

User Modeling for Information Access Based on Implicit Feedback

Jinmook Kim^{*}, Douglas W. Oard⁺, and Kathleen Romanik[†]

^{*}College of Information Studies
University of Maryland, College Park, MD 20742-4345
jinmook@glue.umd.edu
phone: (301) 405-2033
fax: (301) 314-9145

⁺Institute for Advanced Computer Studies and
College of Information Studies
University of Maryland, College Park, MD 20742-4345
oard@glue.umd.edu

[†]Powerize.com
901 Elkridge Landing Road, Suite 350
Linthicum, MD 21090
kromanik@powerize.com

Abstract

User modeling can be used in information filtering and retrieval systems to improve the representation of a user's information needs. User models can be constructed by hand, or learned automatically based on feedback provided by the user about the relevance of documents that they have examined. By observing user behavior, it is possible to infer implicit feedback without requiring explicit relevance judgments. Previous studies based on Internet discussion groups (USENET news) have shown reading time to be a useful source of implicit feedback for predicting a user's preferences. The study reported in this paper extends that work by providing a framework for considering alternative sources of implicit feedback, examining whether reading time is useful for predicting a user's preferences for academic and professional journal articles, and exploring whether retention behavior can usefully augment the information that reading time provides. Two user studies were conducted in which undergraduate students examined articles and abstracts related to the telecommunications and pharmaceutical industries. The results showed that reading time could be used to predict the user's assessment of relevance, although reading time for journal articles and technical abstracts are longer than has been reported for USENET news documents. Observation of printing events, a type of retention behavior, was found to provide additional useful evidence about relevance beyond that which could be inferred from reading time. The paper concludes with a brief discussion of the implications of the reported results.

1. Introduction

Internet searchers are faced with the classic needle-in-a-haystack problem, but the haystacks are growing so rapidly that there is continued demand for improved search technology. Information filtering and retrieval systems could provide them with a support for information access. Information retrieval is a "pull" service that users can search for information they need, whereas information filtering is a "push" service that finds new information and presents it to the user (Kim et al., 2000). Content-based filtering systems select documents based on the characteristics of the documents such as the words they contain (Sheth, 1994; Oard, 1997). An alternative, now commonly referred to as recommender systems, is to base the search at least in part on annotations made to the documents by the users (CACM, 1997). A user model that represents some aspect of a user's information needs and/or preferences can be useful in any information access system design, and in the case of information filtering it is clearly a central component. User models can be hand-crafted, but machine learning techniques offer the potential to automatically develop or continuously refine a user model. The usual approach in research systems has been to assemble a set of training instances that have been labeled by the user as relevant (either absolutely or to some degree) or as not relevant. Studies have shown that such explicit feedback from the user is clearly useful (Yan & Garcia-Molina, 1995; Goldberg et al., 1992), but obtaining explicit feedback would likely be problematic in many information access applications. It is well known that users of commercial

information retrieval systems make little use of explicit relevance feedback mechanisms when they are provided, at least in part because providing feedback takes time and may increase the cognitive load on the user. Implicit feedback, in which the system learns by observing the user's behavior, offers an attractive alternative that has received increased attention in recent years (Stevens, 1993; Morita & Shinoda, 1994; Konstan et al., 1997; Nichols, 1997; Oard & Kim, 1998; Kim et al., 2000).

In the next section, we review the state of the art on the use of implicit feedback in information access systems, drawing together what has evolved over time as a diverse set of fields to assemble a coherent picture of the sources of information that can be exploited. We then present the results of a pair of user studies that explore how two such sources, observation of reading time and observations of printing behavior, might be used jointly to build a better user model than could be built using either source alone. The paper concludes with some observations on the limitations of our study, future work that is needed, and the larger implications of work on implicit feedback.

2. A Framework for Implicit Feedback

Implicit feedback may bear only an indirect relationship to the user's assessment of the usefulness of any individual document. But because it can be collected ubiquitously (and thus potentially in great quantities), the potential impact of implicit feedback might ultimately be even greater than that of explicit feedback. InfoScope, a system for filtering Internet discussion groups (USENET), utilized both implicit and explicit feedback for modeling users (Stevens, 1993). Three sources of implicit evidence were used: whether a message was read or ignored, whether it was saved or deleted, and whether or not a follow-up message was posted. In summarizing this groundbreaking study, Stevens observed that implicit feedback was effective for tracking long-term interests because it operates consistently without being intrusive.

