Attention-Based Information Retrieval

Georg Buscher
German Research Center for Artificial Intelligence (DFKI)
Kaiserslautern, Germany
georg.buscher@dfki.de

ABSTRACT

In the proposed PhD thesis, it will be examined how attention data from the user can be exploited in order to enhance and personalize information retrieval.

Up to now, nearly all implicit feedback sources that are used for information retrieval are based on mouse and keyboard input like clickthrough, scrolling and annotation behavior. In this work, an unobtrusive eye tracker will be used as an attention evidence source being able to precisely detect read or skimmed document passages. This information will be stored in attention-annotated documents (e.g., containing read, skimmed, highlighted, commented passages). Based on such annotated documents, the user's current thematic context will be estimated.

First, this attention-based context will be utilized for attention-based changes concerning the index of vector-space based IR methods. It is intended to regard the contexts as virtual documents, which will be included in the index. In this way, local desktop or enterprise-wide search engines could be enhanced by allowing new types of queries, e.g., "Find a set of documents on my computer concerning the current topic that I have formerly used in a similar context".

Second, the thematic context will be utilized for pre- and postprocessing steps concerning the retrieval process, e.g., for query expansion and result reranking. Therefore, attention-enhanced keyword extraction techniques will be developed that take the degree of attention from the user into account.

Categories and Subject Descriptors

H.3.1 [**Content Analysis and Indexing**]: Indexing methods, H.3.3 [**Information Search and Retrieval**]: Relevance feedback

General Terms

Algorithms, Human Factors

Keywords

Eye tracking, implicit feedback, attention-based index

1. INTRODUCTION

In the last years it could be observed that the human-centered perspective came more and more into the focus of IR research. In this regard, a recent trend is of strong importance: the environment of the user gets more and more personalized; the narrower and broader context of a user is considered in order to better understand the user's needs (e.g., see [2], [3], [21]). This context is increasingly taken into account in information retrieval systems.

However, context is a relatively vague term. For instance, context can be generated implicitly or explicitly, it can be meant in a thematic way (based on content) or rather in an environmental way (e.g., based on application structures, people worked with, etc.), or it can be seen as a short-term or as a long-term context, etc.

One of the current challenges is to elicit the context of a user. This can be done explicitly, for example by asking the user, whether a document is currently relevant or not (i.e., explicit relevance feedback). To use such explicitly generated context is suggestive and yields better results in IR than without considering any user context. However, asking the user about explicit feedback requires a higher effort on the user's side and should therefore be avoided. Thus, implicit feedback recently gained in importance, i.e., observing the user's actions and environment and trying to infer what might be relevant for him.

Up to now, the main sources for implicitly generating thematic context are the user's click-through, scrolling and typing behavior (see [12]); thus, data, which is provided by all normally available input devices that are used to interact with a computer.

A very interesting new evidence source for implicit feedback are the user's eye movements, because mostly they reflect the user's visual attention directly. The eye trackers of today are unobtrusive and are able to identify the user's gaze with high precision (e.g., see [24]). Therefore, applying an eye tracker as a new evidence source for the user's attention introduces a potentially very valuable new dimension of contextual information in information retrieval. It is clear that eye trackers will not be wide spread in the near future due to their expensiveness. However, if becoming less expensive, they might well be interesting for knowledge workers in middle- or large-sized enterprises.

Motivated by these thoughts, in this proposed dissertation, the impact of attention evidence data especially obtained by an eye tracker on information retrieval methods will be examined. Therefore, the focus lies on local desktop and enterprise-wide search. It primarily shall be concentrated on attention-based changes concerning the index of vector-space based IR methods, but also on attention-enhanced pre- and postprocessing steps like query expansion and reranking mechanisms. As an eye tracker is not the only source for attention evidence, a model shall be developed, which integrates different attention evidence sources so that a standardized overall level of attention can be derived for any piece of text in a document.

