

Users' Perception of the Performance of a Filtering System

Raya Fidel
Graduate School of Library and Information Science
University of Washington
Seattle, WA 98195

Michael Crandall
Boeing Technical Libraries
P.O. Box 3707
Seattle, WA 98124-2207

Abstract. Although filtering electronic information is spreading rapidly, very few studies examined users' perceptions about the success of filtering. Users at the Boeing Company participated in a study which collected data through observation, verbal protocols, questionnaire, and interviews. Data analysis used four levels of relevance to assess the importance, and frequency of use, of thirteen criteria for relevance, and fourteen for non relevance, that are not topics or subject matters. Results showed that perceived precision ratios for filtered information were higher than the ratios for non-filtered information, but not significantly and could still be improved even though most respondents were satisfied with these ratios. Developing methods to create and maintain useful profiles, and finding ways to incorporate relevance as well as non-relevance criteria into profiles, are necessary to improve the performance of filtering mechanisms.

the topic (Maure, 1992). These are only a few examples. The essential role of filtering mechanisms has also been recognized by systems designers and developers who have proposed various algorithms and interface agents for information filtering (e.g., Anick et al., 1991; Maes, 1994; Shuldberg et al., 1993; Yan & Garcia-Molina, 1994).

As Belkin & Croft (1992) explain, filtering information differs from information retrieval (IR) for a search request in several ways. Filtering is designed to deal with an incoming stream of unstructured, or semistructured, data and implies removing data, while IR deals with a search of a remote database with highly structured data and implies finding data. Further, filtering is based on individual or group profiles that may change in time but that typically represent continuing interest whereas IR is based on a momentary information need.

Filtering electronic information is a vital component of contemporary information work. The rapid development of information technology affected not only the way users look for information and retrieve it, but also organizational design, intelligence, and management decision making (Huber, 1990; Dou, et al., 1993; Goodman, 1993). In addition, with the increasing amount of information that is available electronically, the need to develop mechanisms to filter what a user receives has become of paramount importance.

To filter information requires building models of users' interests, or filtering profiles. Such models are difficult to build because of semantic and contextual complexities, and because users' interests are constantly changing (Stadnyk & Kass, 1992). It is important, therefore, to study the effectiveness of filtering systems from the users' point of view. To date, very few user studies have been carried out (Gant, 1995). The study reported here examined users' perceptions of the performance of a filtering system in a real-life situation.

This need has been recognized by professionals in various fields and on all levels: *The Wall Street Journal* alerted its readers to fight back data overload (Bulkely, 1989); medical librarians have been examining filtering clinical information for a number of years (Marshall, 1993); and the ACM dedicated a whole issue of its *Communications* to

The main goal of the study was to explore how users selected their preferred method for filtering electronic information. Among the specific objective were: (a) to compare users' perceptions of system performance for filtered and unfiltered information; and (b) to explore those filtering criteria users employed that did not relate directly to subject matter, and therefore might be relatively stable.

Permission to make digital/hard copies of all or part of this material for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copyright is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires specific permission and/or fee

SIGIR 97 Philadelphia PA, USA

Copyright 1997 ACM 0-89791-836-3/97/7...\$3.50

The study was carried out at the Boeing Company in the Puget Sound area of Washington. During 1991-93 the Technical Library worked with the publisher of the *Gartner Group Reports* and various Boeing user groups to

establish a company-wide contract for the *Reports* which included electronic delivery of the text to any Boeing employee.

The *Gartner Group Reports* is a monthly publication, much like a professional magazine, with news and analyses relating to the computer industry. During the study period, the *Reports* were delivered to the Company via the Internet each month, and distributed throughout the Puget Sound area through two mechanisms. The first was an unfiltered bulletin-board type system using a Boeing internal variation of the Unix Newsgroup model, which enabled users to subscribe to the *Gartner Group Reports* and to other publications. The second distribution allowed users to establish a subject profile (a model of their interests) through an intermediary librarian in the Technical Library. They then received only those reports which met their profile criteria via their e-mail system as ordinary mail messages. Profiles were developed and maintained using Verity's TOPIC information filtering software which uses term weighting, Boolean operators, proximity operators and field searching (Verity, 1997).

Research Method

The study employed a variety of methods and instruments to investigate the same process. It used a combination of qualitative and quantitative methods, including data collection through observation, verbal protocols, questionnaires, and interviews.

