context of the sensemaker.

The use of context to inform search and to retrieve information automatically has much broader applicability than supporting technical writers. For example, in the earlier example of an attorney preparing arguments for a Markham hearing, the automatic retrieval could be useful in finding both examples and counter-examples that are relevant to the case. For another example, consider a knowledge worker with many ongoing projects. As the knowledge worker shifts his focus – from one task to another, from one account to another, or from one topic to another – a query-free retrieval could carry out a search of e-mail and bring up the headers or sections from the most recent messages.

A variation in the just-in-time and just-in-place value propositions involves wireless connectivity and handheld devices. In a sense, cell phones have radically changed just-in-time and just-in-place conversation. In the same way, handheld computers – even location-sensing computers – will become important in the next few years as the public expectations of the Internet evolve away from the desktop. Increasingly, people will be able to send and receive messages from mobile devices and retrieve visual or audio information from mobile devices. One could ask a mobile device for the location of the nearest Thai restaurant or the nearest theatre playing a particular movie. When bandwidth and screen real estate constraints dominate, the information delivered just-in-time or just-in-place may also be just-at-the-right-level. For example, e-mail may be summarized or filtered to be especially relevant to the time and location of the user.

The potential underlying technology supporting just-in-time and just-in-place information includes segmenters, matchers, topic-driven term equivalence dictionaries, as well as glue for integrating this technology into a work and mobile devices.

Catching Up

One of the familiar difficulties of coming back from a few days absence is that e-mail and other messages pile up. Given recently measured figures that knowledge workers receive twenty to seventy messages per day, the pile of e-mail can quickly become quite daunting. E-mail can pile up even if a person does not go away – such as when there is just a lot of traffic on a given topic. When the traffic on a topic is intense, one does not have to go a way for a few days to need help keeping up. A related challenge is in making sense of an archive of e-mail on a topic.

Current e-mail tools provide some help for this by sorting messages by time and sometimes classifying them according to user-provided rules that check names of senders and for key words in the subject field. Nonetheless, especially when there is an extensive series of messages on a topic, there is much more that could be done.

The mail stitcher idea is to have a program that goes through a pile of e-mail and makes a personal newsletter for the messages. It removes redundant messages, removes repeated message content, and creates visualizations and summaries of the content. It could also fetch attachments and linked URL's for rapid presentation of the information. A mail stitcher might work differently when summarizing messages for a day versus summarizing messages for a week – in terms of creating single or multiple newsletters and levels of summarization.

The idea raises several questions: When should you get multiple newsletters? Do you get one article per topic thread? Presumably there should always be a one-page summary. How do you best do a semantic zoom? How could a newsletter be used for introducing a new member of a group to a group discussion list? If multiple messages contain elements of information on a given topic,

The underlying technology for mail stitchers would include linguistic technology for parsing, categorizing, and extracting information from e-mail, as well as specialized summarization and visualization technology.

Identifying Document Redundancies

The maturing of the office computational environment has placed new tools in the hands of both Information Technology departments and non-specialist office workers. For example, Microsoft's FrontPage application now ships as part of Office 2000. Such tools make it easier for non-specialists to build web pages and have led to community groups creating and maintaining their own web sites.

Acting on behalf of local communities, workers can now create "grass-roots" knowledge bases that support their local practices, and IT departments can deploy richer versions of such repositories across the enterprise. The result has been a drive to move paper-based systems onto the web and to create new systems specifically for work-practice communities.

Placing vital information on the web in the form of document collections is just a first step in knowledge management. As time passes and the environment changes, documents once regarded as critical knowledge become stale, or are recognized as inaccurate or only to apply to particular cases. Active and ongoing management of online collections is required to enable workers to continue to use them readily and with confidence.

Human editors are often very good at making fine distinctions among similar documents and maintaining the quality of a repository. Deciding what information to retain and what to discard is a complex task, combining domain expertise with the skills of an online librarian.

As Eureka and other such knowledge repositories scale up, these expert editors, who typically have other pressing responsibilities, find it increasingly hard to identify those documents that have either become stale or are in some critical way similar to other documents. By exploiting computation to analyze the document collection extensively in the background, automated systems can periodically present the experts with much smaller sets of tips for further action such as deletion, combination or clarification.

Several technologies are relevant for automating systems that identify document redundancies. Domain-specific common-sense knowledge particular to the task at hand -- repairing photocopiers in the case of Eureka – is needed as a foundation. Natural language technology is needed to connect the variously-phrased text of the documents to the relevant semantic concepts. Recent advances in the state of the art now present an opportunity for tackling this value proposition as a value-added knowledge service.

Coordinating Action

Taking action with information is not always a solitary activity. Sometimes coordinated action is needed. For an example of massively coordinated action by a company, its suppliers, and its distributors, consider the national recall of Tylenol days after the 1982 Chicago area poisonings. Johnson & Johnson found it necessary to recall some 30 million units from drugstores across the nation. To salvage the brand, a massive public relations effort was launched within weeks of the poisonings, and within months Tylenol was back on the shelves, the first over the counter drug to abide by new FDA packaging regulations.

The need for coordinated action also arises in small and medium-sized situations. Workflow systems often support a process of handoffs – where information and documents are passed electronically from one person to the next. Workflow systems provide a high degree of automation when business processes can be planned in detail and where the information is highly regularized and kept in data bases. A major motivation for deploying workflow systems is to increase the speed and decrease the cost of business, often eliminating paper handling.

How can a company flexibly distribute information and coordinate action when business is less routine and regimented? In cases like the Tylenol example, the public expects a corporation to act in a coordinated way.

Nonetheless, a common criticism of large organizations is that they are slow and that “the right hand does not know what the left hand is doing.” The problem is not that big organizations have inept employees. Rather, the problem is that coordination becomes more complicated with large scale. Questions that are simple in small groups are complex in large enterprises. Who needs the information? What is the status of the process now? How do workgroups regulate their priorities and activities when authority is decentralized? How can an organization optimize value for the larger enterprise rather than for an individual business group? Could a value-added model or market-negotiation model facilitate joint decision-making? How could a firm organize its knowledge-sharing activities to achieve an organizational advantage over a competing market place of small, independent firms? How could a firm lower the costs of negotiation and mutual understanding in order to achieve efficient coordination? Companies are answering these questions now with a combination of technology and corporate culture. There is much room for research and modeling informed by cases and experimentation.

Disseminating information to facilitate coordinated action is a difficult task in the best of circumstances, but one that will become ever more critical as firms become ever more geographically and organizationally distributed.

Reusing work

In business school and law school, teaching is organized around cases. As taught by the overpowering curmudgeon, Professor Kingsfield in the 1973 movie *The Paper Chase*, “Expect no final answers to questions in this classroom. Expect cases. After cases expect another question and more cases. … You come in here with a head full of mush and you leave thinking like a lawyer.” The main experience in business school classes and law school classes is in learning to work through cases.

In the same way, the main experience in large business consulting companies like Arthur Anderson, KPMG, and others is also in reasoning about their consulting cases and the main knowledge asset of these companies arises from their case experience. So how do they leverage this expertise in a global enterprise? Part of their response is to put proposals and case reports online and make them accessible to all consultant teams. In principle, this makes it possible for practicing consultants to be more competitive in the efficiency of their analyses – reasoning from prior cases and extracting out business context, industry forces, and best practices. By sharing their case data, consultants avoid repetition of effort and reduce their project completion times.

Even in other companies which may be less case-based, there is considerable interest and potential leverage in building community and making records of work accessible through information portals. This includes not just best practices, but also sample documents – not just templates, but actual filled-in proposals, deals, summaries, reports, and so on that can serve as starting points for new ones.

Practical as this simple sharing approach sounds, there are substantial issues in its implementation. A key issue is security of client information. Consulting companies serve specialized industries and often have on-going projects with companies that are each other’s competitors. As such, a consulting company has a fiduciary responsibility to protect a client’s data. The possibility that one client’s sensitive data might leak to a competitor via another team makes clients and partners very nervous. The problem here is that security is as strong as the company’s weakest link. Careless activity by one ambitious or junior consultant could undermine the reputation of the company. One form of sharing that addresses this concern is simply passing around “sanitized” presentations. This reduces the risk of accidentally releasing confidential information, but it also reduces the quality of the knowledge that the consulting company effectively shares. Furthermore, the work required to sanitize or disguise sensitive data is considerable. If the consulting system is driven solely by billable hours, the work of writing up and sanitizing cases tends to drop to a back burner.

Another issue that arises concerns the real re-usability of proposals. Previous work can be re-used to get a reasonable-sounding proposal out to a client quickly. On the other hand, indiscriminate use of "boilerplate" material can lead to consultants neglecting to do the work to really understand the details of an individual client. In at least one company, such "excessive generalization" from other consultants' materials was seen as contributing to less focused proposals, and fewer of these were accepted by clients.

This is an area where companies are trying various approaches. There is no shortcut to training individuals about the issues involved in protecting client data and the issues of deeply understanding a client’s situation. In situations where companies opt to sanitize the reports or to create and share reports that are auxiliary to the regular work practice, companies have found it crucial to provide incentives for contributing the reports and also to vet the reports.

There is room for research in addressing these issues. With regard to the question of whether the re-use of materials leads to lower quality reports, there may be deep underlying tradeoffs between modes of sensemaking – between synthesis and integration. At this time there is no "science" of sensemaking; restated, there is no practice of what constitutes good sensemaking practice analogous to what constitutes (for example) good accounting practice or good engineering practice. As sensemaking continues to grow as an activity – especially with the increased use of online information of uneven and diverse quality -- tools for partially automating and supporting sensemaking will need to co-evolve with the principled development and evaluation of competing practices.

There are also opportunities for research in finding ways to support the generation and re-use of materials. Some of these are discussed under the "innovate creatively" value proposition. These include ways of matching documents to new situations, to finding ways of uncovering "old" context in a document that may need updating to a new situation, ways of formalizing aspects of the documents to support greater automation, and ways of supporting commentary on the use of the documents as the community gains experience in developing a sharable knowledge base. In the context of consulting companies, there may be room to find ways to automate the process of sanitizing presentations.

Understanding Deeply

Several of the previous value propositions addressed ways for speeding up access and use of information. Some kinds of information are vital for an organization and are used for crucial decisions. In a large and evolving organization, a substantial challenge to understanding information from other parts of the organization is that there can be differences in terminology and in practices for categorizing or classifying information. The Gartner Group sees this challenge as leading to the creation of "knowledge maps."

> A "knowledge map" contains a taxonomy (i.e., a classification scheme), cross references from the taxonomy to sources of knowledge both tacit (i.e., people) and explicit, and an operator interface (more and more frequently a Web browser) for navigation. The taxonomy represents the way the enterprise thinks about its world. …
>
> No taxonomy is valid indefinitely. No single individual can know enough to maintain the entire taxonomy for an enterprise, so responsibility must be assigned to multiple individuals, each of whom maintains the taxonomy for a particular subject area of domain. (from *Enterprise 2003: The Technology-Enabled Enterprise*, January 29, 1998)

At the time of this writing, the viability and limitations of the "knowledge map" idea are still in question. For example, the existence of a knowledge map – even developed by corporate teams – does not by itself assure that it will be useful or used. Furthermore, much of the knowledge used by a corporation necessarily comes from outside its borders and is therefore influenced by how others use terms. Nonetheless, the development of knowledge maps reflects the considerable difficulty that people are experiencing in effectively using information from varied sources, as well as the substantial technical challenges in automating access and fusion of information given linguistic diversity.

The "understand deeply" value proposition is about technology that supports people in working with confidence with detailed and vital information, assured that it is correct and up to date.

