Collaborative Aspects of Information Retrieval Tools:

Summarising three action case studies

Dick Stenmark
it.dixi@memo.volvo.se
Volvo Information Technology, Department 9734 HD3N,
SE-40508 Göteborg, Sweden, phone: +46 31 765 7037
Knowledge Management Group, Viktoria Institute, Göteborg University,
P. O. Box 620, SE-40530 Göteborg, Sweden, phone: +46 31 773 5566

Abstract

Information retrieval (IR) tools are engaged to find information in electronic form, in databases, digital libraries, and networks. However, the advent of the web technology has boosted the amount of information available to us way beyond our cognitive capacity. The importance of good IR tools has thus increased and the tools have become natural focal points for organisational information seeking. This situation offers a possibility to exploit the collaborative efforts of the individual users. However, is the opportunity utilised? In action case studies of a search engine, a notifier program, and a recommender system on a corporate intranet, this study seeks the answer to that question. The conclusion is that most IR tools have neglected to incorporate support for collaborative action, or allowing awareness of other users with similar information needs. By integrating functionality from all three tools, according to the design implications presented here, future systems should be able to benefit more from the collective information seeking efforts of the organisation.

Keywords: information retrieval, collaboration, awareness, intranet

1. Tools for information awareness

Corporate intranets have become major channels for intra-organisational information, and large sources of corporate-specific knowledge. However, as the amount of information available on the net continues to grow, the level of awareness of what goes on within the organisation decreases. Only a fraction of the available information will be relevant to any particular employee and the problem will therefore be to pan out the nuggets while minimising the need to search through irrelevant information (Foltz & Dumais 1992). Whenever information technology (IT) is used to present information, other IT tools may be engaged to assist the users in their filtering or searching efforts, and various forms of technology have been developed to address these situations. Alongside more well-known tools such as search engines and web notifiers, recommender systems has begun to gain acceptance.

From awareness and collaborative points of view it remains unclear, though, whether these tools cover the same ground, partly overlap, or complement each other. Though information retrieval (IR) and information filtering are said to be two sides of the same coin

Proceedings of IRIS 23. Laboratorium for Interaction Technology, University of Trollhättan Uddevalla, 2000.
L. Svensson, U. Snis, C. Sørensen, H. Fägerlind, T. Lindroth, M. Magnusson, C. Östlund (eds.)

(Belkin & Croft 1992), technology performing such operations has thus far been evaluated separately and often only from a strict IR perspective. The intention with this work is thus to compare these approaches not within their respective categories, but relative each other in order to *a) establish how well they support information awareness*; *b) pinpoint their collaborative potential*, and; *c) derive design implications for future development and integration of such systems.*

Since many different types of recommender systems exists today (cf. Resnick & Varian 1997), it might be appropriate to clarify what the term means. In this article a recommender systems is understood as a piece of software able to suggest concrete objects to a known user, based on "knowledge" of that user's interests. Three different approaches may be identified: One is to present each user with a number of objects that the user should *rate or rank*. Once a sufficient number of objects have been rated, they are compared with that of other users until a certain level of likeness is met. Objects found interesting by one user but not yet rated by another user is then "recommended" to the other user. This sort of collaborative filtering was used in HOMR (later known as Firefly) (Shardanand & Maes 1995).

A second approach is the social filtering seen e.g. at amazon.com, where recommendations are based on *actions* of similar anonymous users. If you show an interest in books by Hemmingway you will be recommended books bought by other people who also happened to buy Hemmingway. Thirdly, the recommendations may be based entirely on computed similarity between the object themselves, i.e. *content-based* similarity. Document can then be treated as vectors in a high-dimensional semantic space and similarity is based on the angle between such vectors (Salton 1989), or advanced pattern matching algorithms (e.g. Autonomy 1998). The third approach was chosen for this study since it avoids the "cold start problem" (Maltz & Ehrlich 1995), i.e. it does not require a critical mass of users to pay off but provides benefit also for the very first user.