The study had three phases. The first phase used observation of and interviews with selected users to determine the patterns of searching behavior, the factors perceived to be important for the selection of a filtering method, and the elements of perceived satisfaction. During the second phase, these data were analyzed and used to design a questionnaire which was administered to all users. The third phase included statistical analyses of the data that had been collected from the questionnaires and interviews to support the interpretation of the statistical results.

Typical of qualitative research, we began the study with a comprehensive and in-depth analysis of the filtering preferences of a small sample of users. Thus, in the first phase we observed a total of 15 users, both engineers and managers working in computing support and systems architecture, as they examined the reports they received via e-mail or on the Newsgroup. We observed a total of 34 sessions. Four of the participants received unfiltered information and we observed each one only once. The other participants received filtered reports, and we observed most of them during three consecutive deliveries of the *Reports*.

Participants were asked first to explain why they looked for information and then to think aloud as they viewed the information on the screen. At the end of each session, users evaluated the session and its results. At the end of the observation period, we interviewed each participant to confirm our interpretation of the individual's searching behavior, and to further investigate the reasons that led him or her to select a filtering method. All verbal protocols, think-aloud as well as interviews, were audio taped and transcribed.

Participants were extremely cooperative. They were very generous with their time and answered questions we asked in great detail. Most of them found it comfortable to think aloud while browsing the *Reports*. Generally, they liked to explain what they did and the reasons behind their decisions and actions.

In the second phase, we analyzed the data from the transcribed verbal protocols to identify the various factors that participants considered in the selection of filtering method, and in assessing their satisfaction. The analysis uncovered four levels of relevance that the participants used, and a host of criteria beyond subject matter they employed to determine whether or not a document was relevant.

This analysis guided the development of a questionnaire that was distributed to all users. After a pilot test, the questionnaire was administered online: We attached it to the beginning of the next delivery of the *Gartner Group Reports* for all users who received them through the library filter, and posted it twice on the Newsgroup. On the day we posted the questionnaire, Boeing announced organizational changes that could have affected the potential respondents. Response rate from those users who received filtered information was only 30%, even though previous experience with Boeing users' indicated that under normal conditions we could have achieved over 60% response. In addition, because no list of the subscribers to the Newsgroup is available, we could not determine the accurate response rate for this group which comprised 15% of all the respondents. We received a total of 83 responses, and all were usable.

The third phase of the study included statistical analyses of the data collected through the questionnaires. Some of these results are presented next.

Relevance Levels

Through observation we recognized that the electronic reports had four distinct destinations, depending on their relevance to the participant's work. A participant might:

- Delete a report without reading, or
- Delete it after reading or skimming on the screen, or

- Save it for future, in-depth reading, or
 - Keep it for future use
- No other physical destination was discovered in this study in which participants assessed a total of over a thousand reports¹.

Regardless of the filtering and delivery method, participants always encountered first the titles of the *Reports*. If they received them through e-mail, the titles appeared in the "subject" line. Because the number of characters in the "subject" line is limited on most systems they used, the titles were incomplete at times. On the Newsgroup, the first display included the list of titles; participants had then to enter special commands to view the full text. The "delete without reading" destination refers to cases where participants decided to delete a report immediately after they examined the title. If a participant deleted a document after reading the Summary or Key Issues, after skimming a report, or even reading it on the screen, we denoted the case as "delete after reading or skimming." Sometimes after skimming a report, participants decided that it might require thorough reading, and they saved it for a more leisurely reading time. Such reports were usually saved electronically, although two participants who spent much time traveling printed them "to read on the plane." They might have remained in the original "box" or moved to a "to read" type of file. When a report was deemed essential to a project in which a participant was involved, it was usually moved to a specific folder for future use. Such reports were often printed on paper and kept in file folders designated for a particular project, or even for a certain meeting.

In the study we interpreted these four destinations as levels of relevance. We construed "deleting a report" (the first two destinations) as two distinct levels of "not relevant," and "saving or keeping a report" (the last two destinations) as levels of "relevant." We used these levels of relevance in the questionnaire for two purposes: (a) to find out which relevance/non-relevance criteria were perceived most important; and (b) to assess perceived precision ratios.