Fluid Categories

Categories are an important aid to thinking and a central feature of many information systems. Cognitive scientists have noticed that much of our mental commerce with an environment deals with classes of things ("categories") rather than with unique events and objects. File systems – directories and folders – present implicit categories by virtue of their presentation of information in terms of logical trees. For example, named folders in a file directory implicitly present statements like "This folder is where programs go"; "This is where documents go and this is where presentations go"; "Files for project A go here and files for project B go there." The problem with strict categories – where everything goes in exactly one right place – is that the system almost always breaks down. Where do we put stuff that is about both project A and project B? Do we put all our presentations in one directory and then subdivide by project, or do we put all our project stuff in one directory and subdivide the directory by documents, presentations, and programs? What happens when two people with their own systems for categorizing and organizing things need to work together?

More precisely, the problem is not just with the strictness of categories. One problem is that items can belong to more than one category so that the notion of mutually exclusive categories breaks down. That problem can be addressed by allowing items to belong to more than one category. Another problem is that file and folder systems tend to organize categories in trees, where branching at different levels is based on different kinds of desiderata. This approach forces a user searching for information to know not only about categories, but also about the order in which the desiderata for categories are represented in the tree. That problem can be addressed by using underlying property-based representations for items -- such as the Placeless system -- rather than nested directories. A third problem is that membership in a category may be approximate. That problem can be addressed by having rules for "fuzzy membership" or other kinds of approximation logic rather than a Boolean (true/false) logic. A fourth problem is that there can be a mismatch between categories that a user needs and categories that are represented in the system. Sources of mismatch include using variations in naming and variations in which distinctions are useful between one task and another.

Issues about categories come up generally in knowledge sharing because searchers may need to know about categories or implicit organizations in order to find things, and sharers need to consider how organization can facilitate or get in the way of people using the information. Moving to a flexible category regime can, however, have substantial benefits to understanding. A fundamental truth about the search for the "right" representation is that different representations are better suited for different purposes. Making it possible to easily re-organize a representation means that it becomes easier to reuse it for different purposes. It becomes easier to see different kinds of patterns in the information and to extract and put different parts of the aggregated information to use.

Level-three technologies that support such flexible reorganization of information include clustering systems, visual organizers, multiple view interfaces, and the broad set of linguistic tools for matching and retrieval. One concept from research on sensemaking is the "docu-lens." A docu-lens says roughly – "here are my categories. Organize the information from the corpus of documents this way and then show it to me."

What-now Documents

The world is not static. Information can become dated and reports can become stale. One way to try to keep up is only read the latest news. However, every organization tends to have an abiding interest in certain focused topics, meaning that work needs to be invested if documents and collections are to be kept evergreen.

In accounting and business planning, it is common to develop spreadsheets that capture some of the logic of a variable situation. Planners use dependency-linked spreadsheets to explore scenarios, doing "what-if" analyses that show them the budgetary consequences automatically when they vary their starting assumptions. A related idea is to represent symbolic dependency information for "what-now" or "living documents". By analogy, a what-now report could link conclusions to other sources, to symbolic representations of assumptions, or to measures of belief. In an intelligence organization, it sometimes turns out that information gathered quickly is subsequently shown to be unreliable. A living intelligence report would allow a reader to disqualify sources or assumptions, and then have the system show which sections or conclusions become suspect. A living document in an organization could also link its conclusions to a database, so that it always presents the most recent information.

Multiple Perspectives

Different ways of organizing and displaying information make it easier to see different concepts and patterns. For a simple example, a chart that puts sales of different regions together makes it easier to compare region, and a chart that puts different time periods for the same region together makes it easier to see local trends in sales.

Technologies relevant to this value proposition include systems that show multiple simultaneous views, such as an overview and local view, and systems that help users to create multiple views, such as scatter-gather systems.

Translation for Understanding

More and more, people in multi-national companies and people using the web for information are finding a need to gain a level of comprehension of material not in their native language. Native speakers of languages other than English often have to deal with English to obtain information in a (web)timely fashion. And English speakers are noticing that not all web pages they might be interested in are written in their language. A variety of conveniently accessible translation assists can be of significant utility here.

For people with some knowledge of the language of the text, an intelligent glosser can be a big help. The same type of facility can be used to assist people in dealing with unfamiliar technical terminology in texts written in their own language, in this case by supplementing the text with definitional glosses and with references locating explanatory/tutorial material.

Where the reader has little familiarity with the language of the text, full-scale automated translation comes into play. The results, given current technology, will depend on the degree to which the translation mechanism is adapted to the subject domain, and will not be perfect in any case. But even current results can be very useful, and technology is improving in this area.

Finally, a potential combination of translation technology and summarization should be mentioned. Unless a translation system is very well-tuned to the subject domain, the results will probably be tedious to read. In particular, they will not be as scannable as human-generated text. Yet, for many purposes, either the first, or only, interest of the reader is to determine what the text is about. In these cases, it makes sense to apply a summarizer to the text first, followed by a translation of the summary, so that the reader can determine whether the additional effort of reading the full translated text is worthwhile.

We can also extend this idea to the problem of providing some idea of what a highly technical document is about to people not familiar with the technical area. In this case, it may be possible to produce special-purpose summaries, geared primarily to identifying and defining the major subjects covered by the text.

Mapping a Collection

The need for understanding applies not only to documents, but also to collections. As information portals are coming online, they are adding new capabilities to overcome the shortcomings of directory-oriented file systems for managing corporate information and knowledge. Examples of these portal capabilities are fusion of sources, indexing, and search. As information continues to accumulate online and enterprises continue to build up larger and larger knowledge bases online, observers are predicting that these collections of services and documents will collapse under their own weight.

Xerox customers (such as the Los Alamos National Laboratory) are struggling to gain control of their knowledge management processes. What is needed in managing large collections is more than the current capabilities of portals and also more than library-related concepts like “cataloging.” Cataloging connotes properties of earlier, more static collections in which it was more reasonable to assume that there exists *a priori* a catalog that one catalogs new documents into and a standard classification system that one uses. The term cataloging also misses the possible richness of visual maps, which may be customized for different purposes and different people.

Consider a system where a user can move items into a variety of clusters. An item might appear in more than one cluster and may be reclassified as a user changes the rules for classification. Visualization maps can have the richness of multi-level and multi-color hyperbolic browsers that show multiple relationships among items. There may be dynamic inferential rules that derive labels or properties for library items based on meta-data and content analysis. The properties on items may be sharable and persistent using an underlying technology like the Placeless system at PARC. There may be multi-modal clustering algorithms that assist a user in exploring candidate groupings of items.

When people access information on the web, the immediacy and convenience of a response leads to the expectation that the information is correct and current. This creates a premium for technology that helps an organization to replace or discard information that is no longer current.

Documents for Thought

"In one ear and out the other" describes conversations where we aren't really listening. A child will often check to see if a parent is actively listening by asking the parent to repeat what was just said, to answer a question, or to offer a comment. This tendency to "tune out" can apply to reading as well as to listening. It is all too easy to spend time reading a document without reflecting adequately on what the document says. In the corporate setting, this problem is exacerbated by the fact that more and more documents, presentations, and e-mail are demanding our attention. Improving the quality of reading and listening in a corporation is certain to improve the quality of knowledge sharing within the corporation.

How can we improve the chances that a document will be carefully reflected on by a reader? First, we can give readers tools that help them to focus more clearly and engage more deeply with the documents they are reading. Second, we can give authors tools that will enable them to evaluate how clearly their messages speak to readers and to respond accordingly. When such tools are used by a community of writers and readers, there is a chance to make the document-centered process more conversational and to give feedback about the collective understanding of a community.

Traditional techniques for improving the reading experience work by encouraging active engagement of the reader with the material and by directing the user's attention in reading. Manual techniques for readers include annotating, highlighting, focusing on key sections, and summarizing. Note taking in the margins is often the hallmark of a researcher. Another common strategy for reading is to focus on key parts of the document, including the abstract, topic sentences, and conclusions. Highlighting text during reading is a favorite technique of students. Writing summaries of what was read is less common but often even more useful because it forces a reader to think more abstractly about the material. online counterparts to these techniques are beginning to emerge. Techniques for annotation (including those under development at PARC and FXPAL) provide readers with tools for note taking while reading digital documents. Progressive disclosure techniques (possibly including fluid user interface techniques) can give a reader more kinesthetic engagement with reading. Novel techniques for encouraging readers to focus on key parts of the document might include automatic highlighting of topic sentences or even highlighting of sentences that correspond to a particular perspective on the document. For example, a person in marketing might wish to see marketing-related sentences of a document highlighted in the text. Automatic summarization techniques, while they do not engage the reader in the construction of summaries, can give the reader a chance to pause and perhaps notice the main points more clearly. A special purpose summarizer might be designed to call out main points, due dates, and action items. Finally, an interesting example of an online system that encourages the authoring of summaries is amazon.com, which creates a sense of community around book reviews. The incentive to write a review for amazon.com encourages readers to think perhaps more carefully about what they have read.

Giving authors tools to evaluate how their documents are read is tantamount to making the interchange between author and reader more conversational. Changing from the reader's to the author's perspective corresponds to changing the question from "What's this all about?" to "Is anybody listening?" If an author can determine how a document will be read, the author can tune the document before distributing it. By invoking the same summarization tools that a reader will use, a writer can determine what points would be highlighted or emphasized. This creates an opportunity for tuning the text for clarity or for marking the text for the summarizer to help locate the main points. As was noted earlier in this paper, the traditional publication process has this flavor. An author writes a draft, a reviewer comments on how well the message is delivered (among other things), and the author fine-tunes the document accordingly. Techniques for allowing an author to see how often a document is downloaded, to see which sections of a document are read (tracking of reading becomes possible if readers navigate their way through documents via progressive disclosure in user interfaces), and to view questions posed by readers will all serve as forcing functions for rewriting the document to be more understandable and more engaging for readers. Feedback of this sort can be collected and aggregated from a community of readers. This could give feedback to authors about how their messages are being read, and information to a community of readers about how others are reading a document.

Innovate Creatively

Because of the increased market efficiency brought on by the Internet, company profits must come more from innovation and less from monopoly and market positioning. How can a company respond to rapid changes in its environment and innovate successfully to meet them? In this context, innovation refers not only to technological innovation. Much of the crucial innovation in a successful company is in business innovation – "re-inventing" its own business processes, market models, channels, and partnerships in order to thrive.

Analyzing the emerging critical success factors for corporations, Gartner Group highlights effective innovation.

> The profits of a global economy will come from innovation and not monopoly or market isolation. By 2003, more-perfect information and decreased barriers to entry will drive the marketplace to rewarding innovation over positioning (0.6 probability). The successful enterprise of 2003 will reward innovation and will have developed policies and technological infrastructure to bring innovation to market rapidly. Innovation drives change, and the successful enterprise of 2003 must have change-management capabilities and flexible applications to support continuous process, product, service, and channel innovation. (from *Enterprise 2003: The Technology-Enabled Enterprise*. January 29, 1998).

Collaborative Inventing

Does a large corporation have advantages in innovation? Forward-looking technology-oriented corporations have the capacity for sustained investment in long-term research. Such investment potentially gives a company awareness and access to future technologies not easily available to small companies and start-ups.

Although research is an important aspect of innovation, it is not the only one. Much of the growth of companies in the past decade has come from innovation in business practices – where companies recognize emerging opportunities in markets and opportunities for increased efficiencies in business. Traditionally, the size advantage of big companies leads to an opportunity to exploit economies of scale. For example, in the automated manufacturing of hard goods, the unit cost for manufacturing and marketing a product substantially decreases as production volume increases. This gives an advantage to the first big companies to develop a product – since they benefit from increasing margins and are first to climb the learning curve for more efficient manufacturing processes.