Finally, recommender systems can be combined with *agent technology*. If included, such technology enables a certain degree of autonomy that allows continuos monitoring without user intervention. Unlike the case with amazon.com, an agent-based recommender system would detect updates to the information corpus regardless of whether or not its user is present. Updates would be accumulated and communicated to the user whenever he or she comes online.

In the following sections the research methodology used will first be introduced before reporting from the three case studies, separately. The result as a whole will then be summarised and discussed from an awareness and collaborative point of view.

2. Research approach

My mission is two-fold: First, to understand user behaviour in order to inform the design of new IT artefacts, and second, to intervene in, and thus influence, the way in which IT is used in the researched organisations. To achieve the first objective I have used mainly qualitatively oriented analysis methods, though the data has been collected using a mixture of techniques. The latter includes both site observations and qualitative interviews, and data from questionnaires and application log files. The material has been analysed and categorised using grounded theory approach.

To intervene and initiate organisational changes, a more action research-oriented

approach is required, but since the scale of my intervention is rather small, I find Vidgen and Braa's (1997) notion of "action case" to be a more appropriate description of my approach. Action case is a hybrid of understanding and change, designed to balance the trade-offs between being either an observer capable of making interpretations or a re-searcher involved in creating change in practice (Vidgen & Braa 1997).

Being a single researcher with limited resources, the organisational constraints of a large corporation make full-scale action research projects difficult. *Action case* research acknowledges this dilemma and instead suggests that small-scale interventions based on deep conceptual understanding should provide a more pragmatic and feasible approach.

This paper reports from empirical fieldwork taking place at Volvo Information Technology, an IT service provider within the Volvo Group. The Volvo intranet was implemented in 1995. Three related action case studies were conducted over a period of twenty months. First I studied the use of the previously implemented search engine usage and recorded its strengths and weaknesses. Thereafter, I installed a notifier program that given a URL would notify the user whenever the corresponding page was updated. I finally implemented an agent-based recommender system, which allowed users to monitor areas of interest and detect new or updated information as well as find peers with similar interests. Table 1 summarises the three case studies.

Table 1: Time frames, intranet size, test group sizes, and data collection methods for the three case studies.

	Search engine	**Notifier program**	**Recommender system**
Test period	January-May 1997.	February-March 1998.	August-December 1998.
Intranet size	100 servers, 40,000 documents, 13,000 connected users (May 1997).	210 servers, 130,000 documents, 21,000 connected users (April 1998).	+400 servers, +300,000 documents, +40,000 connected users (Dec. 1998).
Test group size	4,719 company-wide users.	38 registered test users.	48 voluntary test users.
Data collection methods	Analysis of search engine log files from January to April (29,391 queries). Eight in-depth interviews (32-48 minutes) conducted with randomly selected users in May 1997.	Log files from February to March were analysed. I received 21 answers to an email questionnaire asking for their experiences, and any suggestions for design improvements.	Group interview with eight participants. An emailed questionnaire answered by another 12 test users. Seven open-ended interviews lasting between 28 and 66 minutes.

All participating users were experienced computer users with access to individually assigned

PCs, and their roles varied from technicians and system developers to content providers and administrators.

3. Results and discussions

I shall in this section report the results from each of the three cases; the search engine, the notifier program, and the recommender system, and also discuss these results individually.

3.1 Search engine

Search engines were well known IR tools, and every interviewee had personal experi-ences from both public internet search engines, and from using the previously installed Harvest (Bowman *et al*. 1994) search engine

3.1.1 Results

By analysing the search engine log files, I found that queries had been submitted from 4,719 unique IP-addresses during the four months test period. All together 29,391 queries were submitted and though the most active user submitted 213 queries, a large portion of the users (25.9% or 1,221 individuals) only submitted one query. Almost 83% (3,903 users) accessed the search engines less than 10 times.

