How Information Technology Could Benefit from Modern Approaches to Knowledge Management

Christopher Lueg
Artificial Intelligence Laboratory
Department of Information Technology
University of Zurich
lueg@ifi.unizh.ch

Reinhard Riedl
Systems Architecture and Software Group
Department of Information Technology
University of Zurich
riedl@ifi.unizh.ch

Abstract

The ability to find appropriate information in a work situation is a knowledge-intensive activity that is relevant to knowledge management. However, whereas modern knowledge management increasingly adopts knowledge-sharing approaches, such as communities of practice, commercial information technology that is to support knowledge management still tends to aim at traditional document-centered approaches to knowledge management. In this paper, we argue that document-centered approaches tend to neglect essential aspects of information-seeking and illustrate the point by example of a study conducted in a large swiss company. We conclude that a paradigm change is needed and propose to consider communities of practice not only in terms of knowledge management but in terms of information management as well.

1 Introduction

Knowledge management in organizations is widely considered inseparable from document management as documents are recognized as important knowledge resources that are viewed as part of the organizational

memory. Apart from accessing and managing document collections, the ability to find appropriate information in a work situation is increasingly recognized as a knowledge-intensive activity and should therefore be considered as an important aspect of knowledge management as well.

In respect to the importance of an organization's document collections, knowledge management has been undergoing a significant change over the past few years. Whereas early approaches to knowledge management focused on the management of document collections viewed as knowledge repositories that just have to be accessed in an appropriate way, more recent approaches have begun to stress the importance of knowledge *sharing* for knowledge management. This change was induced by a changing conceptualization of knowledge from something that is "contained" in documents to something that can be "generated" or "reproduced" in the interaction with documents which are now understood as representations rather than as containers. The meaning of "knowledge transfer", in the knowledge sharing perspective, is quite different from the traditional understanding of knowledge transfer through front-based teaching or reading books and is nicely illustrated in the following :

> "[...] *objective knowledge* seems possible and through *objective knowledge* the universe appears systematic and predictable. Yet *knowledge* as an experience is something personal and private that cannot be transferred, and that which one believes to be transferable, *objective knowledge*, must always be created by the listener: the listener understands, and *objective knowledge* appears transferred, only if he is prepared to understand." [14, p. 5, original emphasis].

The recognition of the "fluid" nature of knowledge

has fostered the acceptance of socially-oriented approaches to knowledge management. For example, workgroups are increasingly viewed as "communities of practice" [22]. The term was originally coined by Jean Lave and Etienne Wenger at the beginning of the Nineties when they explored apprenticeship in various types of communities ranging from midwives in Mexico and tailors in Liberia to butchers in U.S. supermarkets and Alcoholics Anonymous [9]. Lave and Wenger found that "legitimate peripheral participation" is a prerequisite for effective learning though communication and collaboration.

Information technology provides efficient support for knowledge management viewed as document-management but has difficulties to meet the flexibility demanded by knowledge-sharing approaches to knowledge management. In a sense, information technology is still committed to the document-management perspective but has failed to keep pace with recent changes in knowledge management. Finding appropriate information in a situation is too often equated with retrieving the information from existing document collections. In this paper, we outline that much relevant work done in the human-computer interaction (HCI) and computer-supported cooperative work (CSCW) communities is waiting to make its way into the world of business-oriented information technology.

We proceed as follows. First, we briefly motivate the change information technology has to undergo in order to meet the flexibility demanded by modern knowledge management. In particular, we discuss similarities between the traditional understanding of knowledge —which modern knowledge management is overcoming— and the traditional understanding of information needs which still has significant influence on information technology. Following this more theoretical discussion, we motivate its practical relevance by discussing an information-technology and knowledge-management scenario that we recently investigated in a large swiss corporation. Finally, based on a more appropriate perspective on information needs, we outline a new paradigm for information-seeking support that is inspired by modern approaches to knowledge management. The paper concludes with a brief summary and an outline of future research directions.