In the context of global markets, however, mass production is now taking an interesting twist: mass customization. It is becoming increasingly important to tailor products to the special needs not only of geographic and cultural regions, but even niches at a much smaller scale. Especially for consumer-oriented companies, mass-customization requires developing products that are tailored to customer needs while still benefiting from mass production and dominant branding.

Within this context, what are the sources of strength for large companies? Size is not enough. About innovation it is said that "need is the mother of invention." Different contexts create different needs. Recognizing those needs requires a diversity of eyes and minds. This suggests that an important potential strength of a large company is in its many eyes and its diversity. A global company includes people from different regions and different cultures – leading to many eyes on the opportunities in the world and to cultural diversity in thinking. A company with specialists of many kinds of expertise can potentially bring powerful and diverse points of view to bear.

For example, modifications to a printer to handle paper in high humidity weather – developed in Japan – may be useful in Washington, D.C. For another example, devices with large buttons that appeal to the "disabled" may have a larger market for people who like the ease of use. How can organizations spot the opportunities to re-use designs from one cultural or geographic setting to another, or from one business context to another?

It is also said that "opportunity favors the prepared mind." How can a corporation harness the potential power of its organizational mind? How can it tap the diversity of its people to innovate effectively? Can the many eyes of the company recognize the diverse needs of its potential customers? Can the many minds of a company organize to understand how to satisfy those needs with mass-customization?

Recent studies suggest that a key for a company to exploit the power of its diversity lies in the abilities of its groups to accumulate social capital. In effect, it is crucial for groups in one part of a company to identify and work with the groups in another part. For example, suppose that a company is expanding into a new market in a new region. How can marketers in the region come to recognize how existing products need to be adapted in order to sell better in the region? How can the corporate specialists responsible for design and manufacturing interact effectively with the diverse markets, tuning and re-tuning the products to enhance their utility and appeal? How can product ideas that arise in one part of the world be tested for appeal in other parts of the world?

These questions are all about knowledge sharing for innovation. They are about not only acquiring and sharing knowledge, but putting it to work. They are about collecting observations and generating insights. They are about brainstorming – involving people from different parts of a diverse company. They are about how a company marshals its resources – making informed decisions for goals of the larger enterprise by drawing on the strength of its diverse communities.

In summary, collaborative inventing leverages the power of community by enabling key groups to work together. Relevant technologies include those discussed under the "mix creatively" value proposition – supporting cross-functional teams and shared workspaces. They include those under the "access rapidly" value proposition – for finding expertise. They include communication technologies for groups at a distance. Research questions include developing practices and technologies that enable groups at a distance not only to find each other and work together, but to leverage each other effectively so that the potential strength of a large company can be realized to its advantage.

Digital Sandbox

Effective innovation in a complex world requires more than just people getting together to share ideas. Some problems are inherently complicated – too complicated to work through in a conversation, even using the proverbial "inventor's napkin." The difficulty is that the complexity – either the amount of information or the number of connections and interactions – is just too big to hold in mind. This is where computational tools that support external cognition provide leverage for individual and group visualization, design, and innovation.

The creation of digital tools for complex problems – "digital sandboxes" – take varied forms. One of the most popular tools for working through financial cases is the ubiquitous spreadsheet. In the past few years, the spreadsheet has been augmented to scale to corporate databases – with so-called OLAP (online Analytical Processing) and ROLAP (relational) technology. These tools provide ways of summarizing, tunneling down to details, and carrying out what-if case analysis. In specialized design tasks, specialized computational tools include tools for major branches of engineering design – ranging from chip design to building design.

An open question is whether there are new variants of the digital sandbox to be discovered that can provide substantial leverage for knowledge tasks in large corporations. Such opportunities may arise from a combination of developments: the rapid accumulation of online information; the availability of tools for online linguistic processing; the availability of increasingly powerful tools for visualization; the simplification of protocols for information networking. One promising avenue is the combination of tools for thinking with tools for finding information. In this way, documents and design examples relevant to a context can be searached for or brought to the attention of a designer at work. Another promising direction involves the integration of notation tools with digital sandboxes – making it possible for people to share commentary in a task-specific work context.

What's-needed invention

Different contexts create different constraints and requirements. How can a corporation re-purpose materials created for one context to another? One approach to this problem – a variant of document sharing – is to assume that the biggest part of the problem is in distribution. One need only put case studies, design examples, forms, and other documents online. If people can find them, they will already have a leg up addressing new situations – and can jump start their approach by starting with them. This attitude assumes that the hard part is in just sharing the information; re-deployment is assumed to be easily handled by local people at different sites.

Such an attitude has shaped the current generation of commercially-available knowledge-sharing tools, when companies are still coming up to speed with the idea of information portals. As more companies employ such approaches, effectively this raises the bar for expected performance in using shared information. As rote sharing becomes widespread, attention will turn more to the issues of effectively re-purposing the information. How often are there substantive mistakes in re-deploying information and designs? Are there cases where assumptions and stale information from an old case was inappropriately carried forward in responding to a new context? Are there cases where proposals in the new context were not competitive because local experts did not understand the knowledge deeply enough to tune it to the situation, perhaps eliminating relevant portions that they did not understand?

These questions suggest possible directions for research on the next generations of sharable explicit knowledge. How can design assumptions and context-requirements – and their interactions with other knowledge elements – be highlighted in external representations? How can the dependence of representations on potentially stale information be discovered and corrected when new contexts are considered? In short, how can we make external representations be more-intelligently reusable, so that they fit the real needs of a new context? How can we capture and make re-usable knowledge about re-purposing different kinds of documents for different regions and contexts?

Appendix F. The Sensemaking Cycle: Themes and Variations

Appendix D describes a model of the sensemaking cycle that is based on a conversational metaphor. This appendix explores variations on this model. To review, the conversational model (presented again below) is a three-step model. It divides the sensemaking cycle into acquiring knowledge, sharing knowledge, and acting on knowledge.

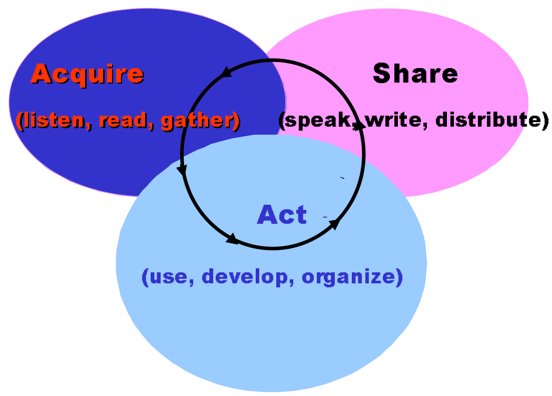

Figure D-12. The document-oriented analogy to the conversational sensemaking steps "listen-speak-act" are "read-write-act". In order to generalize beyond any particular medium, we propose the more generic steps "acquire-share-act."

Acquiring knowledge involves taking action at both the technological and cognitive levels. On the web, acquiring knowledge involves not only reading, but also searching for a relevant set of documents, choosing a document, and then reading the document. The story does not end with reading. Reading and listening involve not only perceiving and processing the symbols, but also interacting with the text in the context of the reader's own experiences and understanding of the world.

Sharing knowledge is likewise a complex process. Sharing can take place not only through speech, but also through documents. The arrows between acquiring and sharing in the acquire/share/act diagram can be understood at multiple levels. At the most straightforward level, the arrow signifies that a person who writes a document is constructing something that another person can acquire. At a deeper level, the arrow signifies that the process of sharing is interleaved with the process of acquiring. The sense that an idea can take shape through the process of writing is familiar to anyone who writes.

Acting, the third step in the conversational model, also includes actions at many different levels. Decision-making is an important action in the corporate setting. Figure D-4 depicts a very concrete type of action: accomplishing work as a result of acquiring and sharing knowledge. Organizing is an action that is fundamentally a sensemaking activity. Consider the note cards that a student might create in the course of researching a term paper. Acquiring knowledge involves finding references and reading them. Sharing knowledge involves actually writing the term paper. Acting involves not only completing the assignment but also organizing the note cards.

Many researchers have discerned and written other descriptions of underlying cyclical processes for sharing or creating knowledge. These cyclical models all show how knowledge and information are used to create more knowledge and information. The models share some of the same steps. They also differ in interesting ways. Models differ in their scope – so that they sometimes complement each other by illuminating different parts of a larger process. Models differ in the degree of detail and focus that they bring to different parts of a process. Finally, it should not be supposed that the different models are all variant descriptions of the same process. Researchers study different people doing information work, and different people work in different ways. The knowledge-work "verbs" listed in the body of this report can be combined in many ways to describe many kinds of work practices. Tasks differ in their "information cost structures" – and these differences lead people to use different strategies. The variability in the descriptions is indicative of the general point that changes in technology can lead to changes in work practices.

This appendix is intended for readers who are curious to understand and compare different models of knowledge sharing. The models featured in this section are drawn from prior art that influenced our own elaboration.[3]

Lederberg: Scientific Discovery Cycle

Figure F-1 presents an unambiguous example of a process for creating, using, and sharing knowledge – a description of scientific discovery through laboratory experimentation. The description is generalized from many examples of laboratory activity. There are several things to notice from a perspective of knowledge. First – it is a cycle with many repeating loops. The outer loop of arrows is a kind of linear process that proceeds from one step to the next. The inner web of arrows is intended to suggest how it is possible to loop back from any step to any previous step as events require.

The twelve-step model refers to a layered community. It specifically calls out an experimental team, which includes students and others who are in training. It includes close colleagues as well as a larger scientific community. In this regard, it makes roughly the same size distinctions as models in the previous section – individuals, work groups, and larger organizations.

The cycle begins with the education and orientation of a researcher to doing scientific work. Each particular experiment begins with establishing a theoretical framework for the experiment. Although the model does not call out the detailed steps for establishing the framework -- finding relevant articles, identifying colleagues, or information foraging – such things would naturally arise at this point in the cycle.

[3] Several more multi-step cycles for sensemaking are documented and compared on the sensemaking white paper web site at http://amberweb.parc.xerox.com/View/Collection-1150.

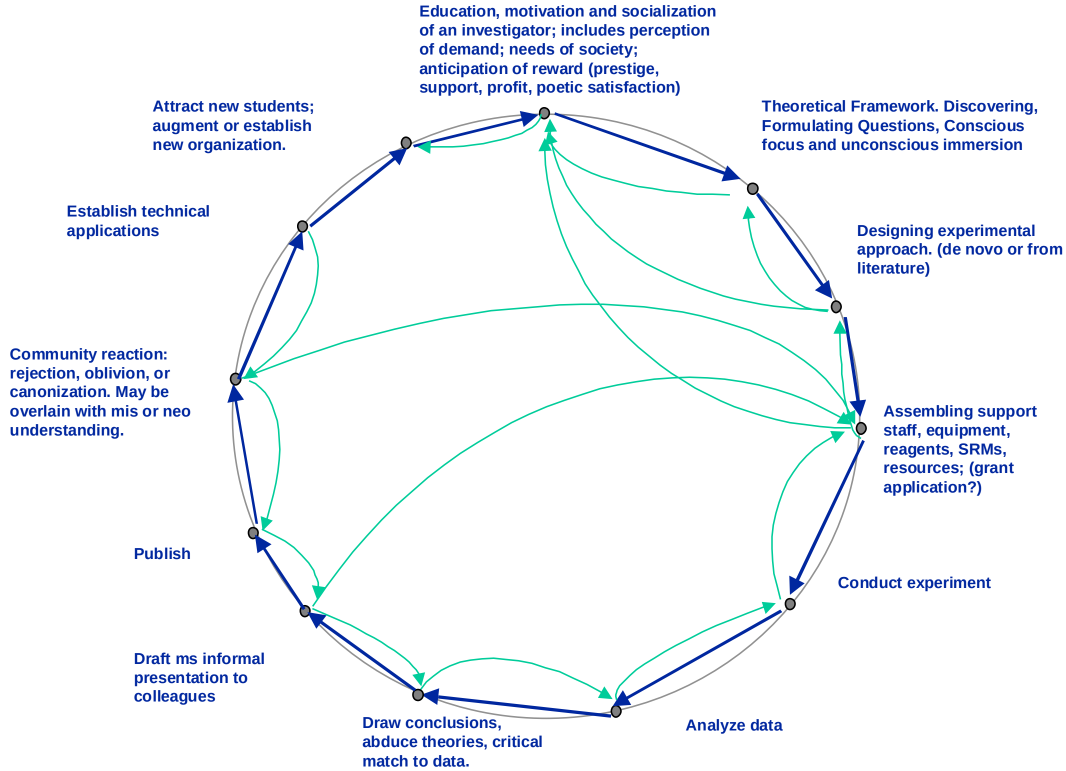

Figure F-1. Twelve-step process for scientific experimentation. Epicycles of scientific discovery. (From Lederberg, J. Preface, *Excitement and Fascination of Science*, Vol 3, Annual Reviews, Inc. Palo Alto, 1989.