Single keyword queries were by far the most common ones with 17,157 entries or 58.4% of the total. The second largest group was the multi keyword group, where two or more keywords or phrases were entered without any Boolean operators. Harvest would interpret such a query as a Boolean AND. This group had 5,259 queries (17.9%). Note that these queries were not “natural language”-style queries (Turtle 1994), but simply several consecutive “keywords”.

Though Harvest only offered rudimentary search capabilities, a large number of users appointed the searching capabilities to be the major benefit of the intranet. The interviews revealed that though claiming to be familiar with search engines, the users were uncertain of what query syntax to use. Despite the fact that six of the eight inter-viewed users admitted that they needed help, the build-in help function was very seldom used. Log file analysis reveals that only 4.0% of the visitors (189 users) requested the help page, which stated that Boolean AND/OR expressions should be used. Several examples, including nested Boolean expressions where also available.

The interviews consistently showed that the users perceived the result set as overwhelmingly large. All eight interviewees commented on the large amount of docu-ments typically being the result of a query, and five of these further complained about the presentation of the results, the ranking in particular, which was described as “invisible”, “mysterious” or “odd”, and thus adding to the problem. The users seldom bothered to look past the first few items in the result list. As seen from the log files, only after 1,552 queries, or 5.3%, of the total 29,391 queries submitted, did the user click on the Next page button.

3.1.2 Discussion

Many queries generate a very large result set, partly due to the heavy use of single key-word questions, which in turn stems from the inability to formulate queries. To be efficient, searching

requires the user to enter a very precise search string, while the true interests are often vague or difficult to express. An IR tool such as a search engine is best suited for quick-and-dirty-queries when the user has a very precise problem at hand (Belkin & Croft 1992). Having "smart" interfaces that allow for alternative query syntax or actively guide the user through the query process would address the problem with ambiguous query interpretations and thus be helpful. Such a dynamic and interactive interface would help the users to understand the correlation between query and result (Shneiderman *et al*. 1998).

The observation that users seldom look beyond the first items in the result set is consistent with the research of Nielsen (1996) who states that "only 10% of users scroll beyond the information that is visible on the screen when a page comes up". This sugg-ests that rating is more important than precision. As long as the "good" result scores highly in the ranking it does not matter if 100,000 results are returned. However, the proprietary ranking process is difficult for the user to understand. Though the use of visualisation techniques to communicate a richer meaning of the ranking to the user has been suggested (Rao *et al*. 1995), it has not been implemented in any commonly used tool.

From a collaborative point of view, search engines are under-utilised. The fact that many users search for the same or similar topics could have been used to enhance the results (by noting what previous users clicked on) or to facilitate the establishing of online communities (by indicating who those previous users were). The fact that the users did not explicitly ask for these features does not necessarily mean that they would not have benefited from them.

3.2 Notifier program

A notifier is a program that monitors a specified web page and notifies the user when the page is changed, updated, removed, or relocated. There are a number of different variants of this technique available. They mainly differ in the way they alert the user (via email or by updating a web page) and where execution takes place (on a central server or on the local client). I used the server- and email-based Enterprise Minder (EM) from NetMinder for my main test. EM was announced on the intranet as a temporary service available free of charge for anyone to try in February 1998.

3.2.1 Results

The 38 users monitored a total of 47 intranet pages and nine internet pages during the test. The page monitored by the highest number of users was Volvo IT's homepage with 13 entries. The "CEO on-line", a page where the Volvo Group's CEO makes strategic declarations once in a while, was second with six and the Volvo Group homepage scored four in third place. The remaining pages were each monitored by one or two users.

The most positive aspect noted by the users was the elimination of the need to check a page just-in-case. Being notified when an interesting site had been updated was highly appreciated by all 21 users answering. Fifteen users claimed that not having to manually check saved them time. Seven users answered that the possibility for long-time coverage was a great addition.