Morita and Shinoda (1994) introduced another source, proposing an information filtering technique based on observations of reading time. They conducted a user study over a six-week period with eight users to determine whether preference for Internet discussion group USENET messages was reflected in the time spent reading those messages. The results showed a strong positive correlation between reading time and explicit feedback provided by those users. They also discovered that treating messages that the user read for more than 20 seconds as relevant actually produced better recall and precision in an information filtering simulation than using the messages explicitly rated by the user as relevant would. Konstan et al. (1997) repeated this study in a more natural setting, distributing modified software that allowed volunteer participants in a recommender system trial in which both explicit feedback and reading time were recorded for a small set of USENET discussion groups. Their results indicated that recommendations based on reading time can be nearly as accurate as recommendations based on explicit feedback. They also suggested some additional observable behaviors, including printing, forwarding, and replying privately to a message, as sources for implicit ratings.

Nichols (1997) began the effort to develop a comprehensive view of implicit feedback, with a focus on its use in information filtering systems. He presented a list of potentially observable behaviors; adding purchase, assess, repeated use, refer, mark, glimpse, associate, and query to those mentioned above. Oard & Kim (1998) extended that work, organizing the behavior into three broad categories (examination, retention, reference). They also presented examples from related fields, for example, using Weblink analysis (Brin & Page, 1998) and indexing based on bibliographic citations (Garfield, 1979) to illustrate the potential of implicit feedback based on reference behavior.

Table 1 shows a further refinement of the framework developed in (Oard & Kim, 1998) in which the behaviors are further sorted by the scale of the information objects being manipulated. These segment levels include operations whose natural scale is a portion of a document (e.g., viewing a screen), the object level includes behaviors whose natural scale is an entire document (e.g., purchase), and the collection level includes behaviors whose natural scale includes more than one document (subscription). By "natural scale" we mean the smallest unit normally associated with the behavior—behaviors thus have analogues at larger scales (e.g., viewing an entire document), but not normally at smaller scales (e.g., purchasing a paragraph). The choice of segment, object and collection as labels is intentionally inclusive, since the ideas captured in the table would apply equally well to non-text modalities such as video or music with only minor variations (e.g., listen rather than view). We have also added a fourth major category, annotation, that reflects our realization that the behaviors in that category do not fit cleanly into any of the other categories. Interestingly, when viewed from this perspective, explicit feedback (rating behavior) is merely one type of user behavior that we might observe. This unification is attractive, since it may be beneficial to include

both explicit and implicit feedback in many applications. We categorized our intuition about typical user behavior, and applications (e.g., users might be able to bookmark segments) find this to be a useful framework within which to consider

based on our assignments of behavior to some adjustments may be needed for specific types of documents in meaningful ways). But we have some potential sources for implicit feedback.

Natural Scale

		Segment	Object	Class
Behavior Category	Examine	View	Select	
	Retain		Bookmark Save Purchase Print Delete	Subscribe
	Reference	Quote Cut&Paste	Cite Link Reply Forward	
	Annotate	Annotate	Rate Publish Organize	

Table 1. Potentially observable user behaviors.

3. Experiment Design

As described above, previous studies have found that prediction based on reading time can be about as accurate for USENET as those based on explicit ratings (Morita & Shinoda, 1994; Konstan et al., 1997), and evidence from practice clearly indicates that some types of reference behavior are valuable as well (Brin & Page, 1998; Garfield, 1979). We know little, however, about the utility of many other types of observable behaviors. We thus chose to focus on retention behavior, both because it was easily measured and because our intuitions suggested that users might spend less time reading a document in cases in which they decided to save it for later use. The system that we used was designed to provide access to scientific and professional journal articles (both full text and abstracts), so we were also interested in how reading time and explicit ratings were related in this case. Because we are interested in a broad range of information access applications, we chose to focus on the relationship between observable behavior and explicit ratings rather than some measures such as filtering effectiveness that is tied more closely to a single task.

