Chapter II
The Social Context of Knowledge

Daniel Memmi
Université du Québec à Montréal, Canada

ABSTRACT

Information and knowledge have become a crucial resource in our knowledge-based, computer-mediated economy. But knowledge is primarily a social phenomenon, on which computer processing has had only a limited impact so far, in spite of impressive advances. In this context have recently appeared various collaborative systems that promise to give access to socially situated information. We argue that a prior analysis of the social context is necessary for a better understanding of the whole domain of collaborative software. We will examine the variety and functions of information in modern society, where collaborative information management is now the dominant type of occupation. In fact, real information is much more complex than its usual technical sense: one should distinguish between information and knowledge, as well as between explicit and tacit knowledge. Because of the notable importance of tacit knowledge, social networks are indispensable in practice for locating relevant information. We then propose a typology of collaborative software, distinguishing between explicit communities supported by groupware systems, task-oriented communities organized around a common data structure, and implicit links exploited by collaborative filtering and social information retrieval. The latter approach is usually implemented by virtually grouping similar users, but there exist many possible variants. Yet much remains to be done by extracting, formalizing, and exploiting implicit social links.

INTRODUCTION

The development of computers and electronic networks has considerably advanced our society's capacity for information processing, and the very scale of this global phenomenon raises quite a few questions. Yet electronic data processing is by now so pervasive in advanced societies that it is easy to forget how recent it all is: computer science started about the time of World War II, but personal computers, the Internet, and the Web only go back a couple of decades in spite of their explosive progress.

As a matter of fact, information processing (i.e., the collection, creation, elaboration, and transmission of useful knowledge) has been around for as long as human history, and has become more and more important with the advent of modern bureaucratic industrial states two centuries ago. Recent technological developments take place within this social framework, which determines their shape, usage, and direction. The interaction between pre-existing social practices and new technologies is then an obvious issue to consider.

So how do human beings and organizations process information in today's technological, computer-mediated environment? How do they interact with each other through electronic networks? How can they put recent technical advances to the best possible use? And what future directions can be foreseen? To try and answer such questions, it would be useful to first analyze human information processing in more detail.

The classical approach, prevalent notably in cognitive psychology, has been to focus on individual information processing capabilities (Neisser, 1967; Mandler, 1985). A body of studies on perception, learning, recall, association and inference, and so forth has been performed on individual subjects in laboratory conditions. Much has been learned in this way on human information processing: for example our limited short-term memory, perceptual schemas, associative recall, probabilistic learning, and inference mechanisms are by now fairly well-established findings.

These studies have however been increasingly criticized for dealing mostly with isolated subjects performing artificial tasks in unrealistic ("non-ecological") environments. One has seen in the past 20 years a gradual shift to the study of situated and collective cognition. There has been more emphasis so far on physically situated rather than socially situated behavior, but the general trend is clear (Clark, 1998; Harnad & Dror, 2006).

Researchers in this growing movement try to understand how human beings perform tasks and solve problems in real physical and social situations. What they may lose in precision and experimental control, they hope to gain in scope and realism. Such an approach seems more relevant to the complex socio-technical environment in which human information processing must take place today.

The recent emergence of virtual communities which has been made possible by the Internet and other electronic networks is also a phenomenon worth investigating. These communities constitute a novel, computer-mediated form of social grouping, combining in variable proportion traditional social relations with more functional, goal-oriented features. Virtual communities should be studied as a collective entity rather than a mere collection of individual participants (Kollock & Smith, 1999; Rheingold, 2000; Memmi, 2006).

Understanding the social and technical context of individual information processing is important for several reasons. Beside the inherent interest of this subject, studying the way human beings use their social skills and social networks to acquire relevant information would help develop better information retrieval systems. As a matter of fact, there has recently appeared a variety of collaborative software systems inspired by human task-oriented social interactions.

Even if socially situated knowledge management cannot be totally reproduced with computers, software systems can be designed to borrow from the most pertinent aspects of human collective processing. Such distributed systems will also fit better the manner in which human beings naturally operate and solve tasks within society, and should thus prove easier to use. More generally, we will see how studying the role and use of knowledge in our society may prove useful to software designers and developers.

Our main thesis will be that information retrieval and information management in general should profit greatly from the study of socially situated information processing by human beings. This text intends to survey fundamental issues more than recent technical solutions. Understanding the nature and functions of knowledge in society appears necessary for long-term advances. We thus hope to bring some order to a fairly diverse range of proposals and to point to new research directions.

In this chapter we will therefore describe in turn: (1) the social and economic context of human information processing, (2) the nature and varieties of knowledge as well as its social pathways, and (3) various technical methods that have been devised to make use of the social aspects of human information processing.

We will resort in rather eclectic fashion to several disciplines, notably cognitive psychology, structural sociology, economics, management theory, and of course computer science and software design. But our main goal throughout will be to replace present work in collaborative software systems within the context of socially situated human cognition.

SOCIAL CONTEXT

We will start by showing more precisely how information processing can be seen as socially situated. This point of view will also have concrete technical consequences for the design of software systems.

The Social Import of Information

Far from being a purely individual phenomenon, information is intimately interwoven with the social and economic fabric of human groups. Social life is not possible without a constant exchange of information within groups and organizations. Because the social functions of information are still largely underestimated, they deserve much more emphasis.

In this respect, one might want to make a distinction between raw information and knowledge acquired by human beings. Whereas information could be formulated objectively, knowledge is inherently a cognitive phenomenon and knowledge acquisition is a complex process. This distinction will prove useful later on, but following common usage, we will use the two terms more or less interchangeably for the time being.

Information can be defined in various ways, notably in probabilistic terms, but its practical function is to reduce uncertainty and to answer questions, allowing us to avoid dangers, fulfill goals, solve problems, and plan for the future. Information obviously has a biological survival function: all life forms, from insects to mammals, need information about their environment in order to find food and mates, avoid predators, and seek appropriate living conditions.

Information comes from the environment, be it physical, biological, or social. But most of our human environment is in fact a social one. Like most primates and many mammals, mankind is a highly social species and social situations are an integral art of our daily life. In modern urban society, moreover, we live in a mostly artificial, man-made environment replete with social functions and meanings.