Relevance Criteria

Like most filtering mechanisms, the one used in the study was based on profiles that expressed topics or subjects. Respondents' profiles included words that might occur in the text of a report, subject categories that designated a subset of the reports, or some combination of both. The

¹ Most participants also forwarded reports to other Boeing employees. Because forwarding a document does not physically remove the document from one's box, we did not consider it a physical destination.

relatively low precision ratios found for some participants in the observation phase suggested the need to look for additional criteria that could be used for filtering, beyond topics or subjects. Supporting this approach were previous studies that were successful in identifying such criteria (e.g., Barry, 1994; Schamber, 1994), and the fact that the participants in the observation phase were highly articulate in expressing their reasons for accepting or rejecting reports. Thus, through the verbal protocols collected during the observation phase we identified the criteria participants used to express the relevance of reports, and those they used to explain why a report was not relevant.

Participants used a total of thirteen criteria to explain why a report was *relevant*. They also identified fourteen criteria they used to explain why a report was *not relevant*. These are listed in the boxes below.

Box 1. Criteria for Judging a Report Relevant

- It was relevant to the Boeing Company
- It was about a product or a service that related directly to a project the participant was working on
- It was about new concepts, products or services
- It was a case study
- It had hard data
- It displayed issues in a classified order and clearly (e.g., in the form of a list)
- It was written on a non-technical level
- It described industry trends or gave predictions
- It was about a specific vendor, product or service
- It confirmed or validated what the participant already knew
- It dealt with something the participant and his or her group had done
- It included background information or general information
- It had information that helped the participant keep up to date about a product with which he or she were familiar

Box 2. Criteria for Judging a Report Non-Relevant

- It was not relevant or applicable to the Boeing Company
- It was about something Boeing was already doing
- The participant had no influence over the issues the report raised
- The participant's group had already made a decision about the product or service that was addressed in the report
- It was about a technology that was not here yet

Box 2. Criteria for Judging a Report Non-Relevant

- It was completely non-technical (e.g., about law suits or company analysis)
- It was about specific vendors
- It raised questions but gave no answers
- It expressed opinions rather than presenting facts
- The participant was not familiar with the product or the service the report was about
- It did not tell the participant anything he or she did not already know
- It took too long to understand what the report was about
- It was too basic or too general
- It was too detailed or too technical

During the observation it became clear that some criteria were used more frequently than others, and that some were more important than others. In the questionnaire we asked respondents to mark all the criteria from the non-relevant

list they would use to delete a report before reading it, or to delete it after skimming or reading it on the screen. We then asked them to indicate which three they used most frequently and which three they considered most important. Similarly, we asked them to mark all the criteria from the relevance list that they would use to save or keep a report, and to indicate the top three in frequency and importance.

While each criteria proved useful to at least some respondents, data analysis showed that some criteria were much more significant than others. Further, when we compared the most popular criteria, that is, those which the largest percentage of respondents reported employing, with those they ranked high in frequency of use and in importance, we found them to be the same. That is, the top criteria measured according to popularity, were also rated by respondents as top in frequency of use and in importance.

Top criteria used to delete a report *before* reading it, and the percent of respondents who reported employing them, were:

Criterion	Frequency
It was not relevant or applicable to the Boeing Company	35%
It included no information that was new to the participant	34%
It was too basic or too general	34%
The participant had no influence over the issues raised	25%

Top criteria used to delete a report *after* skimming it on the screen were different from those participants used to delete before reading:

Criterion	Frequency
It was about a product the participant was not familiar with	52%
It was about a technology not here yet	46%
It was about a specific vendor	46%
It was about something Boeing was already doing	41%

Top criteria for saving a report for further, in-depth reading, or keeping it for future use were the same:

Criterion	Save Frequency	Keep Frequency
It was related directly to the respondent's project	79%	72%
It was about new concepts	73%	40%
It was about industry trends	69%	44%
It kept the respondent up to date	69%	40%
It had hard data	64%	49%
It was relevant to the Boeing Company	61%	40%

Some of the top relevance criteria, as well as those for non-relevance, still relate to topics and subjects; e.g., "it was related to the respondent's project" or "it was not relevant to Boeing." Others bring new attributes of the reports that

determined whether or not a report was relevant to the respondents. All are important to determining relevance, and therefore are promising candidates for filtering.