3.1 Hypotheses

We tested the following hypotheses:

- On average, users spend more time reading relevant full-text journal articles than non-relevant articles.
- On average, users spend more time reading abstracts of relevant journal articles than abstracts of non-relevant articles.
- The combination of reading time and printing behavior will be more useful for predicting explicit ratings than using reading time alone.

3.2 Experimental System

Powerize Server, TM developed by Powerize.com, is a Windows NT text retrieval and filtering system that searches multiple internal and external information sources simultaneously and presents the

retrieved documents to the user in a customized manner that presently uses a manually constructed user model known as profile, she can choose to save the profile and have it re-executed on a regular schedule. Our experiments were done using the Powerize Server 1.0. A custom version of Powerize Server 1.0 was created for our experiments by Powerize.com. It was instrumented to measure reading time and printing behavior and to record user-entered ratings for individual documents.

On the Powerize Server 1.0, users interact with the system through two principal interfaces: Publications and Studio. The Studio interface allows users to select and manage profiles based on their interests of topics and includes five collections of profiles known as "wizard packs:" General, Pharmaceutical, Aerospace, Telecommunications, and Energy. Each wizard pack is designed to serve the needs of a group of users. For example, the Pharmaceutical wizard pack is intended for users in the pharmaceutical industry. The Pharmaceutical and Telecommunications wizard packs were used in our experiments. Each wizard pack consists of several "wizards," and each wizard is designed to help the user complete a particular task. For example, there is a competitive intelligence wizard to help users find information about a competitor. Each wizard is further divided into "topics," which are collections of profile templates designed to retrieve information about a particular subject. For example the competitive intelligence wizard contains topics such as "Mergers and Acquisitions" and "Financial Information." Each profile template encodes the structure of a query for a set of information sources. Users create actual profiles by selecting templates and providing search terms such as a drug or company name. By using templates, users can create powerful queries without being familiar with the individual information sources or their query interfaces. Once, users construct their profiles through the Studio interface, they can browse documents retrieved by the system using the Publications interface.

3.3 Pilot Study

A pilot study was conducted to validate the experimental data collection procedures in order to determine whether information. The pilot study was done using only "Pharmaceutical Wizards," with 4 students who were taking a microbiology course on Drug Action and Design at the University of Maryland. A total of 21 instances of reading time and rating were gathered, which showed the expected pattern of increasing reading time with increasing rating. The data collected from the pilot study also suggested that printing behavior might prove useful. Everyone of the 9 cases in which printing was requested was rated as relevant, and any obvious way of using reading time alone to make predictions would have missed some of those cases.

4. Data Collection

Two experiments were conducted. Eight undergraduate students taking an honors research seminar at the University of Maryland participated in the first experiment. The students were engaged in research for a group project that required examining new products, services, and technologies for wireless Personal Communications Systems (PCS). After conversations with both the students and their instructor to define their information needs, search topics were created by the authors using the "Telecommunications Wizards." A total of 97 full-text articles were retrieved using 5 topics: digital PCS, Iridium, Teledesic, Nextel 1000, and Ricochet. All of these selected information sources were from Dialog, a provider of professional content. The experiment with the Telecommunications user group took place in a single one-hour session. A total of 130 ratings (explicit relevance judgments), with associated reading time and printing behavior observations, were collected. Explicit ratings were collected on a four-point scale: "00" for no interest, "01" for low interest, "02" for moderate interest, and "03" for high interest in both experiments. A rating of "NA" for no comments was also allowed.