Next – the experiment is designed and materials are gathered. This includes all manner of resources including funds and personnel. The next steps – conducting the experiment and analyzing the data – are given short description in the model. Although these steps are exactly the steps most in the public mind about the work of a scientist, Lederberg's rhetorical point in his model may be to emphasize the surrounding work of preparation and publication. Nonetheless, these steps represent the distinction between tacit and explicit knowledge that so dominates the Nonaka and Takeuchi model. Knowledge work is not just "pushing symbols." In this example, there is an actual experiment in a laboratory with actual results of some kind.

The next few steps reach out of the laboratory and into the world of colleagues. This is where the social construction of knowledge in science takes form in Lederberg's model. Conclusions are drawn – influenced by the theoretical framework alluded to earlier. Prior to formal publication, there are informal presentations to colleagues. These steps serve many roles. If there are additional things to try, experimental ambiguities to correct for, or possible errors to correct for – this is an opportunity to loop back to previous steps and fix things up before committing things to the formal public record.

Publishing – an irreversible public act to disclose ideas and discoveries – then follows. Although few scientific papers are read by very many people in actual practice, published papers are potentially open to the entire world. The next step refers to the community reaction. It is worth noting how carefully the publication process is staged in this model. Feedback first comes from informal presentation to colleagues; then follows feedback from the normal peer review process of formal publication; finally, there follows a reaction (if any) from the larger scientific community. The staging serves to assure that the normal accumulation of knowledge in science has an incremental character with infrequent reversals. Multi-stage vetting protects serves not only to protect the reputation of the scientists; it usually avoids wasting the larger community's time from reading or believing inadequately-tested results.

The remaining steps in the cycle close the loop in terms of action and prepare for the next experiment cycle. Establishing technical applications helps to assure funding. The process of publishing and establishing recognition serves to attract new students or colleagues who want to pursue a scientific career.

In comparison with some of the other models of knowledge creation, use, and sharing – Lederberg's model glosses over detailed steps of using explicit knowledge, such as retrieval even "knowledge combination" for scientific hypothesis formation and testing. Nonetheless, one potentially interesting use for this model is as a basis for framing questions about technological leverage and community action. For each step in the model, how could knowledge-management technology or collaboration technology help?

Russel et al: Learning Loop

Like all large manufacturers of technical equipment, the Xerox Corporation has an organization of field service technicians (known as "tech reps") who maintain and service equipment for its customers. Training courses are developed for teaching about new products and also for providing a baseline of experience to new employees. Every year a few thousand employees take these training courses.

In the summer of 1989, a team of experts was organized within the Xerox education division to design a new course on laser printing for training Xerox technicians. Their charge was to create a "generic service training course" to cover a wide range of laser printers including new ones manufactured by Xerox and other companies. Other short courses were intended for field service on particular copiers. The emphasis of the course was to teach basic principles of operation widely applicable across all laser printers.

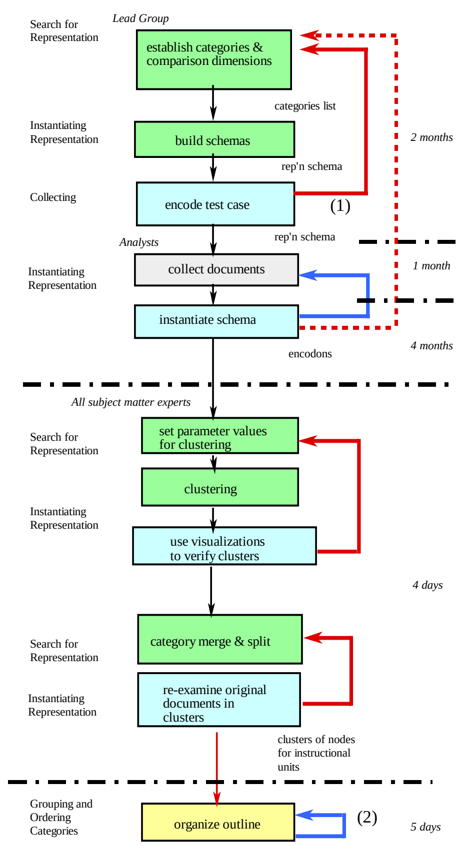

Figure F-2. Eleven-step process model of a collaborative sensemaking task for designing a course about laser printers. (from Russell D.M., Stefik M.J., Card S., and Pirolli, P. The cost structure of sensemaking. *Proceedings of INTERCHI*, Amsterdam, Netherlands (ACM Press), April, 1993).

The required product of the sensemaking was an outline for a course on laser printing. A key part of the task was to understand the commonalties and differences among printers. This involved answering questions such as: what are the basic similarities and differences between laser printers, their operation, ways of failing, and normal functions? What systems should the course cover?

Figure F-2 shows a model of the collaborative sensemaking that was done in this case. The diagram differs from Figure F-1 in that this is not a model of a "generic course" or a "generic experiment," but rather it is a diagram of one specific activity that took place. The figure shows the steps followed to discover clusters of common terms, ideas, subsystems and functions among twenty-one different laser printers. The left column shows what kind of work was being done, while the right hand column shows how long each subtask required.

During the initial two-month period, the lead group identified twenty-one different kinds of laser printers and several different kinds of scanners to be considered. Members of the lead analyst group searched for representations and a vocabulary for describing the devices. Their goal was to provide a common language of description to make it easy to find similarities between devices.

It was not feasible for the lead group to do all of the work. Work was divided among three full-time work groups. During the next five months, the work groups collected and entered information about their assigned printers. To make their results comparable, they encoded their information about printers using the representation schemas provided by the lead group. The results of this encoding effort were a set of databases, one per printer, containing between 300-1000 entries each.

The next phase used computers to compare the different printers, identifying important similarities and differences. Because the printers were designed and documented by different companies, there were substantial variations in the way printers were described. Each printer had subsystems for moving paper, scanning images, scaling images, and so on. Different printers used different designs and consequently had different capabilities, modes of failure, and so on.

The output of the second phase was a set of common categories from each of the databases. These clusters were then used by the sensemakers in the third phase to identify the topics and concepts for the course. By examining the sources on which the clusters were based, the groups identified the educational concepts behind the clusters.

The last phase was to create an outline by organizing the concepts found in the analysis. An outline was represented by relational links that sequenced the educational concepts. The group considered several alternative course outlines. These were compared by a viewing program.

In comparison with our other cyclic models for knowledge creation, there are several things to notice. Like several of the other models, this one involves multiple communities. There was a lead community of teachers and analysts, a larger community of three groups of printer experts, and finally, the community of technicians that would take the developed course.

Every cyclical model has a focus and purpose. This model was created to illustrate the use of external representations for tasks where the information is too voluminous to keep in mind all at once, the variation in the cost structures for reasoning from different representations, and the behavioral pattern involving a learning loop in the development of effective representations. Figure F-3 illustrates the elements of a generic learning loop.

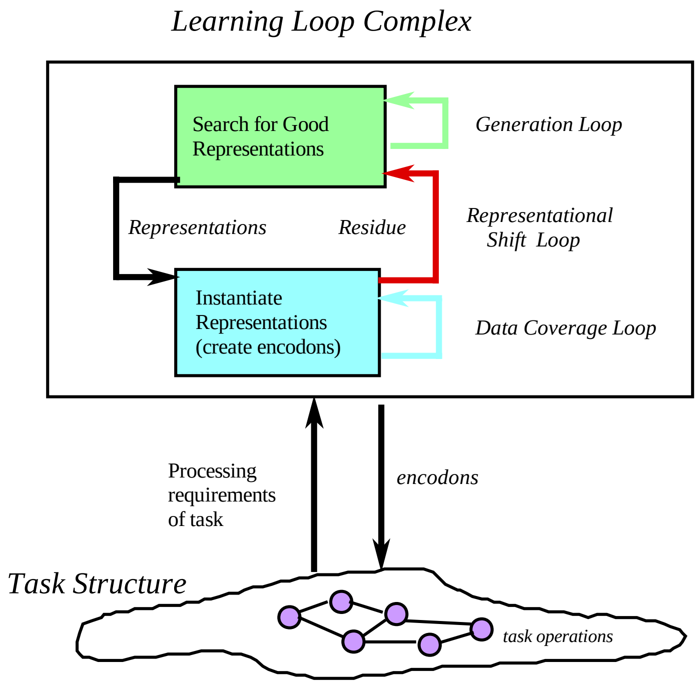

Figure F-3. Two-step learning loop model. (from Russell D.M., Stefik M.J., Card S., and Pirolli, P. The cost structure of sensemaking. *Proceedings of INTERCHI,* Amsterdam, Netherlands (ACM Press), April 1993.

Like many of the other cyclic knowledge processes, the process in Figure F-2 included a corpus. One corpus was an existing corpus of documentation about various printers. The process also created a corpus of comparable documents – the schemas – though an externalization process. In one case, the documents about a Japanese-manufactured printer were created through a process of reverse-engineering on an actual printer from a competitor.

Another interesting feature of the process in Figure F-2 is that there was extensive use of computers in the sensemaking task. Beyond the typical use of computers for search and retrieval, computers were crucial for the creation of schemas and for a massive clustering task for aligning terminology. They incorporated algorithms and dictionaries for recognizing common constructions in the schema. They combined text processing algorithms (such as for stemming) with the use of structural matching on the schemas.

Stefik: Sensemaking by Teams of Analysts

Figure F-4 presents another model of collaborative sensemaking. This model was developed in 1997 after a series of interviews with analysts at various intelligence-gathering organizations.

Like the other models of cyclic knowledge-creating processes, this one involves multiple communities. There is a community (called the “customer”) that wants an intelligence report on a topic, there is a management team, there is a task force of specialists that is recruited to research and write the report, and there is a small community of reviewers that checks over the report before it is delivered. The twelve-step model in figure F-4 is a generic model. Actual cases and actual teams vary. Some reports are created over a six month time frame and others over a few hours.