Performance Evaluation

One of the specific objectives of the study was to uncover the criteria the participants in the observation phase used to evaluate the set of reports they received, and the performance of the filtering mechanism. We were unsuccessful in discovering these criteria because all participants were very satisfied with the service (with or without filtering), and were unable to explain in specific terms what attributes contributed to their satisfaction. In addition, we got the impression that the participants were not concerned with recall at all. Most of them were surprised when we asked them if they thought they missed important reports due to filtering, and they were convinced they received all the reports relevant to them. We decided to focus on precision as a performance measurement because we assumed users asked for information to be filtered for one reason: To reduce the number of non-relevant reports in the sets delivered to them; that is, to increase precision.

To measure performance, therefore, we asked respondents to the questionnaire to estimate the average precision of their monthly delivery of reports. We also asked them if they were satisfied with the precision of the *Reports*'

Results	Person A	Person B	Person C	Person D
Precision (no filter)	13%	8%	2%	5%
Precision (with filter)	18%	20%	16%	15%
Recall (with filter)	13%	36%	50%	88%

These findings showed that precision indeed increased with filtering, for one (Person A) by 40%, and for another (Person C) by 800%. They also showed that even with the increase provided by filtering, precision was rather low.

The findings also revealed a large range of recall ratios, with the lowest at 13%, and the highest at 88%. While such a small sample can suggest no general conclusions, the experiment provided some insight about the reasons for low recall. When the librarian examined the titles which the participants marked relevant, he observed that many were not on topics the participants wanted to include in their profiles. In fact, the lower recall was, the lower the agreement was between the topics of the reports marked relevant and the topics in the participant's profile. This suggested that difficulties in building a useful profile were a significant source for low recall ratios. These difficulties may explain the low precision ratios as well. The results

Relevance Level	With Filter	Without Filter	Significance
Skim on the screen	79%	62%	Yes ($p=.02$)
Save to read	30%	14%	No ($p=.07$)
Keep to use	18%	9%	No ($p=.1$)

The results show that while filtering improves precision, perceived precision ratios can be much improved. Over 20% of filtered reports were deleted based on reading the

delivery. Before administering the questionnaire, we conducted a limited experiment to get preliminary assessment of the actual value added by filtering.

The value added by filtering. A small-scale experiment, which we carried out during the observation phase, tested actual precision and recall values, and the value added by the intermediary librarian who constructed the filtering profile. First we observed the four participants who received the *Reports* without filtering, and marked the reports they considered relevant-- that is, of any interest to them--in that month's delivery. Then the librarian contacted each of them and constructed a profile for each individual through an interview. Next we applied the profiles to that month's delivery of the *Reports*, and printed out the titles of the reports these participants would have received, had they had a filtering profile.

A comparison between the two sets of titles, those received with filtering and those without, and the list of relevant titles, gave us a first look at the added value of filtering. The comparison addressed precision and recall ratios:

also indicate that employing a relevance-feedback mechanism when constructing a profile might be fruitful because at times topics expressed through relevance judgment do not coincide with those the users express verbally, yet both should be considered in the filtering profile.

Perceived levels of precision. In the questionnaire, we first asked users to write the average number of reports they received each month. Starting with the weakest relevance grade, we then asked them: how many they read or skimmed on the screen; how many they saved to read later; and how many they actually kept to use. On all levels of relevance, average precision ratios reported by respondents with filter were higher than those reported by users with no filter but a Oneway ANOVA test with significance level of .05 showed most not significant:

title only, and only 30% of the reports were considered of interest enough to be read thoroughly.

These results represented *perceived* precision ratios, but they were close to observed precision ratios collected during the observation phase of the study where the average precision ratio was 23%. The observed ratio by itself cannot be used as a general precision ratio because it was collected from 14 participants only. In addition, the levels of relevance had not yet been defined when observation took place. A report was considered relevant if a participant thought it was relevant after reading it, decided to save it for in-depth reading or to file it for future use. Therefore, there is no direct equivalent with the ratios of perceived precision. Still, given these limitations, the observed ratios indicated that the perceived ratios are probably only slightly higher than the actual ones.

Level of satisfaction. To assess users' satisfaction with the precision level they received, we asked them to mark in the questionnaire whether they were satisfied with the way the *Reports* were delivered to them, or whether they would like to have them screened according to their needs. We also solicited verbal responses with suggestions for improvements. Most respondents (81%) were satisfied, and most respondents who received the reports unfiltered (87%) were satisfied as well. In addition, most of the respondents who wanted the reports to be screened, already received the *Reports* through a filtering mechanism. This suggested that whether or not a user wanted high precision may not relate to whether or not information was filtered.