As I look out of my window while writing this, I can see mostly buildings, whether residential or commercial, cars and traffic, and people walking by, many of them probably to or from work. This physical urban environment is actually a social environment. In my home, radio, television, telephone, fax machine, the Internet, papers, and magazines keep me informed about the larger world. The workplace is also a place of high informational density, where information is constantly exchanged and elaborated upon so as to perform complex social tasks.

As an ordinary member of modern society, I am extremely well connected with my environment, which turns out to be a highly social one. We could indeed be defined as social beings by the rich pattern of informational interactions we regularly maintain with our surroundings. Sociologists and anthropologists have often remarked that social cohesion is both ensured and demonstrated by regular exchanges of goods and services (Mauss, 1923), and information most probably plays a similar role, from office gossip to the Internet.

More concretely, a constant flow of information is obviously necessary for the coordination of social activities. This is true at all levels of social organization, from small business firms to the highest levels of government. The more complex the social and economic organization, the more important coordination activities become (Mintzberg, 1979). At the same time, communication is often highly ritualized and the practical functions of information blend insensibly with its cohesive role. For instance, office memos carry useful information while reaffirming organizational structure.

Another factor to consider is the economic value of information. It is a fact that is not yet sufficiently recognized, that information (or more accurately, human knowledge) has been the dominant source of growth and wealth in advanced societies for more than half a century. Investment in education, research and development, management, and other intangible factors has now overtaken investment in physical assets both in value and contribution to economic productivity and growth (Kendrick, 1994).

It is knowledge, and not physical investment in plants and machines, that is now the driving force in our post-industrial society. Knowledge-based domains such as electronics, computers and data processing, aeronautics, aerospace, biotechnology, and pharmaceutical companies clearly are the most dynamic, productive, wealthiest, and fastest-growing sector of the economy. And this is not a temporary phenomenon, but a solid long-term trend.

In short, most of our information originates from social situations, fulfills social and economic functions, and knowledge has become crucially important in a modern economy. Information must therefore be considered within its social context in order to really understand its functions and uses, and information processing techniques should also be seen in this context.

Toward the Information Society

Information processing is then not only an individual activity, it is the blood flow that keeps our societies running and prospering. Knowledge-intensive occupations and organizations have accordingly become more and more important: research and education, engineering, high-tech companies, consulting activities, law firms, financial services, health care, and so forth. A whole class of "knowledge workers" has emerged whose jobs consist mostly of handling and elaborating information on a daily basis (Drucker, 1992).

Knowledge workers not only handle information, but also create, transform, acquire and store, transmit and exchange, apply, and teach all forms of knowledge. They usually do so in a highly collaborative manner. For various reasons, the management of knowledge in many modern organizations tends to be a collective, distributed activity.

Information being an intangible asset, it is easily duplicated (especially with electronic techniques) and lends itself to cumulative development, a fact that encourages its dissemination and collective production and use. Network effects reinforce this tendency: it is often all the more advantageous to use an informational product (such as software) when it has many more users. Knowledge workers value collaborative work accordingly.

So information and knowledge are used mainly in social situations, even when processed by individuals. Information processing in real life is socially situated, and individual uses are secondary and derived from social goals. Not only most of the information we handle fulfills social functions, it is also managed collectively. As a consequence, useful or necessary information is to be found as much (if not more) in social circles as in libraries or databases.

The growing importance of information in a knowledge-oriented society has also been considerably accelerated by the recent developments in electronic information processing. Social and professional changes have gone hand in hand with technological advances—progress in one area taking place in synergy with evolutions in another. What is striking is not only the enormous increase in computing power available on the job in many professions, but its distributed character and the connectivity between individual computers.

Centralized mainframes have been replaced by cohorts of ubiquitous personal computers, and everybody is now connected to everybody and everything else by the Internet. More than the arrival of computers, the prominent fact of our time is the advent and rapid spread of electronic networks. They have made possible an amazing acceleration in the speed and quantity of information exchanged in our society.

At the same time, and this is of course no coincidence, sociologists have noticed an evolution toward a "network society" of loose, temporary, flexible relationships (Castells, 1996; Wellman, 1999). Instead of staying within closed groups, many social actors tend to shift from one connection to another as required by a different tasks or objectives. Traditional organizations give way to more flexible arrangements, and the Internet has proven to be the obvious tool to switch between diverse social links, regardless of time and distance.

The powerful conjunction between social changes and technological advances makes the information flow ever more important and significant. A network society can only function by constantly exchanging information, and a network structure is the appropriate organization for an information society (Shapiro & Varian, 1999). Computers, electronic networks, urban life, as well as rapid transit systems provide the technical infrastructure for this kind of social life.

The recent movement known as "Web 2.0" is characteristic of this socio-technical evolution (O'Reilly, 2005). This encompasses a loose collection of software tools and applications fostering social relations and collaborative work on the Internet. In this approach, the Web is seen as a platform for various social communication applications. Such tools accelerate even more the present trend toward a network society.

One may speculate about the causes and effects in this global evolution, and whether social changes or technical advances have been the dominant factor. But is clear that changes in different areas have reinforced one another, forming a coherent system that is reshaping our whole society. *Collective, distributed knowledge processing is now the prototypical occupation in today's information society.*

Technical Consequences

Because of these various social, cultural, and technical changes, human information processing is thus becoming more and more a collective, collaborative activity. Information can still be accessed individually in books, libraries, databases, or on the Web, but the sheer volume of accessible information makes social guidance or filtering practically inevitable. And more often than not, pertinent information resides partly in people's heads or expertise, and not in explicit documents, whether physical or electronic.

The constantly increasing complexity of tasks and problems makes it necessary to first locate the right person in order to perform a given task or solve a problem, and this requires a particular kind of social expertise. The diversity and dispersion of information, the fact that various sources of information must be put together and reformulated to become relevant, usually require some human collaboration. And one cannot stay within a small familiar circle of close colleagues or acquaintances to find all the required answers.

The information needed is often to be found somewhere within or by way of a larger social network of professionally related people. These networks may be formal (employees of a firm, professional organizations) or informal (personal address book, casual professional contacts), but they must be searched to locate information or knowledge that could not be found otherwise. Information retrieval thus becomes a social problem.