On the negative side 16 users complained about being notified by email, since they felt they already received too much email. Instead, they suggested that their bookmarks should be automatically monitored and marked with "Updated"-icons whenever a change had occurred, or that links should be added to a "personal news page". Ten users complained about having to

go to a particular “Add URL” page to have a page monitored. Eight of these instead suggested that a “Click here to have this page monitored”-button should be added to “interesting” pages. From examining the log files, it can be noted that many typing errors were made while entering the URLs.

3.2.2 Discussion

As noted in previous research, not being able to detect changes is one of the biggest obstacles for the web to function as the only information channel (Stenmark 1998, 1999a), and a notifier program thus servers a purpose. One limitation with the notifier approach, though, is that it requires the user to first find the interesting page before it can be monitored.

Pages updated very frequently, i.e. daily or even more often, are not well suited for being monitored with a notifier program. Users monitoring the Volvo IT homepage, which contained news flashes updated many times daily, received up to 10 emails per day from that site only. The fact that the IT company’s homepage remained the most frequently monitored page during the test indicates that more information and training about how to use new tools is important.

The most useful setting for a notifier program is a web site with a lower update frequency and a more narrow audience, which suggests that information should be sepa-rated to different pages, each covering a topic of its own. This will enable the users to select only the page(s) that matches their interests thus providing a more targeted in-formation flow. Such an approach will however add to the burden of the content pro-vider, and Ackerman *et al*. (1997) instead suggest that the notifier program itself should be able to deduct whether or not the change is relevant before alerting.

A third and final problem is that of notification. Email is not always desired due to the large amount of emails the users are already receiving (Stenmark 1998), and there are less obtrusive alternatives. For example, the client- and web-based WBI (Barrett *et al*. 1997) instead hooks into the browser and notifies the user by flashing an icon. However, such a solution requires additional software to be installed on the client, which some organisations find undesirable. However, a full utilisation of a notifier tool would pro-bably reduce the total number of email, since the monitoring capabilities would enable much information today communicated via email to be transferred to the intranet.

Some users also suggested that a “Have me monitored”-button should be placed on pages of general interest. WBI, which works as a proxy, does actually add such a button to every page, making sub-mission very easy. Such a solution also eliminates the very common typing errors and does thus provide an added value.

Though being a server-based solution, EM did not exploit the fact that several users were monitoring the same pages. The fact that at least one human being (other than the author) has found a page interesting could be used as an indication of quality. Users monitoring the same page can further be assumed to have common interests, which might be a starting point for collaboration. As for search engines, no user mentioned this possibility or asked for this kind of support.

3.3 Recommender system

The design of my prototype system was influenced by DICA, the “Do-I-Care Agent” (Ackerman *et al*. 1997). DICA uses the natural incentive that exists for a user to train and

maintain a profile in order to receive a higher degree of "interesting" information, and combines this with an effortless collaboration mechanism. The users can improve the agents' predictions by providing relevance feedback. These features are available also in my prototype. However, my approach differs in two aspects: First, I concentrate on intranet usage, and due to its limited size (relative to the Internet), a spider can crawl the entire information corpus and thus monitor not only known information sources but also detect new information. To define what to search for the user had to provide a text example, either by writing a few sentences or, preferably, by pasting in a (chunk of a) previously found document. A second addition is that my prototype enables users to find other users with similar profiles and does thus facilitate the establishing of communities of interests. Details about the prototype design are found elsewhere (Stenmark 1999b, 2000).

3.3.1 Results

Overall, all 27 responding users were positive. Nineteen users claimed the prototype to be useful and though eight users did however not consider the prototype useful in its current state, they believed that a future version would be able to deliver and referred to the technology as "promising". The most frequently reported reasons for their positive attitude were that it was easier to construct queries and that it saved time not having to search actively.

More specifically, seven users explicitly expressed their appreciation of not having to come up with descriptive keywords. However, the prototype interaction differed from that of the more familiar search engine and the users have problem adjusting. Three users believed that the way it actually worked had to be better appreciated in order for them to get something useful out of it.

Despite the general claims that these sorts of recommender tools were welcomed and appreciated, many users experienced mainly negative results. A majority of the users (15 of 27) reported what they referred to as "strange" or "unexpected" document recom-mendations. However, the users tended to blame these bad results on their own inability rather than on the prototype system.

