UCAIR: Capturing and Exploiting Context for Personalized Search

Xuehua Shen, Bin Tan, ChengXiang Zhai
Department of Computer Science
University of Illinois at Urbana-Champaign

ABSTRACT

Personalized search has much to do with capturing and exploiting user-related context information to improve search accuracy. Existing retrieval systems can not support personalized search well for ignoring a user's search context. In this paper, we describe our ongoing work on the User-Centered Adaptive Information Retrieval (UCAIR) project, which aims at capturing and exploiting naturally available user context for personalized search.

1. INTRODUCTION

Precise understanding of a user's information need is essential for achieving optimal retrieval performance. Most existing retrieval systems take a user's query as the sole source of knowledge about the user's information need. However, a query usually only consists of a few short keywords, which are generally insufficient for giving a complete and accurate picture about what the user is really looking for. Thus using more context information about the user and the query is necessary for improving the retrieval performance. Indeed, personalized search essentially boils down to capturing and exploiting related user context information of a query to improve search accuracy.

	Short term (dynamic)	Long term (static)
Implicit	immediately viewed document	past query log
Explicit	judged relevant documents	occupation, hobbies

Table 1: Typology and examples of user context

As shown in Table 1, many kinds of user context information can be potentially exploited [2]. Explicit context consists of information given by a user explicitly, whereas implicit context refers to any context information naturally available while a user interacts with a retrieval system. While explicit context information is more reliable than implicit context, it is often not available to us because it requires extra effort from the user. Implicit context information is thus more interesting to exploit [4, 5, 9].

The goal of the User-Centered Adaptive Information Retrieval (UCAIR) project at the University of Illinois at Urbana-Champaign is to capture and exploit such implicit context, especially short-term context, to optimize retrieval results for a specific user to achieve personalized search[1]. While the project is still in its early stage, we have already achieved some interesting results: (1) We have developed a general decision-theoretic framework for context-sensitive retrieval. (2) We have developed specific retrieval models for exploiting immediate search context based on statistical language model; experiments show that such models achieve better retrieval performance than those not using the context [6]. (3) We have developed a client-side search agent that implements our proposed models and algorithms for personalized web search. A user study shows that the UCAIR search agent performs consistently better than a popular search engine (Google), on which UCAIR search agent is based. (4) We have obtained some experience with evaluating context-sensitive IR. Below we summarize our current work in these directions.

2. A DECISION-THEORETIC FRAMEWORK FOR CONTEXT-SENSITIVE IR

To exploit context for personalized search in a general way, we view the retrieval problem as a decision problem, in which all contextual information and the normally available query and documents should be considered together to optimize the retrieval decision. In general, in response to every user action, the system would choose an optimal system action to take. For example, a user's action may be submitting a query and the system's response may be returning a list of 10 document summaries.

An advantage of treating retrieval generally as a decision-making problem is that we may also treat a user's viewing a document in the search results as a user action, to which the system can respond with updating its own user model about the user's information need. Although, in this case, such a response does not affect the user immediately, we may imagine that after the user views the document and returns to see more search results, the system can choose to rerank any unseen search results based on the updated user model. Indeed, to bring maximum benefit of context to the user, we would like to exploit context as soon as it is available and respond immediately based on any new piece of context information. Such "eager feedback" is precisely what the UCAIR project is aiming at.

We have developed a decision-theoretic framework for optimizing interactive information retrieval based on eager user model updating [7], in which the system responds to every user action by choosing some system action to optimize a utility function. Specifically, as soon as we observe any new piece of evidence from the user, the system would attempt to perform two tasks: (1) compute

[1]UCAIR project web site: http://sifaka.cs.uiuc.edu/ucair/

the current user model to update its belief about the user's information need (2) choose a response that minimizes a loss function. For example, immediately after the user views a document, we could use the knowledge that the viewed document summary is probably relevant to rerank the unseen results so as to minimize a loss function that favors a decision to rank relevant documents above irrelevant ones.

In the traditional retrieval paradigm, the retrieval problem is cast as matching a query with documents and rank documents according to their relevance values. As a result, the whole retrieval process is a simple *independent* cycle of "query submission" and "result display", which is inadequate for exploiting context. The decision-theoretic framework we developed generalizes this traditional retrieval paradigm and allows us to exploit the user's search context in a quite general way.

3. LANGUAGE MODELS FOR CONTEXT-SENSITIVE IR

When instantiating the general decision-theoretic framework described above with specific retrieval methods, we obtain specific retrieval models that can rank documents based on search context. As a case study, we developed several different language models for using implicit feedback information to improve retrieval accuracy in interactive information retrieval [6]. We use the KL-divergence retrieval model [10] as a basis and propose to treat context-sensitive retrieval generally as estimating a query language model based on the current query and any search context information. We proposed and tested several statistical language models to incorporate query and clickthrough history into the KL-divergence model, including linear interpolation with fixed coefficients, Bayesian interpolation, Online Bayesian updating and Batch Bayesian updating. In general, the experiment results show that using implicit feedback information, especially the clickthrough data, can effectively and efficiently improve retrieval performance without requiring additional effort from the user at all [6].