2. BACKGROUND AND RELATED WORK

Relevance feedback has been considered in information retrieval research for some decades, now. While it has first been focused on explicit relevance feedback, the focus now lies more on implicit feedback due to the lower effort it requires from the user. The

sources for implicit feedback are manifold as examined by [12], for instance.

Nearly all of the implicit feedback sources that are considered in current research are ultimately based on input from the mouse or the keyboard: retention actions like printing, saving and bookmarking; click-through, scrolling, creation and annotation behavior, and viewing time.

As examined by [10], retention actions are "not always a good indicator of document preference". Furthermore, a significant person-dependent difference in the viewing times of documents has been ascertained, so that the viewing time alone does not reflect relevance very well (see also [11]).

However, for the combination of viewing time and scrolling behavior, a strong correlation to explicit relevance feedback has been found by [6]. Also, as examined by [7] the combination of viewing time, clickthrough data and the exiting behavior for a web page has been found to correlate with explicit relevance feedback. Furthermore, in [9] clickthrough data as the only source for implicit feedback in WWW search has been studied and has been found to be rather difficult to interpret concerning absolute relevance.

Up to now, there has been very few work, which determined the read passages of a text with an eye tracker, and which used this information as explicit relevance feedback in information retrieval, e.g., for query expansion.

Concerning the eye movement during reading behavior, there have been many detailed studies in the fields of cognition and psychology (see [17] for a broad overview). Based on the most important results, a reading detection algorithm has been developed in [5], which interprets the eye movements recorded by an eye tracker on the fly to determine, whether the user is currently reading something or not. The passages determined as read are defined as relevant and used for proactive information retrieval [14], [13]. Going a step further, in a feasibility study [18], discriminative hidden Markov models were used to directly infer relevance out of eye movements. This implicit relevance feedback was used for proactive information retrieval and showed to yield promising results [19].

In addition to the above mentioned sources for implicit relevance feedback, user actions like highlighting or commenting of text (e.g., in pdf documents) can be very informative (e.g., compare [23]). Thus, on the whole, quite a lot sources for implicit relevance feedback have been studied.

It is obvious that inferring the user's needs from one feedback source alone might not be very accurate. Therefore, approaches that integrate several sources of implicit feedback for a more accurate user model are interesting. An example for such a system is SUITOR [14], which builds an internal user model out of feedback from an eye tracker, web browsing behavior and application focus. The user model consists of a set of keywords that are extracted out of keyboard input, emails and viewed web pages and that are statistically significant for describing the text. This user model is then used for proactive information retrieval.

Other typical applications for implicit feedback from user observation are query expansion and result reranking mechanisms (e.g., [25], [16]).

However, in nearly all of the approaches to handle implicit feedback, it is tried to infer relevance for the user directly. Though, for this proposed dissertation, implicit feedback from user observation shall primarily be used for deriving degrees of attention for parts of documents. Furthermore, implicit feedback shall not only be used for query expansion and reranking mechanisms, but it shall also be included in the index of vector space IR methods, e.g., in order to be able to pose new kinds of search queries.

3. MAIN RESEARCH QUESTIONS

As stated above it will be focused on how attention data from the user can be integrated in vector space information retrieval models in order to enhance the quality of search. Therefore, four main topics are to be focused (compare also Figure 1):

- Using an eye tracker as a new attention-evidence source in order to determine, which parts of long documents the user has focused
- Elicitation and representation of the current thematic context of the user on the basis of user observation
- Context-based modifications of the term-document index in order to integrate attention data
- Context-enhanced pre- and postprocessing like query expansion and reranking.

These four topics will be motivated and discussed in the following.