The Similar Agents feature, which located users who had defined similar agents (and thus had similar interests), was frequently used: "It's really interesting to see who else is searching for this sort of things", one of the users commented. Many respondents reported that they were surprised to find certain people sharing their interests (6 users), or that the Similar Agents feature returned users whom they had not expected to be in-terested in a particular topic (4 users). Note that these comments were not uttered in a negative way. The users regarded these results as useful new insights and no one questi-oned the correctness of the results. However, six users asked for the opportunity to search not for similar users but for users with other specific profiles, e.g. interests such as data-bases or Java programming.

Five users required better navigation aids within the prototype. Five (partly diffe-rent) users asked for an alternative way to present the results from the prototype, giving a better overview of the findings. One suggestion was to have a personal homepage with the most updated results displayed per agent. The users, however, did not think that a notification via email would be a good solution.

3.3.2 Discussion

The possibility to use a richer query representation, i.e. whole documents rather than a few selected keywords, was said to be appreciated. However, in reality the users often re-turned to

using a few keywords. A feasible explanation for this contradictory behaviour is that the users were more familiar with keyword-based search engines and their less-to-type querying style. This fact may also explain some of the low-quality results from their agents. This shows that the introduction of new IT artefacts not only requires sufficient information and training, but also lower thresholds for new behaviour. It also connects with the previous discussion of interface design for navigation and presentation, and underlines the need for more research in this area.

Having agents constantly monitoring the net for interesting updates obviously off-loaded the users in a positive way, which was not achieved with the previously tested tools. A higher level of awareness (within the fields specified) may thus be obtained to the rather low cost of setting up and (initially) train an agent. The users also obtained awareness of previously unknown colleagues with similar interests, which opens up for collaboration. In addition to this explicit form of collaborative support, the system also benefits from the implicit collaboration that occurs when the users give feedback to their individual agents. This feedback is accumulated by the underlying pattern-matching algorithm and helps improve the quality of the predictions.

4. Conclusions

We have examined three different tools for intranet information awareness, and noted various strengths and weaknesses for each of these tools as summarised in table 2. The conclusion is that no one technology alone can provide the users with sufficient awareness of both changing information and other employees' interests and activities. An integration of functions from the different tools is thus desirable.

Searching and monitoring serves two different purposes that both need to be supported. A search engine is best suited for quick-and-dirty queries not likely to be repeated frequently, but it cannot provide long-term coverage of a field of interest. A notifier pro-gram, on the other hand, can detect updates to particular web pages and notify interested users, who thus need not actively and repeatedly check the page just in case. However, a notifier program requires the user to manually detect the page in the first place - a non-trivial task in itself. Neither of these technologies does today support awareness of know-ledge that resides *within people*, and does neither promote or even acknowledge collabo-ration.

Recommender systems recognise the value of social interactions and are thus better equipped for collaboration, while agent technology enables autonomous operation capability allowing the software to continuously roam the net in search of information on the users' behalf. Pairing these approaches, we get a system that, though far from being perfect, has a large potential for both facilitating collaboration and supporting awareness.

Though the three examined technologies are related, they also represent different approaches to the field. With the rapid changes and high pace in the IT industry, vendors have a busy time exploiting and developing *one* approach or technology. It is therefore unlikely that we will see a complete solution from any one vendor. An alternative and more plausible scenario is the emergence of open standards that enable different vendors to link and integrate their products, similar to what we have seen in the web browser industry.

Table 2: Pros (+) and Cons (-) for Search engines, Notifier programs, and Recommender systems.