The second experiment was done with 85 senior or advanced junior students attending laboratory sessions for a zoology course on Mammalian Physiology at the University of Maryland. Search topics were created by the authors using the "Pharmaceutical Wizards" after interviewing the instructor. A total of 96 articles were returned using 5 topics: beta blocker, antihypertensives, ACE inhibitors, positive inotropic agents, and cardiac sympathomimetics. Again, all of these selected information sources were from Dialog. This experiment was conducted in seven sessions during a single week. Sessions 1 and 2 were


administered following the same procedure that the Telecommunications user group used. There were 18 many simultaneous users on server's hardware assessed to be unreliable measurements of , students were paired in groups of two for sessions 3 ed documents, while the other observed the ed were able to participate in some way, but our ns of a single student. To minimize the potential achine, students were asked not to talk to each ollected during these seven sessions.

communications user group used. There were 18 many simultaneous users on server's hardware assessed to be unreliable measurements of , students were paired in groups of two for sessions 3 ed documents, while the other observed the ed were able to participate in some way, but our ns of a single student. To minimize the potential achine, students were asked not to talk to each ollected during these seven sessions.

5. Data Analysis

A total of 122 cases out of 130 ratings collected from the considered valid for purposes of data analysis. All five cases from the data analysis because that student missed the fi exceeded the Z scores of ± 3 were excluded because they were detected as outliers based on the standardized residuals scores for reading time. One case was excluded 1 shows the descriptive data analysis for the Telecommunications user group. An increase in reading time, tshigher on the scatter plot. The rating of "00," ime, and "02," representing "moderate interests," ere able to identify highly relevant articles ant.

ights subjects in the first experiment were ses collected from one subject were excluded rst half of the experiment. Two other cases that because it had a rating of "no comments." Figure cations user group. An increase in reading time, tshigher on the scatter plot. The rating of "00," ime, and "02," representing "moderate interests," ere able to identify highly relevant articles ant.




Rating	Telecommunications Wizards	
	# of Case	Mean Reading Time
00	20	32.85
01	31	42.84
02	34	57.74
03	37	50.24
Total	122	47.60

Figure 1. Descriptive data analysis for the Telecommunications user group.

In these second experiment, there were 7 sessions. In sessions 1 and 2, 36 subjects provided 166 ratings, but data from those two sessions were not used in this study because of the slow system response time described in the previous section. A total of 532 ratings were gathered from 49 subjects that participated in sessions 3, 4, 5, 6, and 7. A total of 153 cases out of the 532 ratings gathered were considered as valid for data analysis in this study, in part because it was discovered after the experiments ly assembled for presentation to the subjects had e given for the 71 bibliographic citations that lacked did not feel that the bibliographic citations alone Threecases that were detected as outliers and aanalysis. The scatter plot in Figure 2 ted table shows both the number of cases and the mean reading time for each rating.

In these second experiment, there were 7 sessions. In sessions 1 and 2, 36 subjects provided 166 ratings, but data from those two sessions were not used in this study because of the slow system response time described in the previous section. A total of 532 ratings were gathered from 49 subjects that participated in sessions 3, 4, 5, 6, and 7. A total of 153 cases out of the 532 ratings gathered were considered as valid for data analysis in this study, in part because it was discovered after the experiments ly assembled for presentation to the subjects had e given for the 71 bibliographic citations that lacked did not feel that the bibliographic citations alone Threecases that were detected as outliers and aanalysis. The scatter plot in Figure 2 ted table shows both the number of cases and the mean reading time for each rating.



Rating	Pharmaceutical Wizards	
	#ofCase	MeanReadingTime
00	29	42.97
01	51	55.14
02	45	52.00
03	28	51.50
Total	153	51.25

Figure 2. Descriptive data analysis for the Pharmaceutical Wizards user group

5.1 Reading Time as a Source for Implicit Feedback

In both experiments, we noted a decline in mean reading time and those rated as high interest. In fact, a consistent decline in reading time in these second experiment was evident as interest increased. This suggests that we will likely not be able to reliably distinguish between degrees of interest using reading time, so we converted the ratings to a binary scale: “00” to “non-relevant” and “01, 02 and 03” to “relevant” for our subsequent analysis in both experiments.

between articles rated as moderate interest and those rated as high interest. In fact, a consistent decline in reading time in these second experiment was evident as interest increased. This suggests that we will likely not be able to reliably distinguish between degrees of interest using reading time, so we converted the ratings to a binary scale: “00” to “non-relevant” and “01, 02 and 03” to “relevant” for our subsequent analysis in both experiments.