This means that *the whole domain of information retrieval should be fundamentally rethought in the light of the social nature of human knowledge.* Information has too often been thought of as some kind of objective material, detached from its social environment and use. This simplistic approach has probably made possible the first developments of information retrieval techniques, but one will not advance beyond those techniques without considering the ways in which human beings process knowledge in society.

Classical information retrieval has dealt fairly successfully with how to represent texts, how to evaluate semantic proximity, and how to index and retrieve documents efficiently (Salton & McGill, 1983; Baeza-Yates, 1999; Manning & Schütze, 1999). But new questions should now be considered: Who is the most likely person able to answer a request? How can we find this person quickly and efficiently? How can one represent people and social links? How can one use social expertise and distributed knowledge to recommend or filter documents?

This is the general setting in which must be seen the recent developments of collaborative software, social filtering, recommendation systems, and similar work. The present interest in such systems is no accident, but rather a sign of our times. We will describe below concrete technical approaches, but we must discuss beforehand the variety of knowledge forms involved in social processes.

NATURE OF KNOWLEDGE

We will now analyze in more detail how human beings manage information in real social situations and how they handle different varieties of knowledge.

A Simple Example

To illustrate this discussion, let us start with a concrete example. Let us suppose your organization has asked you to write or prepare a report on free and open source software, a subject you might not know too well. So how would you go about it? The first step might be to visit a library, looking up the computer science section directly, or consulting the catalog. But there just are not many books on the subject, they are still unlikely to be found in a public library, and relevant articles are scattered among so many journals.

Nowadays, your first reflex would probably be to use a search engine instead, to find references on the Web. But you will then be flooded with a profusion of references, of various relevance and quality. Which ones should you read and use? Can you trust these references to reflect a consensus in the domain? Or are they unorthodox divagations? Should you start with this long official report by a reputable organization or does this unassuming Web page offer a decent summary?

At this point, you will probably try to locate a knowledgeable colleague or acquaintance, somebody who could give you a leg up by recommending a few basic references or by inspecting your first list of references. He or she might also explain how to best exploit those sources, and tell you things about the domain that are not easily found in written documents. And if he happens to be a practitioner of open software, the discussion could become quite lively and really interesting…

He might assert, for instance, that popular discussions on the subject tend toward wishful thinking and unsubstantiated ideological claims. He could, however, recommend two or three studies in which one can find the real professional status and economic support of free software developers. This would probably help you write a better, more informed report on the matter.

But how can you be sure your colleague really knows what he is talking about? Well, you can never be totally sure (until you become an expert yourself). But if he has been recommended by close colleagues of yours, if he has been involved in this subject for years, if he belongs to an association dealing with free software, you might be reasonably confident. If he does not belong to your organization, you will probably try to evaluate somehow the competence of his organization and his own standing, before you trust his advice.

And how does one locate the right person? In most cases, this is done simply by asking personal acquaintances deemed to be closer than you to the information required. For instance, if you do not know anybody working on free software, you might ask a software engineer or your system manager to recommend somebody else to consult. By following two or three such links, you will quickly find a knowledgeable expert.

Such a simple strategy has been shown to be fairly efficient. In a well-known experiment, people in the United States were asked to forward a letter through personal acquaintances only, in order to reach a target person whose occupation was mentioned, but not the exact address (Travers & Milgram, 1969). People were instructed to hand over the letter to somebody they thought closer to the target, geographically or professionally, and the process would be repeated from one person to the next. Not all letters reached the final target, but those that arrived at their destinations took no more than five steps on average. This is a good example of the "small-world" phenomenon (Watts, 1999).

We often use a similar strategy when looking for preliminary information on a subject we do not know much about yet. In other words, *we first perform a kind of social look-up in order to access relevant information or knowledge.*

This fairly straightforward example illustrates some of the points we will now elaborate upon: the difficulty for an individual to manage socially distributed information on his own, the need for social guidance, the problem of trust, how help can be found by exploiting social links, the importance of tacit knowledge and personal expertise, the role and structure of social groups, and so forth. The issue will then be how to formalize and exploit these social phenomena.

The well-known Internet bookseller Amazon.com offers prospective buyers a simplified version of such social guidance. When a book on a given subject is found through Amazon's search engine, the system displays a list of ratings and comments on this book by former buyers and users. This is still very crude (the trustworthiness of the ratings is questionable), but this is an effort to help individual online buyers with social advice.

Varieties of Knowledge

Yet to fully understand human information processing, it must be realized that we are actually dealing with different forms of information or knowledge which are managed in different ways. To begin with, one should distinguish between *information* and *knowledge,* a distinction we have glossed over so far. Although usage varies somewhat, information is basically the raw material of information processing, whereas knowledge has been acquired by human beings through a learning process.

Information can be found in physical form, for instance in written documents, databases, images, and recordings. Information may be defined objectively in probabilistic terms according to information theory: the quantity of information contained in a message is inversely proportional to (the logarithm of) its probability of occurrence. This mathematical approach has proven its worth in signal processing and telecommunications, but its application to human cognition is debatable, as it proves hard to separate information from its practical context of use.

Knowledge, for its part, is inherently personal or social: knowledge is information acquired by human beings. Knowledge must be learned in context, individually or collectively, before being put to use to accomplish human goals and functions. The very notion of knowledge is inseparable from cognitive and social processes, while information could be defined more narrowly as a property of the physical world.

The point is that even if information can be objectively quantified for engineering purposes, *only knowledge is of real social and economic importance.* But knowledge is also difficult to acquire. Information may be copied or reproduced mechanically, but knowledge must be assimilated by humans before it can be used. And specialized knowledge can only be acquired by well-prepared specialists, restricting its effective social range of application.

The increasing division of labor, the complexity of technical knowledge, and the pace of innovation make it more and more difficult to ensure the transmission of knowledge within organizations and firms. Training or tutoring mechanisms may be devised, but bringing together the appropriate people remains a problem for learning to succeed. One must find both adequate experts and well-prepared apprentices. This is very much a social problem, which must first be solved for knowledge transmission to take place.