3.1 An Eye Tracker as Attention Evidence Source

In order to get valuable implicit information about the user's current context, it is necessary to register what the user does with his

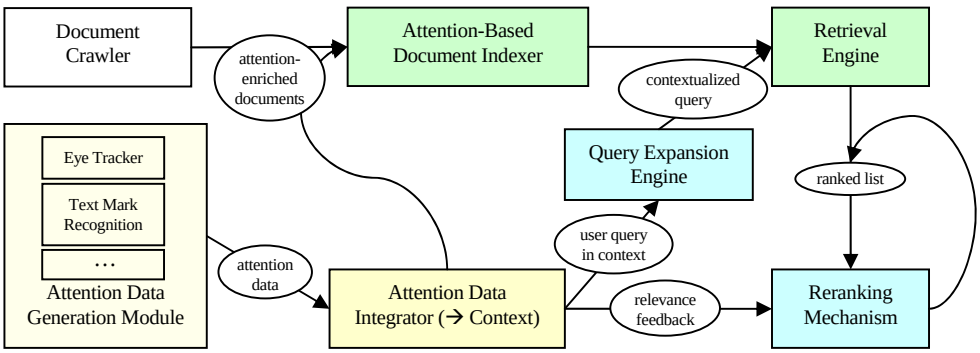

Figure 1: Influence of attention-based user-feedback.

documents as precise as possible. As mentioned in section 1, eye trackers are rarely used, but can provide valuable information about what is interesting for the user. Here, the eye tracker will be used to detect those parts of a document that have been read or skimmed – thus, to which the user has paid attention. By using eye tracking, it can be expected that more precise data about what the user has read can be obtained, e.g., in contrast to merely analyzing the scrolling behavior [8].

During reading or skimming, the eye shows a characteristic behavior composed of fixations, saccades (movements) to the right and regressions (movements to the left). In the last 30 years a lot of research has been done in the fields of cognitive science to study this behavior. As summed up in [17], a fixation is time of about 300ms on average where the eye is steadily gazing at one point. During reading, text can only be perceived up to ca. 9 letters to the right of the fixation point. Furthermore, approximately 10% of the eye moves during reading are regressions. It is important to note, that the eye movement patterns during reading are normally directly associated with cognitive processes.

Considering those eye movement characteristics, an algorithm will be developed, which can detect and differentiate reading and skimming behavior.

An evaluation is planned, which should give information about the degree of correlation of several eye tracking measures (e.g. considering fixation durations and saccade sizes) and explicit relevance feedback from the user for some document: For some query, the user will be asked to read several documents and judge them as either relevant or not relevant. The eye movements during reading will be recorded, and it will be evaluated, if the relevance of a document can be inferred by the user's reading intensity. This question differs from that addressed in [20] in that the relevance of a document is not previously known (judged by others), but defined by the user himself.

What has been read and skimmed will automatically be annotated to the documents, so that this attention-information can be used also at a later time.

3.2 The User's Thematic Context

Identifying the current thematic context of the user is a prerequisite for all further work within this PhD thesis. The thematic context can be elicited by analyzing a variety of different user's actions especially concerning the usage of documents. A document (or parts of it) can be read or skimmed; passages or lines can be commented or highlighted; parts of the document can be created or modified. In contrast to many other works, which directly try to infer relevance out of such implicit feedback, here, it will be focused in the first place on deriving the degree of attention from the user according to implicit feedback: All the interactions with documents imply different degrees of attention from the user. They can be observed in order to get a clue about the user's current thematic context.

As mentioned in section 2 there exist some approaches that try to elicit the thematic context by analyzing the user's browsing behavior, for example. Most of them represent this context by a set of descriptive keywords for the viewed pages [1], [14].

Within this PhD thesis, a technique shall be developed to uniformly integrate several evidence data sources (e.g., read or skimmed passages, comments, …). This multimodality implies that such a thematic context may not only contain a set of descriptive keywords, but rather of weighted keywords giving information about the degree of attention from the user.

A possible solution, which will be studied, is to adapt the Dempster-Shafer theory of evidence [22] in order to combine different attention evidence sources so that a uniform degree of attention can be derived for any text passage of a document.

The Dempster-Shafer theory of evidence is a mathematical theory for the combination of separate information (evidence) sources in order to calculate the probability of an event. It uses belief and plausibility functions to handle uncertainty and vagueness.