It is plausible to assume that respondents who perceived a relatively high precision ratio were more likely to be satisfied with their ratios than those who perceived lower ratios. Results showed that this was not the case. Results pointed to a slight difference in perceived precision ratio, with satisfied respondents' average ratios slightly higher than the ratio perceived by those who believed delivery could have been improved, but a Oneway ANOVA test found these differences not significant.

This suggests that the perceived ratio of precision did not affect respondents' satisfaction with the level of precision they perceived. Most respondents were satisfied with precision ratios, no matter how low they perceived them to be. In addition, the verbal responses suggested that most respondents were not concerned with precision but rather with: (a) improved delivery method; (b) better browsing capabilities, particularly those that would allow them to do some of the filtering by themselves; and (c) better mechanisms to construct their profiles.

The fact that most respondents were satisfied with sets of relatively low precision may explain why none of the participants in the observation phase was concerned about recall. It is possible that with low precision, users were convinced they received all the relevant information. Having to wade through so many reports that were not

relevant to them gave them the impression that their sets were exhaustive and complete.

Discussion

The study described here was conducted at the Boeing Company with engineers and managers receiving the *Gartner Group Reports*, and the results cannot be directly generalized to other settings. Nevertheless, they provide general suggestions about relevance criteria that can be used for filtering.

"Not relevant" is relevant too. Traditionally, information retrieval research has put most emphasis on the relevance of documents, and very little on their other property: non relevance. While some relevance-feedback mechanisms incorporate non-relevance information, this is used to complement relevance assessments, and is rarely expected to introduce new aspects.

Findings from this study show that non-relevance criteria could be as important as the relevance ones. First, some criteria were expressed only as non-relevance criteria. For example, some participants found a report non relevant because: they had no influence over the issues raised in the report; it was about technology that was not there yet; or it raised questions but gave no answers. If limited to positive relevance judgments, these criteria would not have surfaced. None of the participants claimed to be interested in a report because he or she had influence in the matter; because the technology was already there; or because it gave answers, rather than just asking questions. Thus, while these are not likely to be reasons for someone being interested in a report, their opposites are reasons for rejecting a report.

Second, juxtaposing non-relevance criteria against relevance ones may provide more insight about relevance criteria because some were common to the relevance and non relevance lists, and others were the opposites (Fidel & Crandall, 1997). For example, some participants wanted to read reports because they were non-technical, about a specific vendor, or basic and general, while others decided to delete reports for the very same reasons. Unlike other criteria, they have no *absolute* relevance-related value because for some participants they indicated relevance and for others non relevance. Integrating these criteria into a user profile would require two steps. The user first decides whether or not to apply a criterion and then what value to attach to it, relevant or non relevant.

Much stronger indicators of relevance were criteria that appeared in both lists but in opposite directions. These were attributes which caused respondents to conclude that a report was relevant, and the opposite of the same attribute served as a reason for non relevance. For example, respondents might be interested in a report if in

their opinion it was relevant to Boeing, but would delete it if it was not relevant to Boeing. Similarly, they would read a report that presented facts and would delete one without facts, or they were likely to be interested in a report about a product familiar to them but would delete one about an unfamiliar product. These criteria are important indicators of relevance because their presence implies relevance, and their absence non relevance. As such, they can be used in filtering and their application would require only one step. The user decides whether or not to incorporate the criterion in the profile. Once a decision is made, only one value is applicable.

The nature of relevance criteria. Relevance criteria, as well as non-relevance ones, reflected various facets (Fidel & Crandall, 1997). Some were attributes of subject matter. For example, when respondents explained they were interested in a report because it was relevant to the Boeing Company, or because it directly related to their project, they based their observations on the subject matter of the report. Other criteria related to the form of a report, e.g., whether it was a case study, whether or not it included hard data, or how basic or technical it was. The content of a report was another facet. For example, participants noticed when a report described industry trends or gave predictions, when it was about a specific vendor, product or service, or when it included background information or general information. The style of writing played a role as well: Reports that displayed issues in a classified order and clearly were favored, and those who took long to understand were deleted.

While the subject matter, the form of a report, its content and the writing style are stable and somewhat objective facets, some criteria were situational and subjective. Participants rejected reports that dealt with unfamiliar products, and decided to examine those that dealt with familiar products. They considered highly relevant reports about a product or service related directly to their project, and rejected reports dealing with issues about which the respondents' group had already made decisions. Similarly, they rejected reports that included no new information, and considered those with new information relevant, but would at times consider a report relevant if it confirmed or validated what they already knew.