Another important distinction is between *explicit* and *tacit* knowledge, or perhaps more accurately between explicit information and tacit knowledge (usage is unfortunately not coherent here). Explicit knowledge or information is public and formalized, in linguistic or mathematical form notably. Books, journals, textual documents of all kinds, Web sites, databases, and so forth—all contain explicit knowledge, as long as one knows the linguistic or formal conventions necessary to interpret their content.

Information retrieval and computer science deal mostly with explicit information, so that it is too easy to forget that this is only one kind of knowledge. Real social life, however, makes frequent use of other forms of knowledge as well, which can be grouped together under the general label of tacit or implicit knowledge (Polanyi, 1966; Baumard, 1999). There is in fact a variety of forms of tacit knowledge (such as body language, common sense, work expertise, procedural knowledge, etc.), and one might distinguish further between unformulated and unconscious knowledge, but we will not attempt a more detailed analysis here.

Tacit knowledge is knowledge that has been acquired from practical experience: medical expertise, technical know-how, teaching experience, and management skills are forms of tacit knowledge. This cannot be learned from books alone, as learning by doing is a necessary component. Organized tutoring may help, but transmission will then be from person to person, which proves to be a slow and cumbersome process. *Tacit knowledge remains a serious bottleneck in the information society.*

One should also notice that tacit knowledge is often collective. Many organizations perform (more or less adequately) thanks to collective routines and procedures that are distributed among many actors and are often left unformalized. The knowledge inherent in organizational functions is not expressed publicly, and no single actor knows the whole picture. This lack of clarity may lead to serious inefficiencies.

Tacit or implicit knowledge is thus hard to learn and to pass on, and the computer revolution has so far not helped very much in this respect. As tacit knowledge is unfortunately an essential part of social life and economic performance, this is an area that begs for more consideration from knowledge management in general and information retrieval in particular. We feel that serious advances could be expected in this domain.

Last but not least, *social knowledge* is the (largely implicit) knowledge necessary to make use of social relationships so as to perform tasks and solve problems. It is an important component of most professions, but one that is usually learned by long practice and experience. The social skills and expertise necessary to find information needed for a given task, ensure social cooperation, and negotiate common rules are crucial to task performance in most lines of work.

Social knowledge has not yet been given sufficient recognition, however, and is rarely discussed, described, or formalized. Sociologists have been interested in the social structure of groups and how this constrains individual choices and strategies (e.g., Lazega, 2001). But there has been much less emphasis on individual knowledge of these constraints, on how they might be represented and processed cognitively. This calls for more research in social psychology.

To be able to access or use social knowledge would be quite useful for information retrieval systems. Finding the appropriate expert most likely to answer a technical question, for example, is often a better idea than searching the Web by oneself. Though the issue is usually not presented directly in this way, we will see below that collaborative software systems have started to incorporate elements of social expertise.

Social Networks

It should be obvious by now that an important part of human knowledge management takes place by way of social links and requires appropriate social expertise. Social networks have fortunately been studied and formalized by structural sociology, and there is a sizable body of methods and techniques to draw upon (Wassermann & Faust, 1994).

Social networks are a simplified model of social relationships, schematic enough to be represented and handled mathematically on a computer. The basic data structure is a graph, where nodes stand for social actors (individuals or groups) and links represent social relations. Links are usually not labeled, but may have an associated numerical value (standing for the strength or frequency of the relation). This graph is in turn implemented as a matrix on which various operations can be performed.

If the matrix represents direct links, indirect relations (requiring several steps through the network) can be found by computing successive powers of the basic matrix. For instance the square of the matrix will show two-step relations, the cube of the matrix three-step relations, and so on. Many other operations are also possible, and there are various algorithms for extracting from the social graph densely linked subgroups of nodes.

This approach is obviously a drastic simplification of the complexity of real human relationships, but the formal structure of such models can already be very revealing. In particular, the structural subgroups that can be extracted automatically from the graph correspond to social groupings of actors, working on similar tasks and exchanging information about common concerns. Structural subgroups are usually functional groups as well.

For example, after mapping the network of collaboration relationships between 71 lawyers in an American law firm, it is possible to find 11 dense subgroups corresponding to specific locations or specialties (see Lazega, 2001). As these subgroups also interact with one another, they can be seen as forming a higher-level network with fewer nodes, a kind of summary of the basic network. The whole process requires some human interpretation, but reveals social facts that are simply not obvious to the naked eye.

The position of an actor within the social network is usually significant: it shows the centrality or prominence of the actor, and the resources and information he has immediate access to. The network also shows the nodes and paths an actor would have to follow in order to access more remote information. Sociologists tend to interpret structural positions in terms of power relationships: central positions are strategic while actors located at the margins have to go through others to access various resources (Burt, 1992).

From our point of view, however, the main issue to consider is that *structural networks determine social access to information*. Central actors have quick and easy access to socially embedded knowledge, while marginal actors might have to contend with longer access routes. The social expertise necessary to retrieve socially situated information comprises social skills (such as diplomacy or bargaining tactics), but also the basic ability to perceive and exploit the social structure as such.

Social expertise may remain more or less unconscious, but the deliberate "networking" behavior of the ambitious professional is also quite common. Many professionals know the importance of "weak ties": useful information and opportunities are often obtained through casual relations which thus deserve to be strenuously cultivated (Granovetter, 1973). At the same time, developing and using a network of close contacts in the workplace is often a prerequisite to successful work performance.

Social information retrieval and problem solving by human beings is thus achieved through social networks, which govern information circulation and information flow. Formalizing this structure should be very helpful in order to model human knowledge management skills and capabilities, and possibly to design better collaborative software systems.

Now the development of electronic transmission networks has made it possible to extract automatically many social relations, as they leave electronic traces. For instance one may note the pattern of e-mail messages exchanged within an organization and formalize it as a graph. Of course, not all social interactions are reflected in electronic messaging, but e-mail traffic is obviously significant in many modern organizations. Web browsing is also a more indirect source of social affinities, which can be exploited to retrieve social information.

As a matter of fact, collaborative software systems make use of social links and social information, directly or indirectly. They might have been consciously designed in this way, but this may also be the result of practical attempts to solve an informational problem.

TECHNICAL APPROACHES

After this review of human information processing in social context, it is now time to consider how the insights gained during this study can be used to design social information systems. This should also help us put in perspective recent work in collaborative software.