2 On the Similarity of Knowledge and Information Needs

The conceptualization of knowledge has changed significantly of the past few years. Documents are no longer regarded as knowledge "containers" but merely as representations of knowledge that can be used to *reproduce* and to *generate* knowledge. Modern approaches to knowledge management increasingly adopt the modern cognitive-science perspective that knowledge is not so much facts and rules as in the classical cognitivistic perspective but that knowledge is best conceptualized as the capacity to coordinate past experiences with current experiences.

We argue that a similar development from traditional and rather "static" conceptualizations to more "fluid" conceptualizations is required in information technology as well if information technology is expected to support modern approaches to knowledge management. Previously, we argued that the ability to find the right information in a situation is increasingly recognized as a knowledge-intensive activity and that it should be considered as an important aspect of knowledge management. An investigation of information-seeking support reveals that most current support is inspired by the concept of information needs (see [10] for a broad discussion of information needs). The very concept originates from information-retrieval research (e.g., [8,]) and means that users engage in information-seeking activities because they dispose of information needs:

> "Regardless of the terminology used and the motivation, the information problem is the trigger for information seeking and [...] it evolves and changes as the search and the overall situation evolve." [12, p. 36].

The goal of information-seeking, in this perspective, is to find information that satisfies these needs. Accordingly, much support focuses on modeling these information needs and on finding information that is appropriate to satisfy the needs. As information needs cannot be investigated directly, information judged to be relevant by the user is used to infer underlying information needs. Relevance is the central concept in information-retrieval that underlies almost all search activities (see [4] for a broad discussion of relevance).

The similarity of knowledge and information needs addressed is in that both are concepts "in the head of the user", both can be described by users to a certain extent, and the existence of both knowledge and information needs can be empirically tested. In the case of knowledge, this lead researchers to the assumption that the essence of knowledge consists of facts and rules, and that these facts and rules can be extracted and be stored in so-called knowledge representations. However, it turned out that the expressibility as well as the testability revealed little about the nature of knowledge. In the case of information needs, the perspective is that information needs can be described in such a way that the resulting description can be matched to documents in order to find the documents that would satisfy the information need.

Furthermore, the concept of information retrieval as the satisfaction of abstract information needs resembles —in a way— the expert system approach to support problem-solving [10]. In both cases, the approach is to model concepts assumed to be in the user's head (information need or problems) in order to automatically find something appropriate that helps resolve the problematic situation. In the case of information seeking it is information that satisfies the user's information need; in the case of problem-solving it is a solution to the user's problem. In the expert system approach, failure experiences have lead to a lasting re-conceptualization of problem-solving and related concepts, such as problems, knowledge, and expertise. It turned out that the view of problem-solving as matching problem-descriptions to solutions did not adequately capture the nature of real-world problem-solving. So-called problem-solving in the real world is not an abstract activity that takes a problem-description as input and that produces a solution as output as suggested by the information-processing metaphor of human cognition. Instead, real-world problem-solving requires the interaction of the human with the problem situation in order to bring knowledge, expertise, and past experiences to bear. Moreover, real-world problem-solving involves the conceptualization of a (problem) situation as something that has to be dealt with. As computers cannot understand situations in this sense, the very understanding of computers as real-world "problem-solvers" is inadequate. The misconception is nicely expressed in the following quote:

"Of course, they [computers] are no problem-solvers, because they do not have any problems in the first place. It is *our* problems they help us solve like any useful tool, say, a hammer which may be dubbed a 'problem-solver' for driving nails into a board. The danger in this subtle semantic twist by which the responsibility for action is shifted from man to machine lies in making us lose sight of the problem of cognition. By making us believe that the issue is how to find solutions to some well-defined problems, we may forget to ask first what constitutes a 'problem', what is its 'solution', and — when a problem is identified — what makes us want to solve it." [21, p. 30].