Almost all the criteria that relate to the subject matter, the form and contents of reports, and to their style of writing can be used in filtering to improve precision. In addition, filtering can employ some of the situational criteria that relate to the activities and interests of the individual user at the time filtering is done.

Conclusions

Findings of the study indicate that precision of filtering in the study's setting was relatively low, even though users

were satisfied. The findings also suggest two ways to improve precision.

First, users employ attributes of documents beyond the topics when they decide which documents to delete and which to keep. Moreover, it is not difficult to uncover these attributes for a group of users. More research is needed to develop a general typology of these attributes, and to find ways to integrate them into filtering mechanisms.

Second, building a "good" profile for filtering is still the central obstacle to achieving reasonable performance. Probing into criteria users employ, as was done in this study, can help improve users' profiles. For example, when constructing a profile for a Boeing employee, descriptions of the projects in which the user is engaged, as well as a general statement about Boeing's interests, can support the construction of the user profile. The profile can be updated by periodic perusal of electronic documents in the user's work-related files. Additionally, a training set of documents of various forms, writing styles and type of content can be used when constructing a profile through a relevance-feedback mechanism in which users explain why each document was relevant or not relevant.

Generally speaking, in addition to developing filtering algorithms and agents, research in this area should focus on methods to create and improve filtering profiles, as well as to maintain them over time.

Acknowledgment

The study was supported by a grant from Special Libraries Association.

References

- Anick, P.G., Fly, R.A., & Hansenn, D.R. (1991). Addressing the requirements of a dynamic corporate textual information base. In A. Bookstein, Y. Chiamarella, G. Salton, & V.V. Salton (Eds.), *Proceedings of the Fourteenth Annual International ACM/SIGIR Conference on Research and Development in Information Retrieval* (pp. 163-172). Chicago, IL.
- Barry, C. (1994). User-defined relevance criteria: An exploratory study. *Journal of the American Society for Information Science*, 3, 149-159.
- Belkin, N.J., & Croft, W.B. (1992). Information filtering and information retrieval: Two sides of the same coin? *Communications of the ACM*, 35(12), 29-38.
- Bulkely, W.M. (1989, October 20). Fighting back against data overload. *The Wall Street Journal*, p. B1.
- Dou, H., Hassanaly, P., Quoniam, L., & La Tela, A. (1993). Technology watch and competitive intelligence: a new challenge in education for

- information. *Education for Information*, 11(1), 35-45.
- Fidel, R. & Crandall, M. (1997). The role of subject access in information filtering. In: Cochran, P.A. & Johnson, E. (Eds.) *Visualizing Subject Access for 21st Century Information Resources*. In press.
- Gant, S.P. (1995). A portrait of potential adopters of information filters. In T. Kinney (Ed.) *Proceedings of the 58th Annual Meeting of the American Society for Information Science* (pp. 167-171). Medford, NJ, Information Today.
- Goodman, S.K. (1993). Information needs for management decision-making. *Records Management Quarterly*, 12-23.
- Huber, G.P. (1990). A theory of the effect of advanced information technologies on organizational design, intelligence, and decision making. *Academy of Management Review*, 15(1), 47-71.
- Maes, P. (1994). Agents that reduce work and information overload. *Communications of the ACM*, 37(7), 31-40.
- Marshall, J.G. (1993). Issues in clinical information delivery. *Library Trends*, 42(1), 83-107.
- Maure, J. (Ed.) (1992, December). *Communications of the ACM*, 35(12).
- Schamber, L. (1994). Relevance and information behavior. In M.E. Williams (ED.), *Annual Review of Information Science and Technology*, 29, 3-48.
- Shuldberg, H.K., Macpherson, M., Humphrey, P., & Corely, J. (1993). Distilling information from text: The EDS Template Filler system. *Journal of the American Society for Information Science*, 44(9), 493-507.
- Stadnyk, I., & Kass, R. (1992). Modeling users' interests in information filters. *Communications of the ACM*, 35(12), 49-50.
- Verity. 1997. URL: <http://www.verity.com>
- Yan, T.W., & Garcia-Molina, H. (1994). Index structures for information filtering under the vector space model. In *Proceedings of the 10th International Conference on Data Engineering* (pp. 337-347). Los Alamitos, CA, IEEE.