Typology of Collaborative Software

There is already a variety of collaborative systems, but one can try to regroup various proposals into a few classes. We would like to propose a general typology of these systems, using a few relevant features to differentiate between them.

What collaborative systems have in common is the modeling of a social environment and use of social expertise to access relevant information. They differ, however, in the manner, explicit or implicit, in which they model the social community that serves as context for information purposes. Some systems provide users with an explicit representation of a social group, which may be consciously accessed as such. Other systems use social links implicitly, and the end users do not have to be aware of the underlying social structure (we prefer calling such links *implicit* rather than tacit because they might be totally unconscious).

Another pertinent distinction is whether the focus of operations is on the group itself or on the informational task being performed. Virtual communities tend to be task oriented and more impersonal than real communities, and some collaborative systems will emphasize the task more than the social group. In such a case, representing the task at hand is the central issue, and the explicit or implicit representation of the group structure becomes of secondary importance.

One should also remember that a collaborative information system does not have to reproduce faithfully every aspect of human information processing. There are fruitful lessons to learn from studying socially situated human cognition, but a software system can do things differently (and more efficiently in some ways) than the human mind. For example, social expertise about how to retrieve relevant knowledge may be implicitly built into a computer system, whereas a human being would have to search his social network consciously.

In fact some software systems stick closely to the structure and functioning of real human groups, and exhibit the same limitations in terms of group size or cognitive load. We would contend that virtual communities may well function differently, and that collaborative software should be designed accordingly. On the other hand, present software is still far from the complexity and capabilities of human social processing, so that there remains much to be learned from real human cognition.

Still, collaborative systems may also be classified in different ways, notably by using more technical criteria. The manner in which individual participants, relationships, and communities are represented and the clustering algorithms are used to regroup similar actors, the data structures and implementation techniques could also be used to differentiate between systems. But the emphasis being here on social issues, a classification based on community type seems more appropriate to this discussion.

To sum up, we think that work on collaborative systems up to now can be roughly classified into three main types: building explicit communities, building task-oriented communities, and using implicit social links.

Building Explicit Communities

This is the most obvious direction, and this research field is often known as *groupware* (Favela & Decouchant, 2003). Such systems try to make as explicit as possible the structure of the group, the biography and interests of participants, their role and status, and the history of interactions. The goals, tasks, common tools, past actions, and current problems can be posted publicly. The rationale is that group awareness and explicit interactions are conducive to better problem solving.

In a hospital setting for instance, there is an intense exchange of information between various medical staff (physicians, nurses, laboratory technicians, etc.), and timely access to correct information is clearly vital. But medical staff is highly mobile, and information is heterogeneous (verbal exchanges, textual records, images, etc.) and rapidly changing. The collective task to be performed (taking care of patients) is therefore highly distributed and in constant evolution.

And the problem is not just information distribution, but rather one of coordination between different actors and collective decision making. Although they often communicate through common objects (such as whiteboards and clipboards), medical personnel must be aware of each other, because the source and time of information may be crucial. A multi-agent architecture can then be used to locate or notify the right person at the right time with the appropriate information (Munoz, Gonzalez, Rodriguez, & Favela, 2003). In this way interactions are made explicit but also kept under tight control.

Yet groupware systems, properly speaking, can only function for small groups of participants. When the number of active members reaches more than 30 or 40 people, personal information and individual interactions may prove overwhelming. On the other hand, high awareness about individual group members may lead to personal relations and allow focusing a search for information on the person most likely to know the answer to a problem.

In this way, groupware systems try to reproduce the functioning of small social groups as we traditionally know them: family and friends, office life, workgroups, neighborhood associations, and so forth. Such systems are often used in a close professional context (e.g., a hospital or a firm) where people already know each other or are likely to meet face to face sooner or later. In this case, groupware will reinforce or assist real or potential social relationships, but will not create unexpected links.

Groupware design presents interesting technical challenges for computer scientists: managing synchronous and asynchronous communication between participants in various and changeable locations, transmission of heterogeneous data (including text files, messages, images, and sound), maintaining the coherence of common data structures, and so on. Sophisticated systems have been developed, notably for healthcare environments and computer-supported collaborative learning. But these systems are not widely used, probably because they are still too cumbersome and not appropriate for many social groups.

Groupware can be useful in professional domains requiring intensive social links with focused interactions dealing with very specific tasks. The density and quality of interactions require fairly elaborate software to update and transmit information in a graceful and readable way, with heterogeneous data and more and more mobile users. But groupware is inadequate and unwieldy for larger groups and casual interactions.

Another possibility is to use the social network that can be inferred from Web pages, social interactions, and common interests to locate experts on a given subject (Kautz, Selman, & Shah, 1987). This might be the only way to find tacit information, which is not publicly available. This approach may also be developed to improve information retrieval by taking advantage of the social links of document authors for instance—well-connected authors are probably more reliable (Kirsch, Gnasa, & Cremers, 2006). But we will see below how to exploit implicit links.

Still another research direction that has not yet been developed much in computer science would be to post an explicit structure for the social network in a given domain. So this would also be an explicit representation, but a more schematic and lighter one.

We have seen that structural sociology has elaborated formal models of social groups considered as networks of relations (Wassermann & Faust, 1994). The complexity of real social interactions is deliberately simplified so as to represent a group by a graph, in which nodes are actors and links are relations. Social interactions are reduced to simple relations, such as collaboration, advice, or influence.

Without going into more detail, the point is that structural sociology is well formalized and sufficiently advanced to offer relevant representation tools for larger communities. Representing groups with hundreds of members is not a problem, and the nature of links (edges in a graph) is simpler and more abstract. For larger communities, these formal methods might be a better source of inspiration than current groupware techniques.

From a practical point of view, structural methods could be used to map the current state of a community and to show participants their position in the network, the coherence of the structure, what the sub-groups are, the dynamic evolution of the network, and so forth. This would be another way to raise group awareness, not in personal terms but from a structural, more abstract perspective.

In a large firm, for example, it might be useful to be able to identify structural subgroups in order to find appropriate contacts on a functional rather than a personal basis. Although this is technically possible and software systems are now available for this purpose, they are not really used in practice, perhaps because they are felt to be too revealing and intrusive.