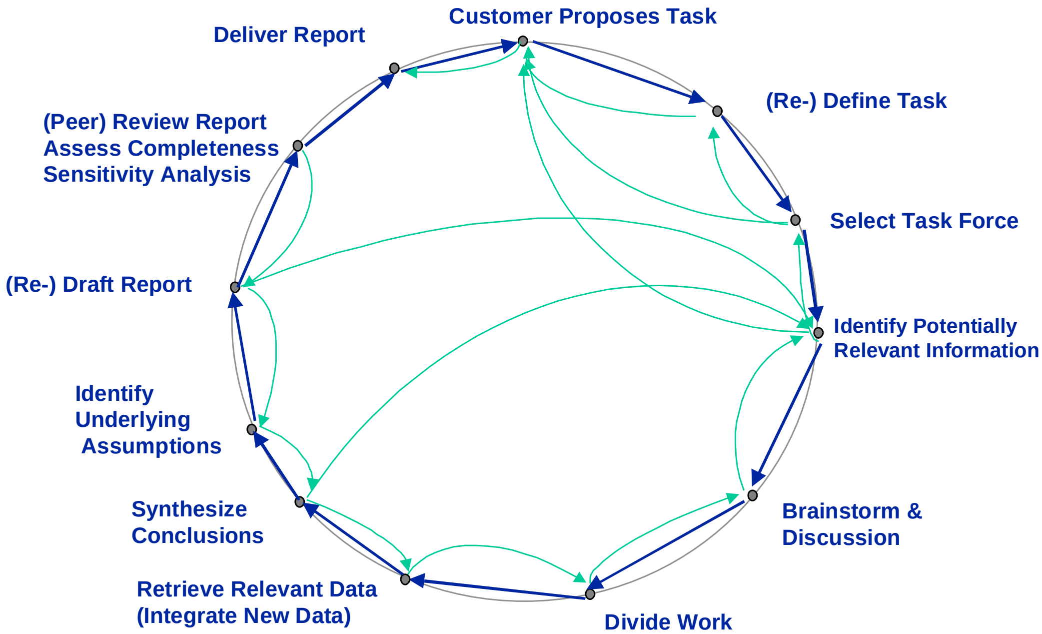

Figure F-4. Twelve-step model for a collaborative sensemaking task for intelligence analysts (from Focusing the Light: Making Sense in the Information Explosion in Stefik, M., *The Internet Edge: Social, Technical, and Legal Challenges for a Networked World,* Cambridge: The MIT Press, 1999).

There are some things to notice relative to the other cyclical knowledge sharing models. One difference in this model is its emphasis on organizational steps – selecting a task force, dividing the work, and so on. These steps surface in this model because of the orientation of the intelligence teams to do their work efficiently – where repeated collaboration by a group of specialists is a routine part of the intelligence community culture.

Another interesting aspect to the cycle is the distinction made between "potentially relevant" information and retrieval of "relevant data". This difference often shows up in models that show detailed steps in information foraging. The sources of information used in intelligence gathering vary from published information (domestic and foreign), intercepted communications, agent reports, satellite photography, and other sources. This is much different from the use of vetted scientific literature in that data vary in their quality and reliability. Indeed, some sources of data may be deliberately misleading in a disinformation exercise.

Another interested aspect of the model in Figure F-4 is the emphasis on steps for managing uncertainty. For example, there are steps for identifying underlying assumptions and for preparing alternative interpretations.

CIA: Intelligence Analysis

Intelligence analysts have themselves studied the intelligence process and put out their own model (Fig. F-5). This model of the sensemaking process comes from an internal report surveying the means by which a particular technology, information visualization, could aid the process. The analysts divide their work into

(1). Collection tasking,
(2). Data monitoring,
(3). Interpretation and analysis,
(4). Drafting/Editing, and
(5). Consumer support

But the important point of the diagram is that one of the analysts' real concerns is the time cost of the different steps. The black line is an estimate of the current situation at the time of the report. Because of the volumes of data in involved (after the end of the cold war, there were a wider set of topics of interest, but fewer analysts), analysts spent most time on data monitoring and report writing, but had little time for interpretation and analysis (as indicated by the black line). Their desire was to move in the direction of the arrows so that most time was spent on interpretation and analysis and less time on data monitoring or report writing as indicated by the red line.

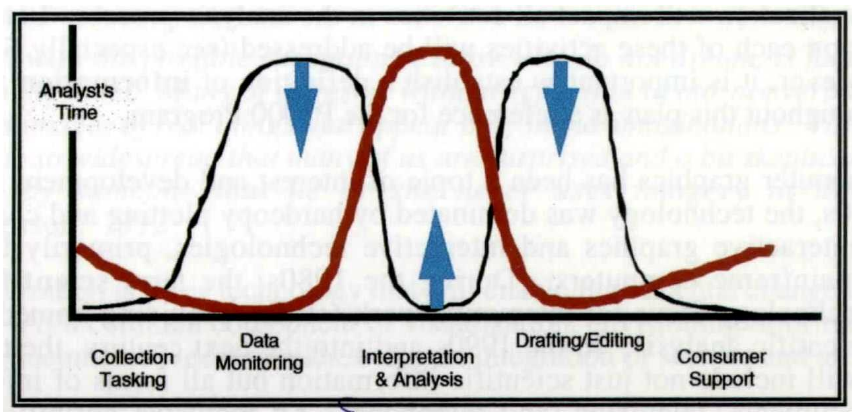

Figure F-5. Five-step model for intelligence analysis (from: P1000 Science and Technology Strategy for Information Visualization: A Roadmap to Provide Visualization Technology Broadly Within the Intelligence Community, Version 2, (CIA) 16 September 1996.

The significant point is that the time for sensemaking operations matters. Other costs besides time also matter. The goal is not just the creation or sharing of knowledge, but the creation or sharing of knowledge at maximum speed and minimum cost. Time or other costs can be expected to trade against quality according to the application. The steps in this model can be mapped more or less into the Stefik model. They are just at a higher level of aggregation. Notice that as in the Stefik model some steps refer to organizational tasks. There is the step of organizing who is responsible for acquiring the collection, as in the Stefik model. An interesting organizational step that is not in the Stefik model is after-report consumer support. This step is particularly interesting because it implies an interactive marshalling of both documentary and human knowledge sharing.

Card, Pirolli, Mackinlay: Knowledge Crystallization

The next model starts with the observation that the sort of tasks studied in the Russell et al and the CIA models above are commonly found in many settings. These tasks are characterized by lots of information, the task itself is ill-structured, interpretation is required, but there is a well-defined output. Examples are buying a printer, writing a newsletter, and business intelligence. Pirolli and Card (1999), have called these "knowledge crystallization tasks." These tasks seem to have steps for (1) foraging for information, (2) sensemaking in the sense of the Russell et al learning loop, and (3) the creation of a knowledge product, which might be a report, a briefing, or even just a decision (Figure F-6). Notice that these steps are more or less the same as the steps in the CIA model of F-7, although derived independently. In this model, the term sensemaking is used for the knowledge creation component, whereas in this report it is used for the more general process.

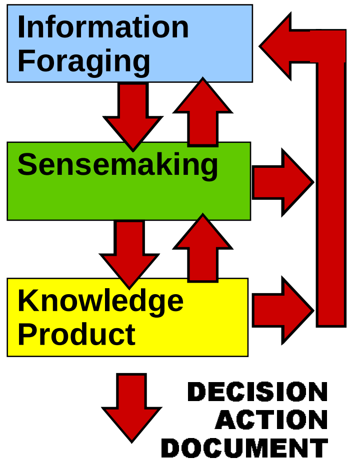

Figure F-6. Three-step model for knowledge crystallization. (from Card, Stuart K., Mackinlay, Jock D., & Shneiderman, Ben. *Readings in Information Visualization: Using Vision to Think.* San Francisco, Morgan Kaufmann Publishing, 1999).

The foraging component of the model has been the subject of several studies. The result has been an "information foraging theory," based on a mathematical theory in biology concerning how animals make choices in seeking food based on maximizing the ratio between gain and cost. By emphasizing benefit per unit cost, this theory gives a different formulation to the information retrieval problem than the usual precision/recall analysis. Applying the theory to information seeking in humans suggests a number of provocative concepts, such as the notion that information is patchy, that information items have different values and costs, that humans are often information rate of gain maximizers, and that information gatherers will tend to organize their information diets in a certain way (e.g., they may prefer tables of products by features to single articles about a certain product). The studies show that instead of just searching for information, users engage in a multi-step process, giving cursory examination at high speed to many items before examining fewer items more seriously. An interesting concept that comes out of the theory is the notion of "information scent." Information scent is some perceptual indication to the user that taking a particular route through an information space is, on the average, likely to cause the user to find information in which the user has a high degree of interest. For example, color or size might be used to highlight those regions of a website displayed as a hyperbolic tree, likely to be interesting. Models allow predictions of how important such a cue is to performance.

Figure F-7 is an extended version of this model used to discuss information visualization. It also lists verbs relevant to the use of information visualization for sensemaking. This version breaks the sensemaking component of Figure F-6 into the search for schema and instantiate schema of the Russell et al model and explicitly adds problem solving to emphasize that sensemaking is directed toward some task and may be enmeshed in the operations of solving the task. To see the relationship of sensemaking to problem solving consider the example in Bertin's 1981 book *Graphics and Graphic Information Processing*. Bertin describes a manager who wants to improve he hotel's importance. He shows how by graphical manipulation of occupancy data, the manager can discover that his hotel actually has distinct seasons, each with differing clientele who use the hotel for different reasons. This enables the manager to design a promotion and pricing program to improve the performance. His analysis is clearly problem solving—how to improve the performance of his hotel, but in the process of problem solving he uses the learning loop operations of searching for a schema for his data and instantiating that schema. These effectively change the *representation* of the problem and hence the operators of the problem-solving. The graphical representation provides external representations of the problem enabling certain visual problem solving operators and enlarging the working memory available for solving the problem.

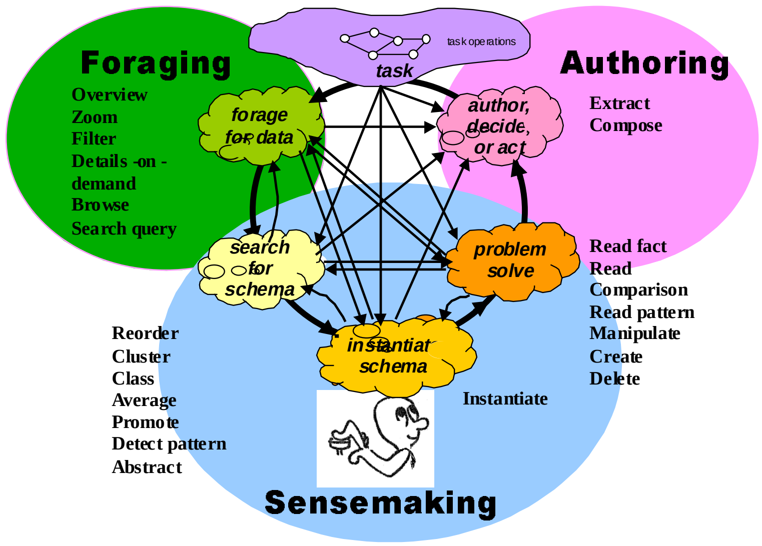

Figure F-7. Five-step model for knowledge crystallization. (From Card, Stuart K., Mackinlay, Jock D., & Shneiderman, Ben. *Readings in Information Visualization: Using Vision to Think.* San Francisco, Morgan Kaufmann Publishing, 1999..)

Unlike the previous models, this model does not call out any explicit operations to organize a collaborative process. Indeed, the communities involved are implicit in this model; they can be inferred from knowing that writing is for somebody or that information foraging is done over documents that somebody else may have written.

Paepcke: Using a Digital Library

The term "digital library" has emerged in the past few years to refer to the development of online collections of materials and services. Research on digital libraries is sponsored jointly by several federal agencies, including the National Science Foundation, the Defense Advanced Research Projects Agency, and the National Library of Medicine. The research funding for libraries has been extended in the later rounds of funding to go beyond the development of technology and to include topics involving collaborative learning environments and community access.