Figure 3 presents the descriptive data analysis on reading time with this binary ratings scale for data collected from the Telecommunications user group. An increase in mean reading time was observed from non-relevant to relevant documents on the graph. Ratings made on non-relevant documents and on relevant documents were normally distributed below and above the mean reading times of 32.85 and 50.49 seconds, respectively

between articles rated as moderate interest and those rated as high interest. In fact, a consistent decline in reading time in these second experiment was evident as interest increased. This suggests that we will likely not be able to reliably distinguish between degrees of interest using reading time, so we converted the ratings to a binary scale: “00” to “non-relevant” and “01, 02 and 03” to “relevant” for our subsequent analysis in both experiments.

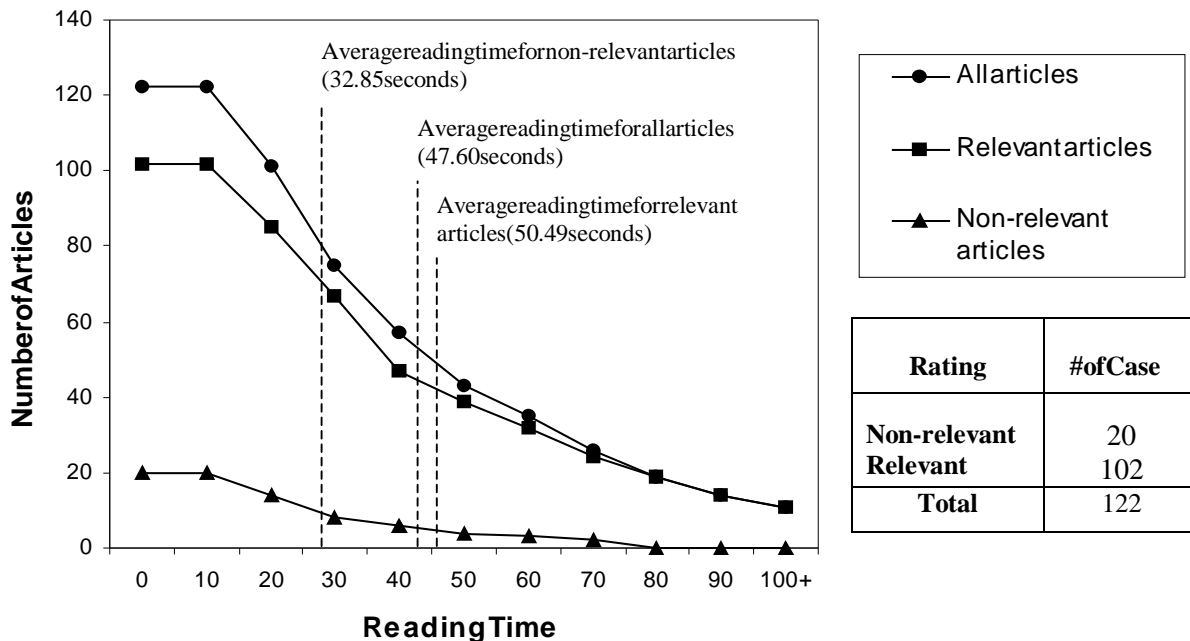


Figure 3. Number of articles read for at least the given duration (Telecommunications user group).

An independent-sample t-test, comparing the mean reading time on relevant documents with non-relevant ones, was done to test our first hypothesis. A statistically significant difference between the two mean reading times was found at $\alpha = .05$. We therefore conclude that users tend to spend a longer time reading relevant articles than non-relevant articles, which is a consistent result with the two previous studies by Morita and Shinoda (1994) and by Konstan et al. (1997). Morita and Shinoda, in their study in 1994, concluded that preference of a user for an article was the dominating factor that affected time spent reading it, and they suggested using a threshold on reading time to detect relevant articles. Their results showed that 30% of interesting articles could be retrieved with precision of 70% by using a threshold of 20 seconds. A much higher threshold would be required in our first experiment to reach a similar recall level. This comports with our intuition, since Morita and Shinoda used USENET messages, while our first experiment was conducted with academic and professional journal articles. Several factors, such as the length of the article, levels of difficulty for understanding the contents, and differences in language skills, could affect the reading time. Subjects in our study might also require longer reading time to understand the content of an article because none of them were experts in the field. Figure 4 shows the recall and precision for different ranges of reading time. For example, the recall and precision that would result from treating articles with reading time of at least 40 seconds as relevant were 0.418 and 0.894, respectively. The horizontal line at a precision of 0.836 shows the value that would be achieved if the user selected articles randomly, since 102 of the 122 articles were judged as relevant.