	Search engines	Notifier programs	Recommender systems
Active awareness	+ support for electronic sources	- no support	+ support for both human and electronic sources
Passive awareness	- no support	+ support for selected resources	+ full support for selected areas
Collaboration	- no collaborative support	- no collaborative support	+ both direct and indirect utilisation
Ad hoc queries	+ quick and dirty queries	- no support	- no support
Monitoring	- no long term coverage	+ supports page monitoring	+ supports area monitoring
Input interface	- proprietary query syntax	+ can be "buttonised"	+ rich query representation
Result presentation	- massive output	- email output	+ relevant output, email notification optional
Information detection	- manual work	- manual work	+ automatic detection (agent version)
Ease of use	+ intuitive (on "simple" level)	+ intuitive	- requires training (agent version)

4.1 Support of information awareness

Search engines supports awareness inasmuch as they return and display information matching a keyword query. However, the facts that users seldom examine material beyond the first ten or so items, that information is only gathered and presented as the result of a direct query, and that information about other users with the same needs or interests is not communicated makes the usefulness limited.

Notifier programs support awareness by providing autonomously monitoring of requested resources. Though useful, the limitation of this approach is obviously that the user has to be aware of the resource in the first place, and that awareness of like-minded peers is not supported.

Agent-based recommender systems provides awareness of *information* by keeping the user up-to-date regarding added or updated information within a user-specified field of interests. In addition, it fosters an awareness of other *employees* with similar interests, by comparing search profiles.

It may be argued that the line between monitoring behaviour to promote awareness on the one hand and surveillance and loss of personal integrity on the other is thin and requires a fine-tuned balancing act. This is true, but less of a problem on a corporate intranet than it would have been on the Internet. First, employees have access to the intranet during working hours, from their office location, while performing paid services. They are not supposed to do anything controversial. Second, the intranet contains (almost) exclusively work-related material, which the employees not only are *allowed* to browse but *encouraged* to read. No illegal, immoral, or unethical material is to be found. The participants in our study expressed no concern for being caught with an interest in COBOL or BPR. However, they *did* express concern for being overloaded with questions as the result of having a non-anonymous agent.

4.2 Collaborative capabilities

Being server-based, all three technologies have obvious potentials for collaboration, but neither search engines nor notifier programs do generally commodify this potential. However, the users did not comment on this missing feature or in any way suggest that they experienced this absence as a problem. The users did in fact not seem to think of search-ing or monitoring as collaborative tasks, nor realise the added value such support might be able to offer. Their thinking seemed to stay entirely within the established single-user tradition of such tools.

Recommender systems on the other hand usually offer both direct and indirect collaborative capabilities. Features such as sharing of results and cloning of agents support direct collaboration, while indirect collaboration is achieved by the systems ability to increase its prediction accuracy based on the collective feedback from its users. By also fostering an awareness of other colleagues the possibilities and probabilities for collaboration increases.

However, the literature is full of examples of how technically excellent solutions have failed due to not having pay enough attention to the organisational culture in which they were implemented. In a competitive environment people are likely to keep their information and knowledge to themselves regardless of any collaboration-supporting IT devices.

4.3 Design implications for future systems.

The implications drawn from this study are the following:

1. The technologies examined do partly overlap, but each has functionality not entirely provided by the others. Instead of forcing the users to chose between the tools, the technologies should be integrated to provide a single point of interaction. Future tools should ideally satisfy both short-time and long-time information needs by allowing both ad hoc quick-and-dirty querying and more structured field-of-interest coverage.
2. IT systems for information awareness should address not only the explicit information electronically available in online documents but also the implicit information and knowledge that exists in non-textual form. Facilitating and supporting identification of, and collaboration between, organisational members will effectively leverage the organisational knowledge base.
3. Adaptive software such as agents-based systems behaves differently from the traditi-onal direct-manipulative programs with which most users are familiar. When designing new tools care must thus be taken not to clash with existing mindsets of the users and not to increase their cognitive load by adding new and confusing interfaces and interaction modes. Unless a richer representation is actively afforded, single keyword queries will continue to be used

4. Navigation and presentation interfaces on the web are still issues that remain problem-atic. Results need to be presented in a way that gives a good overview and is intuitive, while yet being non-intrusive. This presents a delicate balancing act not mastered by any of the examined technologies.