Still, when participation is only occasional or unique, and when interactions are mostly impersonal, the notion of structural network loses significance. If all interactions take place through a common workspace, the most one could probably hope for is to make it easy for users to enter the system and to deal with common objects. A good data structure and convenient access and modification procedures are then necessary.

Building Task-Oriented Communities

Virtual communities are frequently task oriented. Computer-mediated communities are often quite different from traditional social groups, a fact that is too rarely acknowledged in the literature. By comparison with traditional groups, participation in virtual communities is more impersonal, often temporary or anonymous, with a lower level of emotional involvement. These communities are mostly goal oriented: participants contribute to a common goal or task, but are less interested in personal relationships.

In such a case, group activities revolve around a common data structure (forum, discussion thread, Web site, wiki, database, etc.) that shows the current state of the task in progress and is regularly updated. This is a *blackboard* model, where all interactions go through a central data structure rather than by means of particular links.

Such an architecture was originally proposed for the Hearsay-II speech understanding system as an efficient method to coordinate the operation of various modules: all communication between modules takes place through the blackboard (Lesser & Erman, 1977). In our domain, this can be seen as a form of situated cognition, determined by a common public environment which is represented here by a central blackboard.

Since most of the information necessary for group activities is posted on this blackboard, *information retrieval can be done by accessing the common workspace.* Information management is collective, in the sense that the blackboard somehow summarizes the whole history of group interactions and contains all the information deemed relevant by the group. This is another form of collaborative retrieval, but of an indirect and impersonal kind.

One reason that may explain the prevalence of this type of communication is simply that it minimizes the complexity of interactions. The number of potential point-to-point links between n actors is n(n-1)/2, which grows like the square of the number of actors. But the number of interactions with a common data structure only increases linearly with the number of participants, a much more manageable proposition for larger groups.

There is in fact no sharp boundary between explicit communities and blackboard-mediated groups, and the distinction is not always clear. For example, in hospital wards, the “blackboard” (actually a whiteboard) is only one source of information among others. Yet there is a strong tendency in modern life, notably in virtual communities, toward more impersonal, functional, flexible social groups organized around a common task or goal. Such groups have their own informational requirements, which must be served by access to a simple, robust, easily maintained blackboard structure.

The recent *wiki* technique is a good example of user-friendly blackboard management system. A wiki is basically an interactive Web site with simple and easy editing procedures. Registered participants may post text messages on the site, and they can also augment, comment on, or modify previous messages. So everybody can contribute to the site, but interventions must be signed and the history of modifications is kept automatically. In practice, a moderator is useful to check interventions before they are posted.

The well-known online encyclopedia Wikipedia has been (and still is) developed in this way with very good results overall (*www.wikipedia.org*). The quality of entries is not always consistent, and there have been a few problems with inaccuracies or vandalism (hence the importance of competent moderators). But on the whole Wikipedia has proven to be a successful collective, collaborative enterprise and a model of what could be accomplished online.

Although it is in fact a more complex phenomenon, the development of free or open source software may also be seen as a task-oriented activity (Feller, Fitzgerald, Hissam, & Lakhnani, 2005). A software project under development serves as a common object which is repeatedly corrected and improved by a wide community of programmers and testers, many of whom do not interact on a personal basis. This community is strongly structured, however, with a small inner core of project leaders surrounded by concentric circles of contributors and critics, so that this would really be a hybrid example between personal and impersonal relations.

Using Implicit Social Links

Other software systems do not post group structure or common data. This is usually the case with collaborative information retrieval, collaborative filtering, and recommender systems. There exist many variants, but the basic idea consists of exploiting the implicit structure of a group of users in order to find relevant documents, filter search results, or recommend information or products. The grouping may be made public in some systems, but is usually not handled by the users themselves, who might remain totally unaware of this virtual structure.

These collaborative systems work by computing similarities between human users and by taking advantage of the resemblance to share information between similar users (Resnick, Iacovou, Suchak, Bergstrom, & Riedl, 1994; Shardanand & Maes, 1995; Adomavicius & Tuzhilin, 2005). For example one may recommend movies, books, music, or other products to a given user by finding "similar" users and quoting their best choices. Or one may retrieve or filter documents by noting which documents have been retrieved or used by groups of similar users.

To throw some light on the variety of such systems, one may want to make several distinctions between them. Although real systems often blur these distinctions, the following categories of collaborative systems may be useful:

- **Collaborative filtering (recommender systems):** These systems recommend (or rank) products, services, or documents for the benefit of an individual user by collecting the preferences of similar users.
- **Collaborative retrieval systems:** These retrieve (or filter) relevant documents by using the profiles of similar users. Poor initial queries can thus be augmented with more expert information.
- **Active (explicit) rating:** Users explicitly take the time to rate or recommend products. People are amazingly willing to do so (probably as a form of self-expression, in order to promote a product they like, out of sheer sociability, etc.), but their active intervention is required.
- **Passive (implicit) rating:** Information on user preferences is collected by noting significant user actions (buying products, Web browsing, bookmarking, downloading files, etc.). This can be done automatically, but user tastes are inferred, not directly measured.

This general approach requires establishing an interest profile for each end user, and choosing a similarity measure so as to be able to compare users in a coherent way. By analogy with classical information retrieval methods, each user is usually characterized by a vector of relevant features, and users are compared by computing their proximity in vector space. The group profile used as a basis for recommendations can then simply be the average of member profiles.

There have been quite a few variations, such as employing statistical correlation, angle or distance between vectors, or various clustering algorithms to estimate user resemblance, but the determination of a user profile is of course crucial to the operation of the system. One may want to compare different methods, but results depend on the task and the nature of the data (Breese, Heckerman, & Kadie, 1998).

Instead of comparing users to find subgroups of users with similar interests, it is also possible to compare and cluster items with regard to user preferences (this is what Amazon.com does). If you like a particular item, the system can then recommend similar items. But the latter method is in fact a dual representation of the former: one may equivalently represent users in item space or items in user space, but a choice can be made for reasons of implementation.

Collaborative systems unfortunately suffer from a "cold-start" problem: a critical mass of users and user preferences is needed for the system to prove valuable. There is then little incentive for initial users to join the club, and some way must be found to attract them in order to build this critical mass. Symbolic rewards might help in this regard (the pleasure of participating in an innovative experiment for example).