In rather well-defined situations, expert systems as well as information-retrieval systems perform well and provide valuable support to their users. The problem is that these systems are too often considered general solutions in the sense that problems and information needs are considered as something that is rather context-independent. However, empirical evidence suggests that information needs only exist in specific contexts and that information is typically needed to "do" something with the information, such as becoming informed or solving a problem:

"A central finding of our study is that it is dangerous to blur the concepts of 'information' and 'knowledge'. Information must be put into context in order to be useful, and those contexts must be explicated, experimented with, and confirmed. [...] information retrieval must be recognized as part of a larger enterprise of building up knowledge to 'solve problems' not merely 'answer questions'. Knowledge, moreover, is grounded in collective as well as individual experience." [3, p. 23].

The criticism of the prevailing information-processing perspective on information seeking resembles the criticisms of the human factors-perspective in the human-computer interaction (HCI) field. The human factors-perspective being an information-processing perspective focuses on the analysis of isolated aspects of the human-computer interaction, such as limited short-term memory span and limited memory (thus, the term human "factors"). Human factors has been criticized as connotating the human as a "user" and thus as a "passive, fragmented, depersonalized, unmotivated individual" [1, p. 27]. Instead, the term "human actor" has been proposed to connotate an active, controlling one. The term "human actor" also emphasizes humans being autonomous agents having capacities to regulate and coordinate their own behavior rather than simply being a passive element in the human-machine interface [1, p. 28].

In the information-processing view on both information seeking and human-computer interaction, the term "user" denotes a human using a computer for some specific purpose, such as retrieving a document or doing document processing, respectively. However, in both cases the term fails to denote that the interaction between a user and a computer system is part of the larger context of the interaction of the human with his or her environment. Computer tools are used by humans to reach meaningful goals that usually exist beyond the situation of human-computer interaction. Moreover, computer tools often serve a purpose that is more remotely related to the situation of computer use. For example, the goal of using a word processor is typically less the interaction with the program but the writing of a letter and, in a larger context, the communication with a person.

The fact that seeking for information is always embedded in a larger context has severe implications for determining the relevance of documents. In the early years of information-retrieval research, it was assumed that relevance of a document to a given information need can be determined merely on the basis of the document. Modern approaches have accepted that determining relevance is not just a matter of matching an information need and documents but that relevance also depends on situational aspects.

Recently, sociological studies of information use in the context of the International Monetary Fund in Washington DC have shown that relevance can even be socially determined [5]. The study describes how information originating from a recently set-up economic analysis institute in a particular country have been treated as if the information were non-relevant although the sheer content of the information was certainly relevant:

> "It was not enough that information is somewhere out there in the world and simply gathered by some processes of human agency and technology. Rather, information was relevant when it was seen by all the parties involved." [5, p. 85].

Socially-determined relevance has not yet received the attention it deserves in information science and knowledge management. Moreover, all the considerations so far may appear rather academic to practitioners. The next section shows their practical relevance by example of the information-technology management in a major swiss company.

3 An Investigation of Information-Technology Support in a Large Swiss Company

As part of a large research project, members of the Systems Architecture and Software Group have analyzed the Intranet of a major, world-wide operating Swiss company, whose business success essentially depends on the IM services provided by its IT departments. The results are depicted in this section.

The company's knowledge management is based on documents and on training courses. Document-focused knowledge management is genuinely embedded into the communication Intranet, which is based on the pull paradigm. Users have to fetch information from the internal Web. For that purpose, they have to find out about the existence and the address of Web pages providing the information they need. This causes severe problems, since even if users are given recommendations for useful Web pages, it is hard for

them to access these pages, as long as they are not given the precise address. Unfortunately, in many cases personal recommendations are of the form "there is a Web page with this type of information somewhere ...". Although there is a vague matrix structure, guided experiments with users demonstrate that users often fail in tracing down such recommendations, one of the reasons for which is the frequent breach of the organizing matrix structure and another reason is the the limited recall of the search engines. Even members of the company's team for the development and management of retrieval systems failed to search by analogy in those parts of the Intranet, which were unknown to them. In fact, users usually have considerable problems to find pages visited before, if they did not bookmark them. Of course, without concrete recommendations or hints the situation is even worse. It has to be noted here, that the company is commonly considered as good example for successful introduction of Web-technology by higher management people and its Intranet design serves as a model for other companies. Moreover, its experts for retrieval support and content management are indeed highly aware of the state of the art in information retrieval. Nevertheless, the true situation is as frustrating as depicted.