Figure F-8 comes from a paper written in the first round of work on digital libraries. Like the three-step model of the previous section, the verbs focus on largely document-centric actions and groups and people (beyond the primary actor) are represented implicitly.

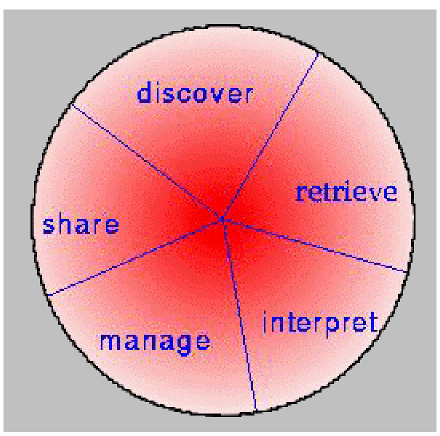

Figure F-8. Five-step knowledge-sharing model for knowledge work using a digital library. (From Paepcke, Andreas. Digital Libraries: Searching is not Enough. What we learned On-Site. *D-Lib Magazine*, May 1996.)

The five-step model in figure F-8 was developed largely to counter the perception that digital libraries were mainly about search. Each of the steps in the cycle represent particular activities by users of a digital library. The discover step emphasizes the notion that resources such as online collections may be a challenge to locate. The next step – retrieve – is the usual focus of digital library research. Even after some documents have been retrieved, work in the digital library is not complete. The interpret step includes finding useful bits of information in the retrieved documents and analyzing them for such things as correlations and trends. The manage step involves collecting and pulling together materials for a particular task and making it ready for sharing by others. Sharing may involve re-representations of information so that it can be understood by others.

In comparison with some of the other models, there are several things to notice. The separation of "discover" and "retrieve" is analogous to the distinction in the intelligence analyst's cycle between "potentially relevant" and "relevant" documents and the multi-step process used in the information foraging model. In contrast with the laser printing course example, this model does not make explicit any of the symbol-manipulation steps of analysis. Indeed, it is not clear how operations such as clustering or annotating or extraction should be distributed under the steps of interpret, manage, or share. Finally, like many of the models, it does not make explicit any of the operations for visualizing information of building decision models based on the information.

Probst: Practical Knowledge Management

Probst's model is a model of knowledge management rather than sensemaking. The model developed from consulting in several industries. The goal is the management of knowledge capital at the enterprise level, for example, to ensure that the research results of the Market Research Department are available to the Product Department. The goal therefore emphasizes knowledge sharing more than knowledge creation. The process begins with a definition of knowledge goals. Then the process proceeds to the identification of the knowledge that exists within the enterprise and the acquisition of additional knowledge sources to supplement what is available locally. Note again the multi-step process as in Stefik's model. Knowledge development is the creation of new knowledge. This is the process of prime concern to the Russell et al and Card et al models, but here it is just another component and one only vaguely defined. Knowledge distribution is explicitly knowledge sharing. Interesting Probst does not mention explicitly either the sharing of documents or mechanisms on contact among people. Knowledge use to Probst means "the productive deployment of organizational knowledge in the production process," which again emphasizes sharing. He mentions knowledge fairs and advisory boards as ways to enhance knowledge use, for example. Two processes that Probst mentions that are not treated in the models reviewed so far are knowledge preservation and knowledge measurement. Knowledge preservation refers to preserving organizational memory. Knowledge measurement means coming up with some indicators to tell how well the process is doing so that it can be managed.

Overall, Probst takes an MIS approach to the management of knowledge. The general process is to be managed like a steady-state process with flows of explicit knowledge from one part of the organization to another and indicators to identify problems that require management.

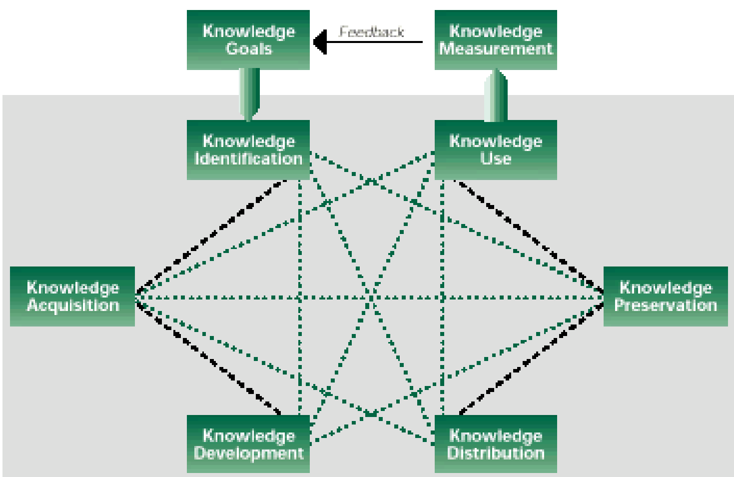

Figure F-9. Eight-step model for knowledge management. (From Probst, Gilbert J.B., Prism, Second Quarter 1998.)

Sprague: Corporate Knowledge Sharing

The knowledge sharing model in figure F-10 was developed by Ralph Sprague in his capacity as a consultant to DSG. Its purpose was to suggest ways that Xerox component technologies might be combined in various ways to addressing customer needs in knowledge sharing.

A key feature of this model is its visual simplicity. The core activity is reduced to three abstract steps – gather, mine and refine. The term gather refers to pulling together knowledge resources from a variety of disparate sources. Particular component functions subsumed under gather include search, extract, store, and convert. The term mine refers to analyzing, understanding, making sense, and structuring knowledge resources. The functions subsumed under mine include lexical analysis, summarize, visualize, categorize, cluster, identify genre, extract meta-data, and identify language. The term refine refers to using and reusing knowledge resources, converting implicit into explicit knowledge, authoring and interpreting knowledge to share with others. Particular component functions subsumed under refine include: author, rewrite, present, revise, refine, redirect, refocus, reinterpret, analogize, synthesize, repurpose, edit, validate, communicate, evaluate, warrant, and improve. Sprague intended this model to describe knowledge dissemination, knowledge management, and knowledge sharing.

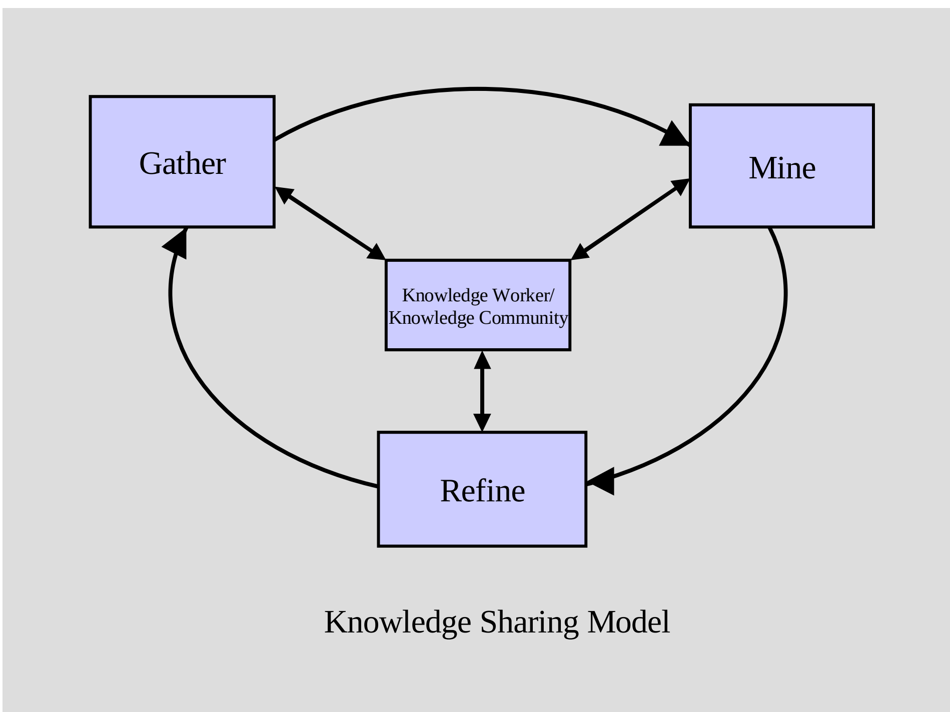

Figure F-10. Three-step knowledge-sharing model from Ralph Sprague.

In comparison with other cyclic knowledge processing models, there are several things to notice. The reduction of the process to three steps makes an abstract pattern easier to see but obscures some other things. Whenever the number of steps in a model is reduced, they become more abstract – moving distinctly different operations under the same category. This makes it more difficult to locate or perceive operations that potentially fit beneath within more than one abstract step. In this case, the terms “mine” and “refine” potentially share operations for extracting and using high-value information – so that mine includes operations like cluster and extract meta-data, but refine includes revise, edit, and evaluate. Similarly operations for searching form information might be found other either gather or mine. Operations for creating knowledge to share seem to fall under refine.

Collapsing both the knowledge worker and the knowledge community to the center square also keeps the diagram simple. It also obscures processes that involve multiple parties at once, such as peer review or document validation where authoring by a knowledge worker and validation by a community are simultaneously represented in the same verb box “refine”.

Nonaka & Takeuchi: Knowledge Spiral

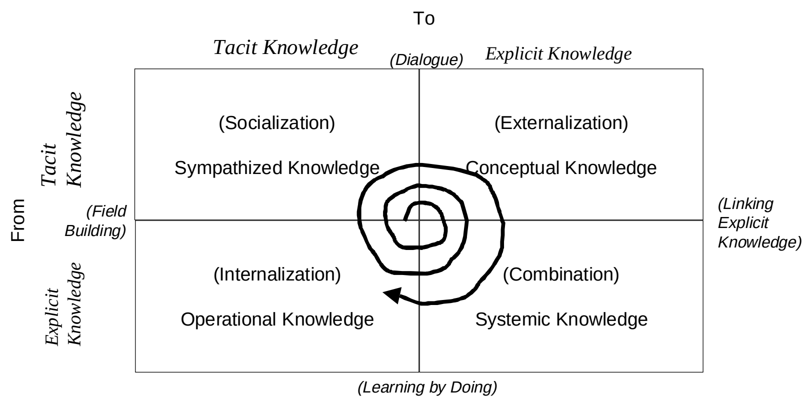

Figure **F-11**. The “SECI” knowledge spiral diagram showing four models of knowledge conversion (adapted from Figure 3-2 (page 62), Figure 3-3 (page 71), and Figure 3-4 (page 72) in Nonaka and Takeuchi, *The Knowledge Creating Company*. Oxford: The Oxford University Press, 1995).