The approach for applying this theory to derive a uniform degree of attention for a text passage is to assign belief and plausibility functions to every attention evidence source. For each evidence source, the minimum degree of attention is determined by the belief function, while the maximum possible degree of attention is specified by the plausibility function. In this way, [min, max]-intervals are defined for the evidence sources. For example, if attention is measured on a scale from 0 to 1, skimming generally infers a lower degree of attention than reading. However, there is some uncertainty, which is expressed by the length of the intervals. If there are two (or more) attention evidence sources available for the same text paragraph (e.g., paragraph has been "read" and "commented in margin"), they can be combined according to Dempster-Shafer to calculate a new, more precise interval for the degree of attention.

The difficulty of this approach is that the definitions of the belief and plausibility functions require an explicit statement of the minimum and maximum degree of attention for each evidence source. For example, the maximum degree of attention for a read text passage is clearly 1, but what is the minimum degree: 0.5? The usefulness of such concrete values will be shown by an indirect evaluation, e.g., by attention-based keyword extraction methods that are used for query expansion.

Besides information about the user's attention, it is also important to know, which currently used documents belong to the same thematic context. As the focused target group are knowledge workers, it is very probable that the user examines several similar documents. To automatically detect switches in the thematic context, which could be determined by analyzing changes in the content of the viewed documents (e.g., compare [4]), might be very useful, here.

3.3 Attention-Based Index

Typically, a vector space index is a term-document matrix where the values for the term-document relationships are calculated according to local measures (e.g., term frequency) and global measures (e.g., inverse document frequency; compare [1]). There are mainly two general observations, which motivate attention-based modifications in an index for local desktop search:

Firstly, document structures can be rather artificial for the reader because they are created by one or more authors with specific subjective views on a topic. In a way, the specific composition of a document reflects the mental models of the authors which might be very different from those of the readers. A reader of the document might only regard some parts of it with different intensities, namely those parts, which are of interest to him in his current thematic context (e.g., for his current task). Therefore, a document index for local desktop search, which supports retrieval of already used documents, should consider the user's degree of attention on the different parts of used documents.

Secondly, there usually exists more than one document dealing with a specific topic being of interest to the user. Often, if a knowledge worker is making some document-based inquiry about some topic, he will not only regard one but several documents. To find these used documents at a later time (e.g., if a user remembers something interesting about a topic read some months ago), some connection between them should be maintained in the index.

Based on these thoughts, having properly elicited switches in context and having inferred degrees of attention from the user (see section 3.2), the idea is to index virtual attention-annotated context documents. These context documents consist of those parts of the documents, which belong to one thematic context, and which the user has paid attention to (compare Figure 2). In other words, they are attention-annotated aggregations of attended subdocuments.

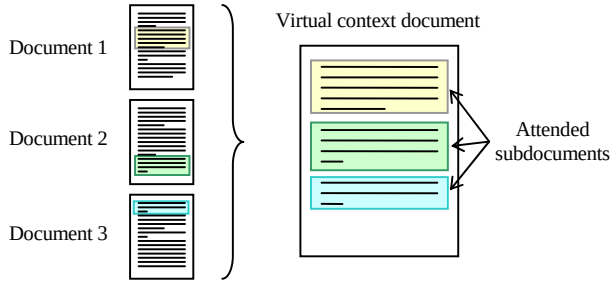

Figure 2: Virtual context documents consist of the attended parts of real documents.

The expectation concerning the virtual context documents is that they represent parts of the user's mental model much better than the single entire documents, out of which they are composed. And having evidence from an eye tracker, which parts of a document have been read or skimmed, the accuracy of the virtual context documents is expected to be very high.

The typical methods for calculating term weights used in the vector space index of documents only consider statistical properties of the documents and document corpus themselves. They do not take into account what the user actually did with the documents. But, having virtual attention-annotated context documents, it will be investigated how the degrees of attention for specific text passages can be incorporated into the term-weighting schema. It can be assumed that such attention-based indexing methods substantially enhance the quality of local search and retrieval processes.