In the following we shall first summarize the results of a statistical analysis of the Intranet log-files as well as the results of interviews with employees. Then we shall deduce the generic problems encountered. And finally, we shall draw our conclusions on the system management re-engineering.

We have analyzed the log-files of Web servers and of an internal information-retrieval system. The analysis yielded that the most common request for information incorporated a single search term. Search requests with multiple search term were seldom and Boolean compositions of search terms were rare. Duration of search sessions, i.e., sequences of interactions with a search engine, was usually very short, and they lacked thesaurus-like variations of search terms. In fact, as with respect to the distribution of session lengths, we observed some similarity between the interaction of the Intranet users with the search engine and the page viewing behavior of customers in marketing Web-sites of another large Swiss company, which we analyzed [13]. Moreover, apart from human users there were lots of information agents interacting with the search engine, and most of the sessions incorporating more than three search requests were created by non-human agents. Furthermore, the analysis revealed some most remarkable resemblance of the search behavior in the company's Intranet to the structural patterns of natural languages, which is not yet understood. See [19] for a broad discussion language-like patterns in log- and trace-files.

Interviews with the employees revealed some sincere frustration with the Intranet as a whole. Employees criticized that they were not informed about updates and that the information could be trusted as there is no "garbage collection" applied to the Intranet. As a consequence, outdated, explicitly false information is still available in parallel with "correct" information. Moreover, there is no notice which describes the confidentiality of the information presented on a page. Another reason for dissatisfaction of the users is the fact that the search engine does not meet its service goals as it only provides limited recall, which is partly due to the technical impossibility to spider the whole company, which in turn results to some extent from internal company politics. Anyway, there are no satisfactory education programs teaching the employees how to use the available tools, although some of these tools do assume much more user awareness of their functionality than is commonly given.

A survey of available tools showed that only limited functionality was available as compared with the current engineering state of art. Most surprisingly, the availability of tools was limited in practical day work. The actual problems could be traced back to a missing task in the business technology workflow, which lead to a wrong design, which in turn was successfully exploited by the dissemination procedure eventually resulting in a denial of service. During our investigations the launch of a meta-search engine was considered, which was designed, such that unsuited usage lead to considerable reduced precision of result presentations and to the creation of critical overload on connected retrieval engines. On the other hand, a project, which intended to overcome the drawbacks of the intrinsic pull principle of the current Intranet organization with the development of a customizable Web-monitor was stopped by the management. Finally, we observed some general resistance from single mid-level managers against the exchange of knowledge among different parts of the company.

Summing up, users were neither treated properly as valuable customers nor as estimable employees, as there was neither sufficient service nor sufficient teaching how to use the available tools. The pull principle dominated all activities in the sense that it was the duty and the responsibility of the employee to find out about updates and changes as well as to find out the place where the available company knowledge was published. The quality of the knowledge presentations was low as there was neither a garbage collection nor a signing procedure in use. Web addresses changed quickly for non-reproducible reasons, and responsibilities for knowledge presentations got lost due to the permanent change in personal resources. In fact, responsibility tracing was a vague, non-monitored task

of the heads of work-groups lacking guidelines and tool support. Finally the system management was totally fragmented, missing links in control chains were the normal case rather than exceptions, and no knowledge management strategy could be identified.

It seems as if there is little hope for the management to regain the control of the system as long as they continue to pursue a rigorous top-down strategy and they try to implement organizational structures in a one-to-one manner on the Intranet. As there is already a lot self-organization going on right now with positive and negative effects for the users, monitored and supervised self-organization might be a possible strategy to control the emerging chaos, but this requires a change of paradigms. Supervision has to respect the fact that value chains in knowledge management are not necessarily profit chains and quality criteris applied have to focus on the benefit for the user. The latter implies, that the value of generated knowledge representations and knowledge push actions has to be considered in the environmental context, in particular with respect to the users' social and problem context and the service offers provided by the information technology of the company.