Figure 4. Precision vs. reading time (Telecommunications user group).

Figure 5 shows the descriptive data analysis for our experiment with the Pharmaceutical user group. There was a 10.22 second difference between the mean reading times on relevant and non-relevant documents, but no statistical significance was found at $\alpha = .05$, based on the independent-sample t-test. The mean reading time on relevant documents was 53.19 seconds, which was close to the one (50.49 seconds) for the Telecommunications user group in our first experiment. The mean reading time on non-relevant documents, however, was 42.97 seconds, which was 10.12 seconds more than was observed with the Telecommunications user group. We suspect that this unexpected outcome resulted at least in part from the different setting in which we paired two students together. As we mentioned in Section 4, one student in each group was observing the session, while the other was browsing retrieved articles. In this case, the student doing the browsing might have sometimes chosen to wait until the other student had also examined the article before clicking on the feedback button. Figure 6 presents the observed recall and precision for different ranges of reading time. Only for extremely long times (over 100 seconds) does reading time provide any clear improvement over random selection (shown by the horizontal line at a precision of 0.810).

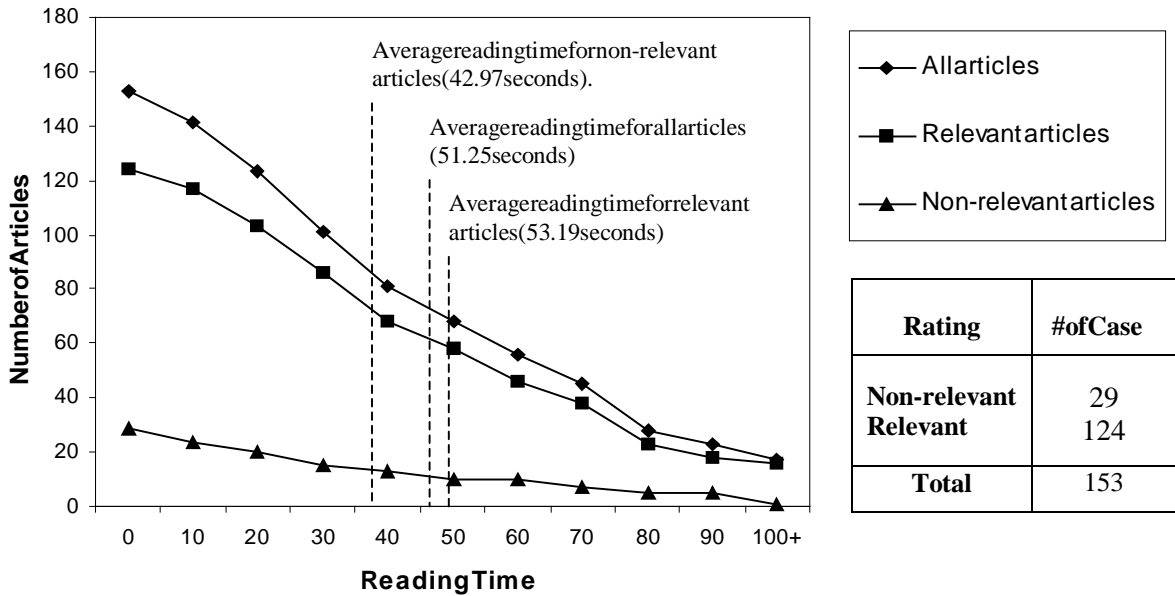


Figure 5. Number of articles read for at least the given duration (Pharmaceutical user group).




Figure 6. Precision vs. reading time (Pharmaceutical user group).