Nonaka & Takeuchi’s model of knowledge sharing was discussed at length in Appendix E, where we introduced their knowledge spiral (Figure F-11). Figure F-11 emphasizes their four modes of knowledge. The steps of knowledge creation they described as a five-phase process (Figure F-12):

[Socialization]	(1) Sharing tacit knowledge
[Externalization]	(2) Creating concepts
[Combination]	(3) Building an archetype
	(4) Justifying concepts
[Internalization]	(5) Cross-leveling knowledge

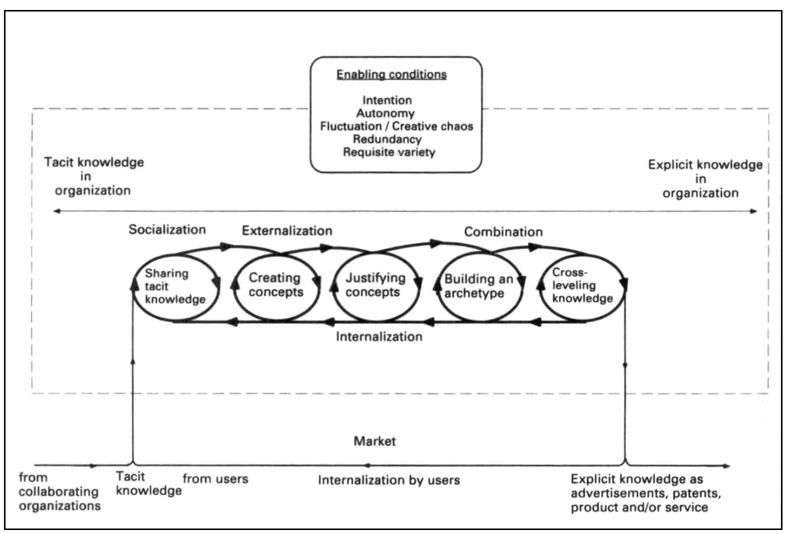

Figure **F-12**. The five-step model of the organizational knowledge-creation process (adapted from Figure 3-9 (page 84) of Nonaka and Takeuchi, *The Knowledge Creating Company*. Oxford: The Oxford University Press, 1995).

In their model, tacit knowledge is held by individuals and so needs to be shared by direct contact. To effect the sharing, the individuals need a "field" or place in which to interact. This field could be a work team or it could be a regular get-together among a group in a bar or coffee house. Tacit knowledge can be externalized by a dialogue among team members in which the knowledge is verbalized and fit into explicit concepts. These concepts are defended and justified as appropriate by vetting them against organizational goals or challenges from other organization members. Justified concepts are converted into a tangible form, such as a new-product prototype or an idealized model of a service. The final phase of the model takes a concept created in one place and shares or applies it in other places. This may start a new cycle of knowledge creation.

Nonaka and Takeuchi's model has a very different focus than the previous models. It is a model of innovation at the company division or enterprise level. The purpose is to explain the corporate processes underlying the innovation success of Japanese companies. As such, it is at a larger scope and coarser grain than many of the other models. It is particularly noteworthy for its concentration on tacit vs. explicit knowledge. Whereas, Lederberg's and Stefik's models recognize organizational steps, this model explicates how people to people knowledge sharing is part of a larger process that involves people to document or people to artifact interactions in other parts of the process.

FX: KR (Knowledge Representation) Model

Fuji Xerox Corporate Research Laboratories have a model of the sensemaking process, the KR model to guide research investment that is a variant of the Nonaka and Takeuchi model. Like Nonaka and Takeuchi, they are strongly concerned with tacit and explicit knowledge. Whereas in the Nonaka and Takeuchi model, documents are part of the explicit world, in the KR model, documents affect both tacit and explicit knowledge. The steps in their model are the movements between squares on the diagram of Figure F-13. Each step points to a class of tools that could be made to augment sensemaking. For example, going from creation to expression involves putting tacit knowledge onto some medium. Going from expression to perception involves transferring media, storing, and various types of conversions. The next step is perceiving knowledge from the medium, and the last step is reorganizing knowledge. Then the cycle is ready to start again.

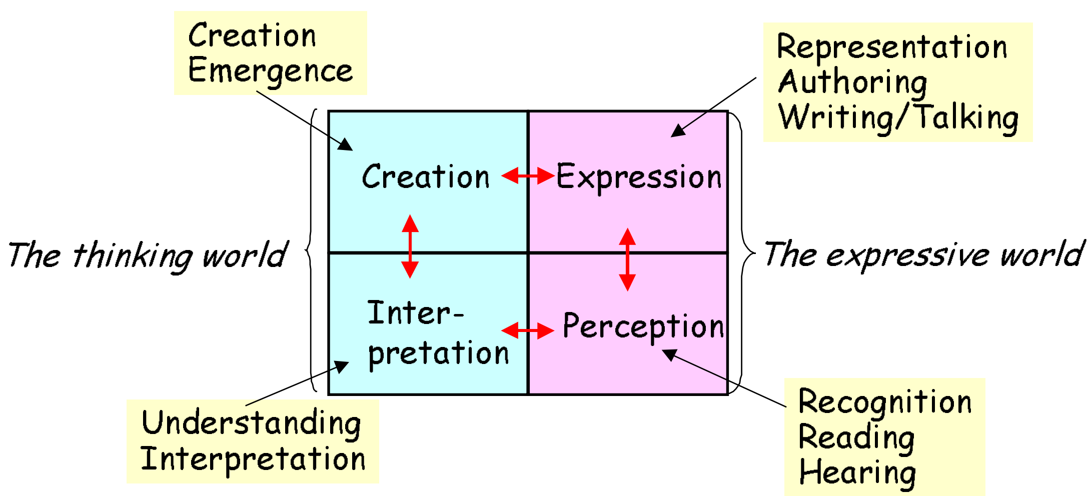

Figure F-13. Creation of knowledge in business processes (from Ueda, M., "Social Space and Knowledge Representation: Knowledge Research Concepts at CRL, Fuji Xerox". Forum talk given at Xerox PARC, August 26, 1999).

Like the Russell et al or Card et al models, this model is interested in transformations of representations, although these are only given at a high level. But interestingly, unlike the Lederberg or Stefik models, there are no explicit references to organization coordination activities. These organizational activities are tacitly assumed, but off stage.

Smoliar: Knowledge Sharing

Smoliar's model (Figure F-14) is another model related to Nonaka and Takeuchi's model. The three main steps are

(1). Figuring things out,
(2). Taking action, and
(3). Sharing with others.

Steps can connect with other steps in any order. Substeps are given for each of the major steps. This model attempts to reduce the complexities of the Nonaka and Takeuchi model to its essential features and to apply it to business processes on a smaller scale. Whereas Nonaka and Takeuchi's model is an innovation process that might take 5 to 10 years and involve thousands of workers, Smoliar's model could apply to individuals or workgroups and take weeks to months.

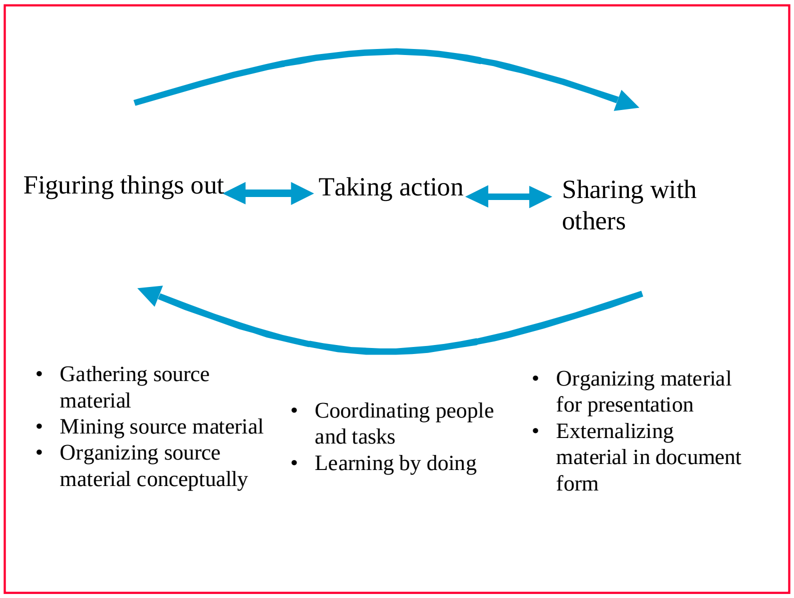

Figure F-14. Three-step model for knowledge sharing. (From Smoliar, Stephen. (personal communication), 1999.)

Although the model does not emphasize the difference between tacit and explicit knowledge to the extent of Nonaka and Takeuchi's or the Fuji CRL model does, it none-the-less contains relevant operators for externalizing information (going from tacit to explicit knowledge) and learning by doing (going from explicit to tacit knowledge). In this model the acquisition activities of gather source material and the learning loop activities of mining and organizing the material conceptually are all included in the figuring things out step. Instead of action being the result of the sensemaking activity, as in the Card et al model, action is seen as a full partner in the activity.

XR&T: Knowledge Operators

Figure F-15 contains an eight-step model for understanding customer needs in knowledge sharing to use in guiding internal Xerox research contracting. These are more in the nature of functions and capabilities than a cyclical model of a knowledge sharing or creation process. Nonetheless, it is interesting that many of the steps have a strong correspondence with previous models. It is the first model we have considered having three types of acquisition: search, monitoring, and capture. Most other models either have a generic operation (e.g., Sprague's gather) or one type (e.g., the CIA model's data monitoring). Its creation operator mixes together knowledge analysis and authoring. But it does have other operators for knowledge creation: organization, transformation, and consolidation. Interestingly for a Xerox model, there do not seem to be operators that transform information back and forth between the physical world (e.g., scanning, printing) other than the event streams under the captures operation. Also there is no mention of organizational issues or representation. Whereas Lederberg's model was the model of a generic task and Russell et al's model was the model of a specific historical task, this model does not reference even a generic task.

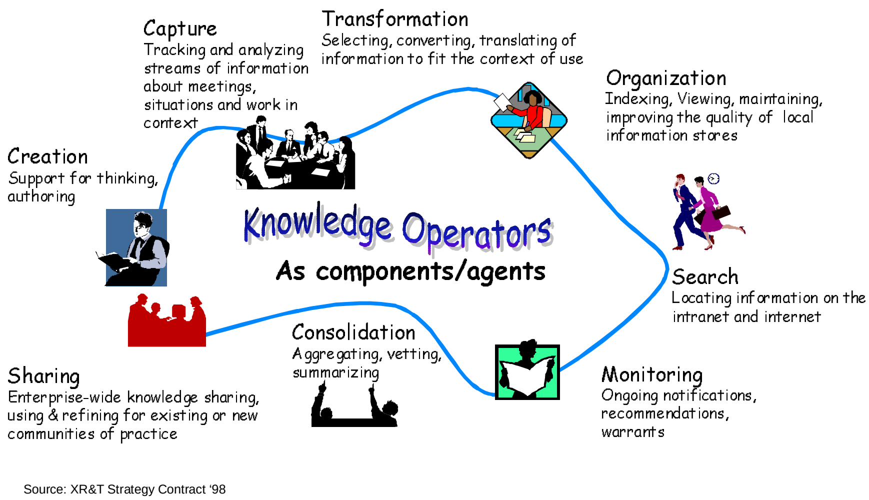

Figure F-15. Eight-step model for understanding customer needs in knowledge sharing. (From Spiteri, P., XR&T Strategy Contract, 1998)

DSG: Knowledge Sharing Vision

The DSG model in Figure F-16 is a variant of the XR&T model, this time used is a customer presentation to give coherence to Xerox's product programs in knowledge sharing line. Creation, organization of information, and consolidation have disappeared as steps. Organizational memory and accessing have appeared in their place. This is the only one of the models reviewed to include accessing (in the security sense of property rights protection) as part of a model of the knowledge sharing process.