5. Acknowledgement

I am grateful to Volvo Information Technology for their support during the fieldwork. This work has been partly funded by Volvo Research Foundation, Volvo Educational Foundation, and Dr. P.G. Gyllenhammar Research Foundation through the project "Genres of Organizational Learning". Thanks are also due to Jan Ljungberg at Viktoria and to anonymous IRIS reviewers for useful comments on earlier versions of this text.

6. References

Ackerman, M., Starr, B., and Pazzani, M. (1997). "The Do-I-Care Agent: Effective Social Discovery and Filtering on the Web", In Proceedings of RIAO '97.

Autonomy, Inc. (1998). "Agentware Technology Whitepaper", (http://www.autonomy.com/tech/whitepaper.pdf).

Barrett, R., Maglio, P. P., and Kellem, D. C. (1997). "How to personalize the web", In Proceedings of CHI '97, ACM Press: Atlanta, GA.

Bowman, C. M., Danzig, P. B., Manber, U., and Schwartz, M. E. (1994). "Scalable Internet Resource Discovery: Research Problems and Approaches", Communication of the ACM, vol. 37, no. 8.

Belkin, N. J. and Croft, W. B. (1992). "Information Retrieval and Information Filtering: Two Sides of the Same Coin?", Communications of the ACM, vol. 35, no. 12.

Foltz, P. and Dumais, S. (1992). "Personalized Information Delivery: An Analysis of Information Filtering Methods", Communications of the ACM, vol. 35, no. 12.

Maltz, D. and Ehrlich, K. (1995). "Pointing the way: Active collaborative filtering", In Proceedings of CHI '95, ACM Press: Denver, CO.

Nielsen, J. (1996). "Top Ten Mistakes in Web Design", The Alertbox, Jacob Nielsen's bi-weekly web column, May 1996, (http://www.useit.com/alertbox/9605.html).

Rao, R., Pedersen, J., Hearst, M., Mackinlay, J., Card, S., Masinter, L., Halvorsen, P.-K., and Robertson, G. (1995). "Rich interaction in the digital library", Communications of the ACM, vol. 38, no. 4, pp 29-39.

Resnick, P. and Varian, H. R. (1997). "Recommender systems", Communications of the ACM, vol. 40, no. 3.

Salton, G. (1989). Automatic Text Processing: The Transformation, Analysis, and Retrieval of Information by Computer, Addison-Wesley: Reading, MA.

Shardanand, U. and Maes, P. (1995). "Social Information Filtering: Algorithms for Automating 'Word of Mouth' ", In Proceedings of the CHI-95, ACM Press: Denver, CO.

Shneiderman, B., Byrd, D., and Croft, W. B. (1998). "Sorting Out Searching. A User-Interface Framework for Text Searches", Communications of the ACM, Vol. 41, No. 4, pp 95-98.

Stenmark, D. (1998). "Identifying Problems with Email-based Information Sharing", In Buch *et al.* (Eds.), Proceedings of IRIS21, Department of Computer Science. Denmark: Aalborg University.

Stenmark, D. (1999a). "The duality of email as corporate information channel", Internal Communication, Issue 49, October 1999, pp 6-8.

Stenmark, D. (1999b). “Using Agents to Capture Tacit Knowledge”, In Proceedings of WebNet ’99, AACE Press: Oahu, HI.
Stenmark, D. (2000). “Turning Tacit Knowledge Tangible”, In Proceedings of HICSS-33, IEEE Press: Maui, HI.
Turtle, H. R. (1994). “Natural Language vs. Boolean query evaluation: A comparison of retrieval performance”, In Proceedings of SIGIR ‘94, Springer Verlag: Dublin, Ireland, pp 212-220
Vidgen, R. and Braa, K. (1997), “Balancing interpretation and intervention in information systems research: the “action case” approach”, In Lee *et al.* (Eds.) Proceedings of IFIP WG8.2, Philadelphia, USA.