5.2 Printing Behavior as Evidence of Interest

Printing behavior was examined in this study with the hope that predict explicit ratings beyond those clues given by reading documents that could not be discriminated from non-relevant ones using only reading time in Figures 3 and 5. For example, using 47.60 and 51.25 seconds as thresholds for cutting off non-relevant documents in Figures 3 and 5 will also throw 61 out of 102 (59.80%) and 68 out away, respectively. Can printing behavior provide a clue for detecting those relevant documents that would have been thrown away using reading time alone?

It may provide us with clues that can detect those relevant documents that would have been thrown away using reading time alone. There were a number of relevant documents using only reading time in Figures 3 and 5. For example, using 47.60 and 51.25 seconds as thresholds for cutting off non-relevant documents in Figures 3 and 5 will also throw 61 out of 102 (59.80%) and 68 out away, respectively. Can printing behavior provide a clue for detecting those relevant documents that would have been thrown away using reading time alone?

	TelecommunicationsUserGroup		PharmaceuticalUserGroup	
	ReadingTime	Rating	ReadingTime	Rating
	156	03	100	03
	81	02	58	03
			53	03
			43	03
			38	03
			12	03
			100	02
			67	02
			66	02
			48	02
			36	02
			35	02
			32	02
			8	02
			17	01
			11	01
Mean	118.50		45.25	

Table 2. Reading time and ratings for printed articles

Unfortunately, only two cases of printing behavior were available from the data collected from the experiment with the Telecommunications user group, as shown in Table 2. No meaningful interpretation on the data collected could be made with only two cases. We believe that the low frequency of the printing behavior may have resulted from a disparity of goals among the subjects. The members of that undergraduate research team had previously assigned responsibility for technology research to a few of the team members. As a result, the other members of the team may have treated this session more as a familiarization opportunity than as a directed search for information.

There were 16 cases of printing behavior for the experiment with the Pharmaceutical user group. Although no statistical significance was found between the mean reading times for relevant and non-relevant documents with this user group, an increase in reading time from non-relevant to relevant abstracts was observed that could be used as a source for predicting explicit ratings. Using the reading time alone as the source for implicit feedback, however, could not detect those relevant documents that fell under the threshold reading time. Our second goal was to examine how many more relevant documents could be detected by using the printing behavior than using reading time alone.

In Table 2, the mean reading time for 16 cases with printing behavior was 45.25 seconds, which was 2.28 seconds more than the mean reading time for non-relevant documents (42.97 sec.), but 6.01 seconds less than the one for all articles (51.26 sec.). In many cases, articles that were printed were highly relevant, and users seemed to discriminate them quickly from non-relevant ones, which resulted in reducing the reading time. Printing behavior thus provides a useful clue for predicting explicit ratings over reading time, in that it can detect relevant documents below a established threshold of reading time. As in the pilot study, every printed document was judged to be relevant, and 10 out of 16 printed documents had a reading time of less than the mean reading time for all documents (51.26 seconds). Using printing behavior could identify those 10 relevant documents with short reading times.

6. Conclusion

We have shown that reading time can be a useful source of implicit feedback for systems that search academic and professional journal articles in full text, but we were not able to demonstrate a similar effect for abstracts of similar materials. When retention behavior (printing, in this case) was observed, it was found to contribute complementary information, suggesting that systems which couple both types of observations may be able to better model a user's information seeking behavior than those that rely on

reading time alone. Table 1 suggests additional behavior that might be observed, organized in a way that should help system designers recognize useful sources of implicit feedback that would be practical to obtain in their application.

Implicit feedback could be useful in a broad array of information access applications, including filtering or retrieval using content-based and/or annotation-based techniques. Annotation-based techniques stand to benefit in two ways—by using implicit feedback to develop better user models and by sharing with other users the annotations derived from implicit feedback. Annotation-based techniques that can exploit large sets of simple (and noisy) observations could see the greatest impact, perhaps significantly accelerating the deployment of large-scale recommender systems.