4. A CONTEXT-SENSITIVE IR SYSTEM – UCAIR SEARCH AGENT

We have developed a client-side search agent (called UCAIR) embedded in a web browser which can capture a user's search context and perform implicit feedback [7]. The UCAIR search agent incorporates models and algorithms proposed in section 3 to dynamically rerank the search results to reflect the most updated knowledge of the user's information need whenever any new piece of implicit feedback becomes available.

We chose to do context-sensitive IR at the client side instead of the server side as it has three remarkable advantages. First, the user does not need to worry about privacy infringement, which is a big concern for personalized search [8]. Second, a richer category of user interactions such as mouse movement can be easily captured for implicit feedback. Third, the computation needed for personalization and the storage of the user profile are both done at the client side, so the server is not burdened [3].

We implemented specific techniques to capture and exploit two types of implicit feedback information: (1) identifying any related immediately preceding query and using the query and its corresponding search results to select appropriate terms to expand the current query, and (2) exploiting the viewed document summaries to dynamically rerank any document that has not yet been seen by the user.

User studies show that the UCAIR search agent improves performance over a popular search engine (Google), on which UCAIR search agent is built.

5. EVALUATION OF CONTEXT-SENSITIVE IR

Evaluation of context-sensitive IR poses special challenges due to the difficulty in collecting appropriate user interaction data and cleanly identifying baseline methods. For example, one challenge in evaluating implicit feedback algorithms is that there does not exist any suitable test collection for evaluation. In our study, we used the TREC AP data to create a test collection with implicit feedback information that can be used to quantitatively evaluate implicit feedback algorithms. To the best of our knowledge, this is the first test set for implicit feedback [6].

When evaluating the UCAIR search agent, we conducted a user study involving 6 people. The participants are asked to do a web search on selected query topics from TREC 2004 Terabyte track and TREC 2003 Web track topic distillation task and then make relevance judgments of the search results. By comparing our ranking that incorporates context information and Google's original ranking, we can see whether the use of context information is beneficial. Such a method [7] can be applicable to evaluating similar context-sensitive retrieval systems.

6. FUTURE WORK

The current work can be extended in the following ways. (1) We will further study the retrieval framework for sequential decision making in interactive information retrieval and study how to optimize some of the parameters in the context-sensitive retrieval algorithms. (2) We have only explored some very simple language models for incorporating implicit feedback information. It would be interesting to develop more sophisticated models to better exploit query history and clickthough data. For example, we may treat a clicked document summary differently depending on whether the current query is a generalization or refinement of the previous query. (3) We will study other important user interactions. At the client side, UCAIR search agent will capture and exploit many other user actions such as mouse movement and dwelling time on a document, which may have strong correlation with the document's relevance [1]. (4) Currently, the UCAIR search agent considers the server-side retrieval system as a black box and therefore can not make use of the server's full-text indexing capability. We will study how to make the client-side UCAIR search agent collaborate with the remote retrieval system to provide more powerful contextual search.

7. ACKNOWLEDGEMENT

This work was supported in part by the National Science Foundation grants CAREER-IIS-0347933 and ITR-IIS-0428472. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect those of the National Science Foundation.

8. REFERENCES

[1] M. Claypool, P. Le, M. Waseda, and D. Brown. Implicit interest indicators. In *Proceedings of Intelligent User Interfaces 2001*, pages 33–40, 2001.

[2] P. Ingwersen and N. Belkin. Information retrieval in context – IRiX. *SIGIR Forum*, 38(2), 2004.

[3] G. Jeh and J. Widom. Scaling personalized web search. In *Proceedings of WWW 2003*, 2003.

[4] T. Joachims. Optimizing search engines using clickthrough data. In *Proceedings of SIGKDD 2002*, 2002.

[5] D. Kelly and J. Teevan. Implicit feedback for inferring user preference: A bibliography. *SIGIR Forum*, 2003.

[6] X. Shen, B. Tan, and C. Zhai. Context-sensitive information retrieval using implicit feedback. In *Proceedings of SIGIR 2005*, 2005.

[7] X. Shen, B. Tan, and C. Zhai. Implicit user modeling for personalized search. Technical report, UIUCDCS-R-2005-2575, 2005.

[8] E. Volokh. Personalization and privacy. *Communications of the ACM*, 43(8):84–88, 2000.

[9] R. W. White, J. M. Jose, C. J. van Rijsbergen, and I. Ruthven. A simulated study of implicit feedback models. In *Proceedings of ECIR 2004*, pages 311–326, 2004.

[10] C. Zhai and J. Lafferty. Model-based feedback in KL divergence retrieval model. In *Proceedings of the CIKM 2001*, 2001.