In the research project discussed here, we have pursued two approaches to improve the situation. We have developed a hybrid agent system for experiments and we designed an alternative approach to business organisation based on the media and agents paradigm. While the first gave us a lot of insight into the capabilities and constraints of agent technology [20], its organizational implementation failed and we had to realize that this was the wrong approach. In the second approach we studied how to reconstruct an organization, in particular the bureaucracy, on an Intranet by employing agents. This gave us some insight how the use of a medium like the Intranet could transform existing organizational structures while improving the value and knowledge creation processes within the organization [6], but it turned out that it was much too ambitious and theoretical for a practical implementation. It has to be noted, though, that related agent modeling is applied in the design of British E-government. (See [2] for general policy information.) As a consequence of past, tragic failure of public services, authorities are becoming aware that IT has to nurture and support the formation of communities of practice in public health-care. Right now, the Systems Architecture and Software Group is experimenting with such agent modeling approaches in the organisational prototyping for inter-organisational E-government in a joint project with various European city councils, which are very concerned about public acceptance of digital administration.

In what follows, we depict how a basic paradigm

change concerning the application of information technology could contribute to resolving the muddy situation in Intranets as the one discussed above.

4 Knowledge Management, Information Technology, and Communities of Practice

In the previous sections we outlined that the applicability of information technology is limited in the context of information seeking and knowledge management, and that neglect of these limitations is likely to cause muddy situations. In order to be effective, information technology has to be complemented by organizational procedures that account for the inherent limitations of technology. The problem with outdated information, for example, could be addressed by implementing rigorous procedures where to publish information on the internal Web and how to update or replace outdated information. However, it would be naive to assume that all problems could be solved in such a way. For example, the existence of contradictory or inconsistent information can hardly be resolved in a company as large as the one investigated.

The unity of information technology and organizational procedures is brittle unless embedded in an appropriate social environment. The complexity of creating and maintaining an appropriate social environment, however, is often underestimated. We propose to consider communities of practice as an appropriate context to establish such an embedding.

The "communities of practice" [22] metaphor has become increasingly popular in modern knowledge management as it offers perspectives on sharing knowledge that is hard to express, such as "soft" or "tacit" knowledge, such as knowledge about roles and procedures. What has received relatively little attention so far is regarding communities of practice as a context for sharing information and sharing of information *about* information, such as knowing how to find the right information in a company, knowing places where important information is published, who is publishing important information, and so on. A recent study indicates that communities of practice can exist in multinational, physically distributed companies [7].

Supporting communities of practice imposes different requirements on information technology compared to traditional document-based approaches. Tools are needed that do not only allow the retrieval of documents but that reflect the status of documents within the community. Active collaborative filtering [11], for example, could support the dissemination of important documents. Personalized social navigation [10], i.e., navigational information reflecting the behavior of particular persons within an information ecology

[15], could be implemented to address the requirements. These and similar technologies are around in the HCI and CSCW communities but did not make their way to business-oriented information technology. Furthermore, simple, stationary SW-agents aware of social role structures [17], designed according to the tool paradigm in communication theory [16] and able to share and exchange information [18] might help to bridge context gaps between providers and consumers of information, as they act in effigy of the users. The important point with all these technologies is, however, that it is not only the technology but the intent to share information and to learn about other's information needs, and thus the development of an information-sharing culture.

5 Conclusions

In this paper, we have argued that information seeking is a knowledge-intensive activity that is relevant to knowledge management, and that supporting this activity in the context of knowledge management demands a paradigm change in information technology. Investigations of information technology in a large company support the point and suggest to consider new technologies, such as active collaborative filtering or personalized social navigation.

References

[1] Liam J. Bannon. From human factors to human actors: The role of psychology and human-computer interaction studies in system-design. In Joan Greenbaum and Morten Kyng, editors, *Design at Work: Cooperative Design of Computer Systems*, pages 25–44. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1991.