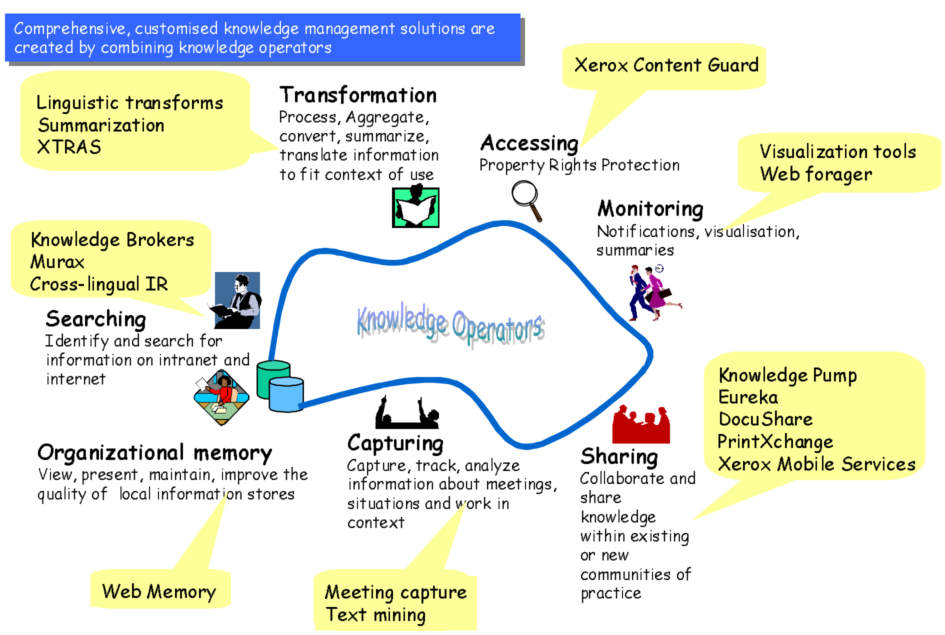

Figure F-16. Seven-step model in vision for knowledge sharing. (From Singh, R., Presentation to Sun, May 1999.)

Comparing the Cyclic Sensemaking Models

The preceding models all are attempts to describe a process whereby knowledge is created or shared. Each is from the perspective of a different context, probably describing a slightly different process, and using descriptions at different levels of granularity. But despite these differences or origin, intent, and description, the models each contain more similar descriptions of the sensemaking cycle than we might have thought. In fact, when we examine them more closely, we find a close concurrence about core processes. We also find interesting variation. Both are significant.

Use

Table F-1 lays the steps in each of these models against each other. Since most models have arrows indicating that the steps may be done in many orders, we have adjusted the order of the steps to make clear their common ground with the other models. Laying the steps out in this way makes visible four common subprocesses that most address. First is the area of use (and context). Knowledge, as we have said, is always about knowledge to some end. This use is part of the goal of knowledge creation or sharing at the beginning and the application o the knowledge at the end. Lederberg, taking the long view, is concerned in his model with the needs of society and theoretical frameworks in establishing goals for his experiments. He is concerned with establishing technical applications for the uses of the new knowledge he creates. Card et al include problem solving as one of their steps to indicate the embedding of knowledge creation as part of its use. Stefik has steps for customer task proposal and task refinement. Probst has explicit steps for knowledge goals and knowledge use, and even proposes a step for the measurement of knowledge as a way of helping to set the goals.

Acquisition

A second subprocess has to do with the acquisition of knowledge. All the models have steps for this subprocess, but the differences are interesting. Essentially five methods of acquiring information or knowledge are described, two about explicit knowledge, three about tacit knowledge:

(1) search,
(2) monitoring,
(3) interacting with people,
(4) capturing from the world,
(5) internalizing from use.

Search itself is complex. Even in Paepcke's library-oriented model, he divides search into discovering the resources and then using these to find information. Pirolli and Card's information foraging model emphasizes the cost structure of the search and methods that can be used to increase information gain per unit time, such as filtering, ordering, physical environment rearrangement, and other forms of enrichment. In other words, the efficiency of search is sensitive to prior investment just as farming investment has higher yields than foraging in the wild for food. Likewise, over-investment reduces returns so there is an issue of balancing preparatory discovery of sources, filtering, planning, etc. against actual assimilation of the materials. Information foraging theory makes clear that monitoring is the flip side of searching, the difference between the lion hunting a gazelle and the spider waiting for a fly to be caught in its web (or between an information retrieval search and an email distribution list). The tradeoff among these methods involves a quantitative tradeoff between the cost of locating an item vs. the number of items (and their information value) that come through the catchment area. Opportunities for machine assistance for are especially great in tasks with large amounts of information and/or deadline pressures.

Acquisition could also come not from assimilation of passive document sources, but from interaction with people who themselves contain the knowledge. Stefik's brainstorm and discussion step is an example.

Several of the models have steps for capturing information directly from the physical world. This is what Lederberg's experiments do. The Fuji CRL model contains a step for perception and the Spiteri and Singh models both contain capture. Capture could be as simple as scanning or more complex, such as making use of video-recorded behavior. Another way of capturing information from the world is Nonaka and Takeuchi's notion of internalization, in which interaction with some product causes the internalization of tacit knowledge. Smoliar's model has a similarly learning by doing step.

Table F-1. Comparison of Descriptions of the Sensemaking Process

	Lederberg (1989)	Russell et al (1993)	Stefik (1999)	CIA (1996)	Card et al (1999)	Paepcke (1996)	Probst (1998)	Sprague (1999)	Nonaka & Takeuchi (1995)	Fuji Xerox CRL (1999)	Fuji Xerox Smoliar (1999)	Xerox XR&T Spiteri (1998)	Xerox DWG Singh (1999)
Coordination	• Assembling support staff, equipment, reagents, SRMs, resources • Attract new students; augment or establish new organization		• Select Task Force • Divide Work	• Collection tasking							- Coordinating people and tasks		
Storing						• Manage	• Knowledge preservation						• Organizational memory • Accessing
Use	• Education, motivation, demand, needs of society, reward • Theoretical framework, Discovering, Formulating Questions, Conscious focus and unconscious immersion • Establish technical applications	• Processing requirements of task	• Customer proposes task • (Re-) Define task	• Consumer support	• Problem solve • (decide, act)		• Knowledge goals • Knowledge measurement • Knowledge use				• Taking action		
Acquisition	• Designing experimental approach • Conduct experiment	• (Collect documents)	• Identify potentially relevant information • Brainstorm and Discussion • Retrieve relevant data (integrate)	• Data monitoring	• Forage for data	• Discover • Retrieve	• Knowledge identification • Knowledge acquisition	• Gather	• Internalization	• Perception - Recognition - Reading - Hearing	- Gathering source material - Learning by doing	• Capture • Search • Monitor	• Capturing • Searching • Monitoring
Sensemaking	• Analyze data • Draw conclusions, adduce theories, critical match to data	• Instantiate Representations • Search for Good Representations	• Synthesize conclusions • Identify underlying conclusions • Assess completeness • Sensitivity analysis	• Interpretation and analysis	• Instantiate schema • Search for schema	• Interpret	• Knowledge development	• Mine • Refine	• Externalization: - Creating concepts - Justifying concepts • Combination: - Building an archetype	• Creation - Creation - Emergence • Interpretation - Understanding - Interpretation	• Figuring things out - Mining source material - Organizing source material conceptually	• Organization • Transformation • Consolidation	• Transformation

Sharing	• Draft ms, Informal presentation to colleagues • Publish • Community reaction	• (Organize outline)	• Draft report • (Peer) Review report • Deliver report	• Drafting / editing	• Author	• Share	• Knowledge distribution	• Knowledge Worker / Knowledge Community?	- Cross-leveling knowledge • Socialization - Sharing tacit knowledge	• Expression - Representation - Authoring - Writing / Talking	• Sharing with others - Organizing material for publication - Externalizing material in document form	• Creation • Sharing	• Share

Sensemaking

Each model has some process by which information or knowledge is turned into new knowledge useful to the target task. The most articulated version of this process is the learning loop of Russell et al, which divides the process into finding a schema and instantiating that schema with information. Information that does not fit the schema is residue, which is used to alter the schema. This view of sensemaking is the basic view informing Card et al and Stefik's models. Lederberg's steps of analyzing data and adducing theories, including critical match to data is roughly the same sort of process. Similar notions appear in the literature, for example, Bertin's (1981) view of how graphics can help problem solving uses a set of graphics transformation to identify and eventually summarize data in terms of a schema. Gell-Mann (1994) describes complexity in terms of the length of a description that can be given. Sensemaking from this point of view is the reduction of complexity by creating a more compact description. The schema represents the concepts used for this description. Smoliar's model, involving mining and organizing, is also similar. A different view comes from Nonaka and Takeuchi and the models that follow their lead. For them, the critical process is bringing tacit knowledge into explicit form, then combinations of this knowledge. Finally, Smoliar's model and Spiteri's model mention organization. This is interesting because it suggests the role of external representations in helping to arrive at schemas for the knowledge

Sharing

All models mention sharing of knowledge in some form. There are two main emphases. One is on sharing information directly with others through interaction. Steps in Lederberg's model involve informal presentation to colleagues and community reaction. Stefik's model includes peer review. Sprague's model involves knowledge communities. Nonaka and Takeuchi's model involves sharing tacit knowledge (e.g., through apprenticeship). Several models just list sharing directly. The other emphasis is on sharing indirectly through the creation of documents. Lederberg's model has steps for publishing. Stefik's model includes drafting and delivering reports. Several other models imply the creation of a document. Such reports would be located through searches or monitoring in an acquisition step. But hybrid strategies are possible in which documents are used to index into people, who then share knowledge directly.

Coordination, Storing

In addition to the above four subprocesses represented in most of the models, several of the models also contain other steps concerning the infrastructure or coordination of the four subprocesses. Knowledge work is sometimes divided into content work and coordination work. Coordination work is the work that must be done in managing the people and resources necessary to do the content work. This work is very relevant to organizations, where much of the work must be done in teams. Lederberg's model, though focused on academic science, is concerned with assembling resources, staff, money, equipment, and students. Implicitly, it is also concerned with the management of these resources. Stefik's model, reflecting the world of a team of analysts, contains steps for selecting a task force and dividing the work among them. The CIA model contains a step for coordinating the locations of information collections for analysis. Smoliar's model also calls out coordination explicitly.

Means of storing and communicating knowledge, that is, concern with the media enabling knowledge accumulation or sharing, is made explicit in a few models. A digital library cannot escape consideration of the storage of knowledge, and so Paepcke's model has a step for managing the collection. Probst's knowledge management model and Singh's model both have steps for preserving knowledge. In addition, Singh's model is the only model to have a step concerned with controlling the access to knowledge. None of the models have an explicit step for communication, although the means and bandwidth of communication are likely to be important enablers of knowledge sharing techniques, for example, for the extent to which knowledge can be accessed anytime or from anywhere.

Differences in Sensemaking Processes

Despite differences in purpose and grain of description, the models of the cyclical sensemaking process have considerable commonality in their description of the process. This commonality is evidence for the existence of a process that can be studied and designed for. As a group, they illuminate parts of this process missed by any particular model. Their differences also show that there are several related but not identical goals for the sensemaking process captured in the different models. One goal is the creation of knowledge and the reduction of complexity through compact description relative to some task. If the goal is to buy a laptop, then the processes is to find information about laptops, reorganize it into increasingly more compact forms (e.g., a product by feature matrix then an even more compact decision diagram), extracting the crucial insight that simplifies a decision. A second goal is to take knowledge available in one part of the organization and move it over to another part of the organization where is can be used. If the goal is to repair a copier and some technician has discovered a new sort of problem with its, then the process is to capture that in a way that is quickly assessable and assimilable by another technician elsewhere. A third goal is making tacit knowledge explicit so that it can be more easily used. If the goal is to develop a new product, the process is to use existing products, developing a tacit knowledge "feel" for them, conceptualizing new approaches and creating artifacts to make explicit the tacit knowledge, then to turn the explicit knowledge into products. Use of these products will induce more tacit knowledge, which starts the cycle anew.

In the communication model used in this white paper, we were interested in each of these goals. Acquire and Share occur in our model explicitly. Act is a combination of Sensemaking and Use. This makes Coordination and Storing (repositories) less visible, which is why we explicitly mention them here.