As the attention-annotated context documents are structurally equal to normal documents if the annotations are ignored, normal documents can also be included in the index (e.g., using simple tf-idf measures as if nothing of it has been paid attention to). In this way, a local attention-based index can be initialized with normal documents and can be continuously supplemented by virtual attention-annotated context documents obtained from user observation.

The vision for the retrieval process is that first, the user puts in a query for some topic, second, via the index, some attention-annotated virtual context documents will be found and third, if the user selects a search result entry (i.e., virtual context document), the appropriate real documents will be opened, where those passages are marked that have been paid attention to some time ago.

In order to find the "real" documents that are part of a virtual context document, it may be necessary to establish a second index pointing from virtual context documents to the appropriate real documents. Furthermore, in order to mark the real documents according to the pattern of attention evidence (parts read, skimmed, highlighted, commented, …), some annotation storage has to exist for the documents.

Generally speaking, having an attention-based index which firstly considers the used parts of documents and secondly allows associating several documents regarded in one thematic context, allows new types of queries:

If the attention-based index is only used for local desktop search, queries can be posed like "Find a set of documents on my computer concerning <some topic>, which I have *formerly used* (e.g., read, commented) in a *similar thematic context*".

As a further use-case, an attention-based index could also be stored centrally in an enterprise, so that all documents regarded by colleagues are also indexed. Then an interesting new type of query becomes possible: "Find a set of documents concerning <some topic>, which I have *not regarded*, yet, but which have been *used by some colleague* in a *similar thematic context*." In this scenario, collaborative information retrieval can be achieved without any further data structures and without much further effort.

An evaluation of these ideas for index modifications is rather difficult, because it has to be worked with highly personalized data. For example, it normally depends on the persons' habits, background knowledge and interests, which parts of a document they read with how much intensity. Therefore, it will be tried to exclude as much subjectivity as possible by the design of the experiments. Generally, three types of evaluation studies are envisioned:

- Simulations with manually acquired attention data showing concrete effects of the attention-based weighting schemes.
- Testing against state of the art desktop search engines (with simulated attention data).
- A field test to get an estimation of the subjective benefit for the users.

3.4 Attention-Enhanced Pre- and Postprocessing

The elicited thematic context (see section 3.2) cannot only be used to create virtual context documents that are indexed, but also for query pre- and postprocessing steps like query expansion and result reranking. Here, the main focus lies on how much attention data from an eye tracker can enhance the quality of search results.

As documented in [15], there is only a slight benefit of taking keywords for query expansion that have been directly fixated a particularly long time. For this PhD thesis, the general idea is to extract characteristic keywords out of the currently elicited context document and to use them for query expansion. Therefore, the choice of keywords for query expansion is not directly influenced by measures coming from eye tracking. It rather will be influenced by more abstract measures like the degree of attention for certain read parts of the context document, which itself might not only be dependent on eye tracking data, but also on manual text annotations like highlighting and commenting. As stated previously, by evaluating query expansion in this way, the methods for generating context documents and for calculating concrete values for attention evidence (see section 3.2) can also be evaluated.

As well as for query expansion, the current context document might also be used for result filtering or reranking mechanisms.

4. NEXT STEPS AND ISSUES TO DISCUSS

Concerning section 3.1, I have already implemented an algorithm for the detection of reading and skimming behavior based on eye movement characteristics. Similar to a reading detection algorithm described in [5], it is based on the accumulation of heuristic scores that originate from the detection of reading features (e.g., forward saccade, backward saccade, etc.).

A first case study has been performed, where eleven people had to perform different reading and skimming tasks (e.g., including difficult texts, comprehension questions, time pressure, …). The results so far showed very promising results concerning the very robust detection and proper differentiation of reading and skimming at different levels of concentration.

Now, there are several issues that are important for further ongoing and that I am interested in discussing:

First, I am very much interested in an independent opinion about the main ideas of this proposal (especially creating and indexing the virtual context documents, which could be seen as a view on the user's mental model). Which directions are promising, on which parts should I primarily focus?