One should also be aware that rankings depend on the particular rating method chosen to evaluate the relevance of documents or products. We have seen that the rating of a particular item could be determined by explicit user evaluations, by semantic proximity to user profiles, or by recording user actions concerning this item. Evaluations may depend both on user profiles and user actions in variable combinations.

In short, *implicit collaborative systems work by setting up groupings of similar users* and then exploiting these virtual groups to retrieve or recommend socially supported items. Collecting individual ratings (whether explicit or not) about items is a prerequisite to calculating their overall social value in the group of reference.

Another example of the implicit use of social structure is offered by PageRank, Google's ranking algorithm for Web pages (Brin & Page, 1998). This famous search engine retrieves pages in classical fashion (by computing their textual similarity to a user query) but then orders them by exploiting the structure of Web links. The page ranking is meant to solve the frequent problem of information overflow with too many answers to a query.

More precisely, Web pages are ranked by the sum of hyperlinks pointing to them from other Web sites, each link being weighted with the value of the pointing site, determined recursively in the same way by considering its own incoming links. The ranking of a site thus increases with the number and value of sites pointing to it. A careful matrix implementation of the graph of hyperlinks speeds up the recursive value computation.

The hyperlink structure used by the PageRank algorithm is in fact the public trace of an implicit social consensus. Web sites with numerous incoming links are better known (and tend to attract even more new links) as they have been judged more relevant by other Web site publishers. This is a measurable form of hyperspace reputation on the Web, which is presumably a good indicator of the interest and trustworthiness of a Web page. The success of Google is largely due to the clever use of this social indicator.

Peer-to-peer file sharing systems such as Napster, Gnutella, or KaZaA have also been very successful, to the horror of major music companies. They work by distributing requests through a network of participants so as to find users with similar tastes. Music files or other documents can then be exchanged among like-minded participants. Napster employed a central server to store and compare user profiles, but in more recent systems both data and processing are totally distributed throughout the network (Memmi & Nérot, 2003; Wang, Pouwelse, Lagendijk, & Reinders, 2006).

Peer-to-peer architectures can be used for file sharing, information retrieval, and collaborative filtering. But the implicit links between users do not have to be made public for the system to work, thus allowing a minimum of privacy.

In spite of their differences, these various collaborative systems all make use of distributed, implicit, socially situated knowledge by building or revealing virtual communities. Relevant information is accessed through social links, and retrieval algorithms embody social expertise about information handling. But individual systems users are not made directly aware of the underlying group structure.

TRENDS AND PERSPECTIVES

Even though our survey has not been exhaustive, the diversity of approaches and collaborative systems is striking. So the question arises whether one can discern general tendencies among recent research work. It is by no means clear at this time that one approach will predominate over the others, but we would like to venture a few general observations and suggest likely developments.

Following the typology proposed above, explicit communities and task-oriented groups are the most obvious phenomena and have probably attracted more initial attention as a basis for computer-aided communication. But it seems to us that *there remains more to discover about implicit social links,* so that interesting novel techniques may be expected to appear in this direction. Because more and more information is becoming available in electronic form about human relationships, new ways will be found to exploit such information.

For example, commercial transactions and work connections often leave electronic traces which can used for informational purposes. Profiling people by their commercial or browsing behavior can also be used to put together virtual groups with similar interests and needs. On the other hand, such techniques could also prove very intrusive, posing difficult ethical and social problems about individual privacy.

We believe that more detailed analyses of human social information processing would be a fruitful source of new techniques. We have tried here to show the wealth and complexity of social information processes, but we still do not know enough about such common social mechanisms. Studying and modeling collective information management should bring about new insights and suggest new approaches.

Unfortunately, interest in this area has traditionally been dispersed among very different disciplines, which do not communicate very well with each other. Sociology, economics, and management studies notably have contributed valuable observations about human knowledge management, but this is too rarely a central concern and approaches vary widely. Fundamental research in this domain is then more likely to be a source of inspiration to computer science than to provide a store of directly applicable models.

Accessing and making use of tacit knowledge has hardly started, and usually only indirectly. In spite of the social and economic importance of this type of knowledge, it only becomes accessible online as a by-product of explicit communication links on the Internet. No systematic effort has been made so far to address this question by computer, although the problem is largely recognized in real life (tutoring relationships and training schemes are basically meant to ensure the transmission of implicit or tacit knowledge).

Profiling individuals by their electronic behavior is the most likely route in order to gain access to the tacit knowledge they might possess, but for privacy reasons this is probably feasible only within work situations and inside organizations. And as to collective tacit knowledge (the kind of knowledge that makes a company more or less efficient), one simply knows very little about how to describe or formalize such distributed information.

We would also like to suggest that a generic platform or general toolbox for collaborative software design would be a good idea for experimenting with various methods. It would help build prototypes and new software systems. Such a platform should contain the main representation and processing techniques we have seen so far, with capacities for exchanging information between different approaches. Common data representations would make it possible to share information among various tools.

A recent example of this kind of open toolbox can be found in the Sakai project (*www.sakaiproject.org*). This is a free collaborative environment which contains many of the communication techniques currently available for virtual communities. The emphasis is on education and e-learning, but the software can easily be extended to other areas.

In our view, such a toolbox should include in particular the following methods:

- Current communication tools (e-mail, chat, forums).
- Blackboard facilities (a wiki structure, for example).
- Social network simulation software.
- Social network analysis software.
- Common profiling and clustering algorithms.

Most of these software tools are already available, but in different domains, and they are rarely employed together. For example, elaborate methods have been developed for social network analysis, and software packages are easily obtainable (e.g., Ucinet or Structure), but they have mostly been used by sociologists. Electronic mail is widely used, but communication patterns are rarely collected and studied. Putting together different methods would make data available for analysis, and help investigate a complex and multidisciplinary field of enquiry.

To sum up, socially situated human information management is an intricate, multi-faceted domain, which we still do not understand well enough to reproduce in all its wealth and power. More fundamental studies are needed, as well as more friendly generic research tools. It is time for a global approach and for comprehensive software tools in order to improve our capacity for useful and efficient collaborative software design.