[2] UK Central IT Unit, Cabinet Office. Channels policy, 2000. web-site: www.citu.gov.uk/channels/channelspolicy.htm.

[3] Kate Ehrlich and Debra Cash. Turning information into knowledge: Information finding as collaborative activity. In *Proceedings of the First Annual Conference on the Theory and Practice of Digital Libraries*, pages 119–125, June 1994.

[4] Thomas J. Froehlich. Relevance reconsidered. *Journal of the American Society for Information Science (JASIS)*, 45(3):124–134, 1994. Introduction of the Special Issue on Relevance Research.

[5] R.H.R. Harper. Information that counts: A sociological view of information navigation. In Alan J. Munro, Kristina Höök, and David Benyon, editors, *Social Navigation of Information Space*, pages 81–89. Springer, London, 1999.

[6] Urs Hengartner, Christos Kefos, Martina Klose, Ulrike Lechner, Elisabeth Maier, Olga Miler, Lutz Richter, Reinhard Riedl, and Beat Schmid. Organization on intranet: An agent based approach. In Bonnie A. Nardi, editor, *Proceedings of the Information Resource Management Conference 2000*, Spring 2000.

[7] Paul Hildreth, Chris Kimble, and Peter Wright. Communities of practice in the distributed international environment. *Journal of Knowledge Management*, 4(1):27–37, 2000.

[8] Robert R. Korfhage. *Information Storage and Retrieval*. John Wiley & Sons, New York, 1997.

[9] Jean Lave and Etienne Wenger. *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press, Cambridge, 1991.

[10] Christopher Lueg. *Supporting Situated Information Seeking: Communication, Interaction, and Collaboration*. PhD thesis, Faculty of Science, University of Zurich, Switzerland, 1999.

[11] David Maltz and Kate Ehrlich. Pointing the way: Active collaborative filtering. In *Proceedings of the Annual ACM SIGCHI Conference on Human Factors in Computing Systems (CHI'95)*, pages 202–209. ACM, May 1995.

[12] Gary Marchionini. *Information Seeking in Electronic Environments*. Cambridge University Press, Cambridge, 1995.

[13] Jan Matousek. *Werkzeuge zum Internet Profiling*. Diplomarbeit, Department of Information Technology, University of Zurich, 2000.

[14] Humberto R. Maturana. Biology of cognition. In H. Maturana and F. Varela, editors, *Autopoiesis and Cognition: The Realization of the Living*, volume 42. D. Reidel Publishing Company, Dordrecht, Holland, 1980.

[15] Bonnie A. Nardi and Vicki L. O'Day. *Information Ecologies*. MIT Press, Cambridge, 1999.

[16] Reinhard Riedl. Agents for customer support in electronic commerce. In *Mile High Expectations: Proceedings of the ACM SIGUCCS Conference, Denver Colorado*. ACM-Press, November 1999.

[17] Reinhard Riedl. Usage of trace data for the deduction of role structures and comparison of knowledge societies. In *Proceedings of the Third International Conference on Cognitive Technology, San Francisco, USA*, August 1999.

[18] Reinhard Riedl. Agent views of information markets. In *Proceedings of the 4th International Conference on systems, Cybernetics, and Informatics, Orlando, Florida*, August 2000.

[19] Reinhard Riedl. "need for database trace benchmarks". In Rudi Eigenmann, editor, *Performance Evaluation with Realistic Applications*. MIT Press, to appear, 2001.

[20] Reinhard Riedl and Takashi Suezawa. Management of information markets with mobile agents. In *Proceedings of EMMSEC '99, Stockholm*, June 1999.

[21] Heinz von Foerster. Thoughts and notes on cognition. In Paul L. Garvin, editor, *Cognition: A Multiple View*, pages 25–48. Spartan Books, New York, NY, 1970.

[22] Etienne Wenger. *Communities of Practice: Learning, Meaning, and Identity*. Cambridge University Press, Cambridge, UK, 1998. First Paperback Edition 1999.