Second, concerning the scientific environment, is there important relevant related work that I am not aware of?

Last, but probably most important, suggestions about how evaluation experiments can properly be designed especially for the attention-based index would be very much appreciated. The main problem is that there is intrinsic subjectivity in the highly personalized virtual context documents so that an objective evaluation of retrieval results in terms of recall and precision is difficult.

5. REFERENCES

[1] Baeza-Yates, R.; Ribeiro-Neto, B.: *Modern Information Retrieval*. Addison-Wesley, 1999

[2] Brown, P. J.; Jones, G. J. F.: *Context-aware Retrieval: Exploring a New Environment for Information Retrieval and Information Filtering*. Personal and Ubiquitous Computing, vol. 5(4), pp. 253-263, Springer, 2001

[3] Brusilovsky, P.: *Adaptive Hypermedia*. User Modeling and User-Adapted Interaction, vol. 11, pp. 87-110, Springer, 2001

[4] Bauer, T. L.; Leake, D. B.: *Detecting context-differentiating terms using competitive learning*. ACM SIGIR Forum, vol. 37(2), pp. 4-17, 2003

[5] Campbell, C. S.; Maglio, P. P.: *A robust algorithm for reading detection.* Proceedings of the 2001 workshop on Perceptive user interfaces, pp. 1-7, 2001

[6] Claypool, M.; Le, P.; Waseda, M.; Brown, D.: *Implicit interest indicators.* Proceedings of Intelligent User Interfaces 2001, pp. 33-40, 2001

[7] Fox, S.; Karnawat, K.; Mydland, M.; Dumais, S. T.; White, T.: *Evaluating implicit measures to improve web search.* ACM Transactions of Information Systems (TOIS), vol. 23(2), pp. 147-168, 2005

[8] Hill, W. C.; Hollan, J. D.; Wroblewski D.; McCandless T.: *Edit wear and read wear*. Proc. of SIGCHI 92, pp. 3-9, 1992

[9] Joachims, T.; Granka, L.; Pan, B.; Hembrooke, H.; Gay, G.: *Accurately interpreting clickthrough data as implicit feedback.* Proceedings of ACM SIGIR Conference 2005, pp. 154-161, 2005

[10] Kelly, D.: *Understanding implicit feedback and document preference: A naturalistic user study*. Ph.D. Dissertation, Rutgers University, 2004.

[11] Kelly, D.; Belkin, N. J.: *Display time as implicit feedback: understanding task effects*. Proceedings of ACM SIGIR Conference 2004, pp. 377-384, 2004

[12] Kelly, D.; Teevan, J.: *Implicit Feedback for Inferring User Preference: a Bibliography*. ACM SIGIR Forum, vol. 37(2), pp. 18-28, 2003

[13] Maglio, P. P.; Campbell, C. S.: *Attentive agents.* Communications of the ACM, vol. 46(3), pp. 47-51, 2003

[14] Maglio, P. P.; Barrett, R.; Campbell, C. S.; Selker, T.: *SUITOR: an attentive information system*. Proceedings of the 5th international conference on Intelligent user interfaces, pp. 169-176, 2000

[15] Miller, T.; Agne, S.: *Attention-based information retrieval using eye tracker data*. Proceedings of KCAP05, pp. 209-210, 2005

[16] Radlinski, F.; Joachims, T.: *Query chains: learning to rank from implicit feedback.* Proceedings SIGKDD 2005, pp. 239-248, 2005

[17] Rayner, K.: *Eye Movements in Reading and Information Processing: 20 Years of Research.* Psychological Bulletin, vol. 124(3), pp. 372-422, 1998

[18] Salojärvi, J.; Kojo, I.; Simola, J.; Kaski, S.: *Can relevance be inferred from eye movements in information retrieval?* In Proceedings of WSOM'03, Workshop on Self-Organizing Maps, pp. 261–266. Kyushu Institute of Technology, Kitakyushu, Japan, 2003