Several important research issues remain, however, if we are to fully capitalize on the potential of implicit feedback to support information access. Our approach leverages prior work on information access using explicit feedback by predicting the feedback that a user would have provided. It remains to be seen whether greater effectiveness could be achieved using more closely coupled techniques. The development of explainable systems is another topic that merits increased effort. Ultimately, the systems we build will be tools that the user understands, that may use them to accomplish things that the tools' developers never envisioned. If we are to exploit this potential, we will need to give serious thought to how users will understand what their systems are doing for them so that they can make most of their potential for intentional action. Our work also suggests specific technical questions that now need to be addressed. Perhaps the most urgent is the question of how to accommodate the uncertainty inherent in implicit feedback. We have, for example, shown the precision improvement that can be achieved at various reading time thresholds, but it is not clear that applying a sharp threshold would be the best approach. And if a threshold does turn out to be about as good as any more nuanced strategy, some guidance on how to select that threshold for particular applications will be needed.

Finally, it is important to realize that our work was conducted in a controlled environment. There is now considerable evidence from practice that implicit feedback from situated users is of value, particularly for examination and reference behavior (e.g., double click.com and Google, respectively). The experiments reported in this paper are a first step towards gaining similar experience with retention behavior as well, but evidence based on observations of situated users will be needed before we can fully understand the potential impact of any combination of techniques in a specific application.

Acknowledgements

The authors wish to thank Nick Carmello of Powerize.com for modifying Powerize Server 1.0, Professors William Higgins and Carol Pontzer at the University of Maryland and for working closely with us to find subjects for our experiments and to craft meaningful tasks for them to perform, and our volunteer participants, without whom our research would not have been possible. This work has been supported in part by the Maryland Industrial Partnerships program and Powerize.com.

References

- Brin, S. and Page, L. (1998) The anatomy of a large-scale hypertextual Web search engine. Dept. of Computer Science, Stanford Univ.
- CACM (1997) Special Issue on Recommender Systems, Communications of the ACM, 40(3), March.
- Garfield, E. (1979) *Citation indexing: Its theory and application in science, technology, and humanities*. New York: Wiley-Interscience.
- Goldberg, D., Nichols, D., Oki, B. M., and Terry, D. (1992) Using collaborative filtering to weave an information Tapestry. *Communication of the ACM*, December, 35(12):61-70.
- Kim, J., Oard, D. W., and Romanik, K. (2000) Using implicit feedback for user modeling in Internet and Intranet searching. Technical Report, College of Library and Information Services, University of Maryland at College Park.

- Konstan, J. A., Miller, B. N., Maltz, D., Herlocker, J. L., Gordon, L. R., and Riedl, J. (1997) GroupLens: Applying collaborative filtering to Usenet News. *Communication of the ACM*, March, 40(3), 77-87.
- Morita, M. and Shinoda, Y. (1994) Information filtering based on user behavior analysis and best match text retrieval. *Proceedings of the Seventeenth Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval*, pp. 272-281.
- Nichols, D. M. (1997) Implicit ratings and filtering. In *Proceedings of the 5th DELOS Workshop on Filtering and Collaborative Filtering*, Budapest, Hungary 10-12, ERCIM. <http://www.ercim.org/publication/ws-proceedings/DELOS5/index.html>
- Oard, D. W. (1997) The state of the art in text filtering. *User Modeling and User-Adapted Interaction*, 7(3), 141-178.
- Oard, D. W., and Kim, J. (1998) Implicit Feedback for Recommender System. In *AAAI Workshop on Recommender Systems*, Madison, WI: 81-83. <http://www.glue.umd.edu/~oard/research.html>
- Sheth, B. D. (1994) "A learning approach to personalized information filtering." Master's thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science. <http://lcs.www.media.mit.edu/groups/agents/publications/>
- Stevens, C. (1993) *Knowledge-based assistance for accessing large, poorly structured information spaces*. Ph.D. thesis, University of Colorado, Department of Computer Science, Boulder. <http://www.holodeck.com/curt/mypapers.html>
- Yan, T. W. and Garcia-Molina, H. (1995) SIFT - A tool for wide-area information dissemination. In *Proceedings of the 1995 USENIX Technical Conference*, pp. 177-186. <ftp://db.stanford.edu/pub/yan/1994/sift.ps>.