CONCLUSION

We have tried to show here how human information processing takes place in a social context and to what extent human beings use this context to retrieve information and solve problems. Shifting the emphasis from individual to social processes greatly improves our ability to understand and reproduce real human abilities. Studying and modeling socially situated information processing is therefore an important source of inspiration for the design of better collaborative information systems.

Of course, technology does not have to imitate life. It has often been the case in the history of computer science that efficient solutions to practical problems were derived mostly from technical considerations. Computers do not work by duplicating human thought processes faithfully, but by exploiting the speed and accuracy of electronic devices. Technical constraints and possibilities may have their own logic.

For high-level abilities, however, and especially when dealing with new areas to model, analyzing human cognitive processes is often both a prerequisite and a good start for system design. In the domain of information retrieval and knowledge management, studying closely the way human society performs its knowledge tasks by using distributed, collaborative processes has proven to be a fruitful approach. We are convinced that useful design ideas are still to be gained in this manner.

REFERENCES

Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering, 17*(6).

Baeza-Yates, R. (1999). *Modern information retrieval.* Boston, MA: Addison Wesley.

Baumard, P. (1999). *Tacit knowledge in organizations.* London: Sage.

Breese, J.S., Heckerman, D., & Kadie, C. (1998). Empirical analysis of predictive algorithms for collaborative filtering. *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence.*

Brin, S., & Page, L. (1998). *The anatomy of a large-scale hypertextual Web search engine.* Computer Science Department, Stanford University, USA.

Burt, R.S. (1992). *Structural holes: The social structure of competition.* Cambridge, MA: Harvard University Press.

Castells, M. (1996). *The rise of the network society.* Oxford: Blackwell.

Clark, A. (1998). *Being there: Putting brain, body, and world together again.* Cambridge, MA: MIT Press.

Drucker, P.F. (1992). *The age of discontinuity.* New York: Harper & Row.

Favela, J., & Decouchant, D. (Eds.). (2003). *Groupware: Design, implementation and use.* Berlin: Springer-Verlag.

Feller, J., Fitzgerald, B., Hissam, S.A., & Lakhnani, K.R. (2005). *Perspectives on free and open source software.* Cambridge, MA: MIT Press.

Granovetter, M.S. (1973). The strength of weak ties. *American Journal of Sociology, 78,* 1360-1380.

Harnad, S., & Dror, I.E. (Eds.). (2006). Distributed cognition—special issue. *Pragmatics and Cognition, 14*(2).

Kautz, H., Selman, B., & Shah, M. (1997). Referral Web: Combining social networks and collaborative filtering. *Communications of the ACM, 40*(3).

Kendrick, J.W. (1994). Total capital and economic growth. *Atlantic Economic Journal, 22*(1).

Kirsch, S., Gnasa, M., & Cremers, A. (2006). Beyond the Web: Retrieval in social information spaces. *Proceedings of the 28th European Conference on Information Retrieval.*

Kollock, P., & Smith, M. (Eds.). (1999). *Communities in cyberspace.* London: Routledge.

Lazega, E. (2001). *The collegial phenomenon.* Oxford: Oxford University Press.

Lesser, V., & Erman, L. (1977). A retrospective view of the Hearsay-II architecture. *Proceedings of the 5th IJCAI* (pp. 790-800).

Mandler, G. (1985). *Cognitive psychology.* Hillsdale, NJ: Lawrence Erlbaum.

Manning, C.D., & Schütze, H. (1999). *Foundations of statistical natural language processing.* Cambridge, MA: MIT Press.

Mauss, M. (1924). Essai sur le don. *Année Sociologique 1923-1924.*

Memmi, D., & Nérot, O (2003). Building virtual communities for information retrieval. In J. Favela & D. Decouchant (Eds.), *Groupware: Design, implementation and use.* Berlin: Springer-Verlag.

Memmi, D. (2006). The nature of virtual communities. *AI and Society, 20*(3).

Mintzberg, H. (1979). *The structuring of organizations.* Englewood Cliffs, NJ: Prentice Hall.

Munoz, M.A, Gonzalez, V.M., Rodriguez, M., & Favela, J. (2003). Supporting context-aware collaboration in a hospital: An ethnographic informed design. In J. Favela & D. Decouchant (Eds.), *Groupware: Design, implementation and use.* Berlin: Springer-Verlag.

Neisser, U. (1967). *Cognitive psychology.* New York: Appleton-Century-Crofts.

O'Reilly, T. (2005) *What is Web 2.0.* Retrieved from *http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html*

Polanyi, M. (1966). *The tacit dimension.* London: Routledge & Kegan Paul.

Rheingold, H. (2000). *The virtual community.* Cambridge, MA: MIT Press.

Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., & Riedl, J. (1994). GroupLens: An open architecture for collaborative filtering of Netnews. *Proceedings of the 1994 Conference on Computer-Supported Cooperative Work.*

Salton, G., & McGill, M. (1983). *Introduction to modern information retrieval.* New York: McGraw-Hill.

Shapiro, C., & Varian, H.R. (1999). *Information rules: A strategic guide to the network economy.* Cambridge, MA: Harvard Business School Press.

Shardanand, U., & Maes, P. (1995). Social information filtering: Algorithms for automating "word of mouth." *Proceedings of the 1995 Conference on Human Factors in Computing Systems.*

Travers, J., & Milgram, S. (1969). An experimental study of the small world problem. *Sociometry, 32,* 425-443.

Wang, J., Pouwelse, J., Lagendijk, R., & Reinders, M. (2006). Distributed collaborative filtering for peer-to-peer file sharing systems. *Proceedings of the 21st Annual ACM Symposium on Applied Computing.*

Wasserman, S., & Faust, K. (1994). *Social network analysis: Methods and applications.* Cambridge: Cambridge University Press.

Watts, D.J. (1999). *Small worlds.* Princeton, NJ: Princeton University Press.

Wellman, B. (Ed.). (1999). *Networks in the global village.* Boulder, CO: Westview Press.

This work was previously published in Social Information Retrieval Systems: Emerging Technologies and Applications for Searching the Web Effectively, edited by D. Goh and S. Foo, pp. 189-208, copyright 2008 by Information Science Reference, formerly known as Idea Group Reference (an imprint of IGI Global).