[19] Salojärvi, J.; Puolamäki, K.; Kaski, S.: *Relevance feedback from eye movements for proactive information retrieval.* In J. Heikkilä, M. Pietikäinen, and O. Silvén (editors), Workshop on Processing Sensory Information for Proactive Systems (PSIPS 2004), Oulu, Finland, 2004

[20] Salojärvi, J.; Puolamäki, K.; Kaski, S.: *Implicit relevance feedback from eye movements*. In Proceedings of ICANN 2005, pp. 513-518, 2005

[21] Sauermann, L.; Grimnes, G. A.; Kiesel, M.; Fluit, F.; Maus, H.; Heim, D.; Nadeem, D.; Horak, B.; Dengel, A.: *Semantic Desktop 2.0: The Gnowsis Experience*. Proceedings of ISWC 2006, pp. 887-900, 2006

[22] Shafer, G.: *A Mathematical Theory of Evidence*. Princeton, NJ, Princeton, University Press, 1976

[23] Shipman, F.; Price, M. N.; Marshall, C. C.; Golovchinsky, G.: *Identifying Useful Passages in Documents based on Annotation Patterns*. In T. Koch and I. T. Sølvberg (editors), Proceedings of ECDL 2003, pp. 101–112, 2003

[24] Stiefelhagen, R.; Yang, J.; Waibel, A.: *Tracking Eyes and Monitoring Eye Gaze*. Proceedings of the Workshop on Perceptual User Interfaces (PUI'97), pp. 98-100, 1997

[25] White, R. W.; Jose, J. M.; van Rijsbergen, C. J.: *A Simulated Study of Implicit Feedback Models*. In S. McDonald and J. Tait (editors), Proceedings of ECIR 2004, pp. 311–326, 2004

Student's Statement

I am currently at the beginning of my PhD thesis. So far, I studied the relevant literature, made a first concept and developed the main ideas in some detail that should be included in the thesis. Besides, a first step concerning implementation and evaluation has been done by implementing and testing a robust reading and skimming detection algorithm.

I would very much like to attend the doctoral consortium at the SIGIR 2007 conference, because of several reasons.

First, the main ideas of this thesis lie in the basic areas of information retrieval (indexing, keyword extraction, and query expansion mechanisms), so that the SIGIR conference fits best. Furthermore, topics that go in a similar direction have already been presented at previous SIGIR conferences as papers in the doctoral consortium track (e.g., "User-Centered Adaptive Information Retrieval" by X. Shen, SIRIR 2006), but also as full papers (e.g. "Learning User Interaction Models for Predicting Web Search Result Preferences" by Agichtein et al., SIGIR 2006).

Second, I would like to get independent and professional opinions and suggestions of experienced IR researchers. I am especially looking forward to discussing concepts for the evaluation of the different described ideas. As individual meetings with experienced IR researchers are planned, the doctoral consortium seems to be a perfect fit.

Third, discussions with other doctoral students might be of high benefit. As it is likely that some other students intend to do research on similar areas (e.g., IR personalization and contextualization), I look forward to exchange and discuss ideas, which might also help the other students to improve their work.

- Georg Buscher

Supervisor's Statement

Georg started to work on his PhD project last year. He showed a high motivation and made rapidly progress. The topic of his thesis is on improving information retrieval by visually measured attention while reading.

In the last year he has developed a nice method for guessing the intensity of reading by capturing the eye movements by an eye tracker. In the near future, he is going to embed these results into the context of user tasks for recording relevant text passages as user experience which may be employed in similar situations for the same or other users. These efforts should be complemented by new concepts for visually driven information retrieval.

It would be very valuable to discus this very new field with international experts and other PhD students. This opportunity would significantly help him in get other views for improving his Ph.D. thesis. Therefore, I strongly recommend Georg Buscher to attend SIGIR 07 Doctoral Consortium.

- Prof. Andreas Dengel