

# Personalization via Collaboration in Web Retrieval Systems: A Context based Approach

**Ayse Goker**

School of Computing, The Robert Gordon University,  
St Andrew Street, Aberdeen AB25 1HG Scotland Email: asg@comp.rgu.ac.uk.

**Daqing He**

Institute for Advanced Computer Studies, The University of Maryland,  
College Park, MD 20742 USA Email: daqing@umiacs.umd.edu.

**World Wide Web is a source of information, and searches on the Web can be analyzed to detect patterns in Web users' search behaviors and information needs to effectively handle the users' subsequent needs. The rationale is that the information need of a user at a particular time point occurs in a particular context, and queries are derived from that need. In this paper, we discuss an extension of our personalization approach that was originally developed for a traditional bibliographic retrieval system but has been adapted and extended with a collaborative model for the Web retrieval environment. We start with a brief introduction of our personalization approach in a traditional information retrieval system. Then, based on the differences in the nature of documents, users and search tasks between traditional and Web retrieval environments, we describe our extensions of integrating collaboration in personalization in the Web retrieval environment. The architecture for the extension integrates machine learning techniques for the purpose of better modeling users' search tasks. Finally, a user-oriented evaluation of Web-based adaptive retrieval systems is presented as an important aspect of the overall strategy for personalization.**

## Introduction

The purpose of information retrieval is to address the information need of a user, at a particular point in time. Because of the difference in information needs between different users, researchers have stated the necessity for information retrieval systems to be adaptive to suit a particular user (Goker & McCluskey, 1991; Matwin & Kubat, 1996). At present, with the increasing amount of

information available and users accessing information via information retrieval systems, it is even more imperative for these systems to improve the search process. This is evident in the Web search environment and will increasingly become a challenge in the ubiquitous computing environment. Retrieval systems need to do this so that users can satisfy their information needs within a reasonable amount of time and effort. One way of achieving this is through the personalization of an information retrieval system enabling it to adapt to individual user's information needs.

Recommendation based on collaboration has been an active research topic for many years. Many systems have been designed to incorporate preference of peers in helping users accessing information (Resnick & Varian, 1997). In this paper, we will present our approach of personalization via collaboration among relevant search contexts in Web retrieval systems. The original idea was first introduced in a user-adaptive information retrieval system (IRS) within a traditional search environment. An adaptive component for a traditional IRS, namely the *user context learner*, was designed and tested to track and apply user specific contexts in subsequent searches. The testing results show the effectiveness of the approach. We then discuss the extension of the original approach with collaboration among relevant search contexts in the Web environment. It includes the description of the relevant architecture in order to enable the application of such user adaptive techniques for the Web. An evaluation methodology for the Web retrieval situation is also proposed.

## Personalization in the Traditional Search Environment

Our idea of context learning based personalization was originated in studying a bibliographic information retrieval system. The IRS was the *Okapi* system, which is based on the probabilistic model of information retrieval

(Robertson & Sparck Jones, 76). Over the years, it has been periodically evaluated via regular experiments such as participation in the TREC (Text REtrieval Conference) program (Robertson *et. al.*, 00). The bibliographic databases comprised of INSPEC, LISA (Library and Information Science database) and the City University Library Catalogue. Users were able to use the Okapi system through a login procedure. Extensive logs comprising of user keystrokes and system responses were kept. For example, it was possible to know content and times of query modification, as well as relevance feedback.

Our user adaptive technique on a traditional IRS is based on the notion of *context*. We believe that an information need of a user at a particular point in time occurs in a particular context, and queries are the results of that need. Goker and McCluskey (1991) identified that often a number of information needs of a user have a common context. We analyzed the search logs of 18 frequent users, and interviewed thirteen of these users subsequently. Our results showed that most users did not appear to have more than 2-3 subject areas when performing their queries. Also, another experiment involving 300 online sessions with 544 queries showed that even when users were specifying that they were starting a new search, often it was a continuation of the preceding one (Walker & Hancock-Beaulieu, 1991).

For the purposes of retrieval, it is important to make some interpretation of context in relation to queries and documents. Via the same context, one query can result in the retrieval of one or more relevant documents and a particular document chosen relevant for one query can also be relevant to others. Thus, it is reasonable to extract terms from one set of relevant documents, viewing them as a representation of the context directly related to these documents, and apply them to retrieve another set of documents that are relevant to subsequent queries. Hence, what may be learnt from the user-IRS interaction for one information need may be of use for further needs. This is our approach of deriving context information for personalization in retrievals.

User context information can theoretically encompass broader aspects of their searches. Cool and Spink (2002) gave a very good discussion on this topic. Here, however, we only use *term* originated from users' queries and any positive relevance feedback provided as an approximation to the context of the information need. The context learner outputs a set of terms that help to form the context representation. They are stored and used in the subsequent queries issued by the user. An important distinction here is that the context derived from one query would be useful for the subsequent queries—irrespective of whether or not they are all in the same session. This is different to traditional relevance feedback, which has a more limited scope pertaining to a single query, i.e., much

within the same session. The approach here is based on what we refer to as *cross-session* learning (learning over sessions).

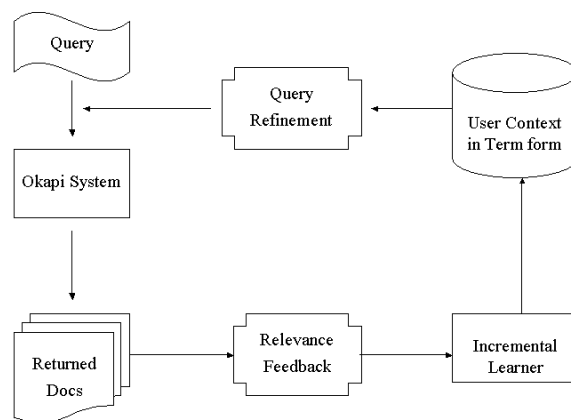


Figure 1: The architecture for applying user adaptive learning in traditional online document search

#### Architecture and Functionality of the Context Learner in Okapi

Our adaptive component is called *User Context Learner* (CL) (Goker, 1999)<sup>1</sup>, which was embedded in an earlier version of Okapi (Walker & Hancock-Beaulieu, 1991).

Figure 1 shows how the CL was added to the Okapi bibliographic retrieval system. The process can be described as follows:<sup>2</sup>

1. A user has an information need which prompts him/her to form a query (and enter a search statement) to the system. The system would then return a ranked list of retrieved documents according to their likelihood of being relevant. The list of documents displayed contains brief record details such as author, part of title and date of publication.
2. When viewing the list of retrieved documents, the user has the option of looking at any of the documents in further detail. For example, the user can view full title, abstract, subject heading and descriptor information for any specified document. If he/she chooses to view the further details for a document, the system requires him/her to provide

<sup>1</sup> This part of the work was performed by Ayse Goker at Centre for Interactive Systems Research, The City University, London.

<sup>2</sup> Further discussions on the CL can be found in (Goker,1997; Goker, 1999).

*relevance feedback* before he/she can resume browsing other documents in the list. Typically, the user is asked the question “*Is this the sort of reference you are looking for?*”—to which the user must answer yes or no.

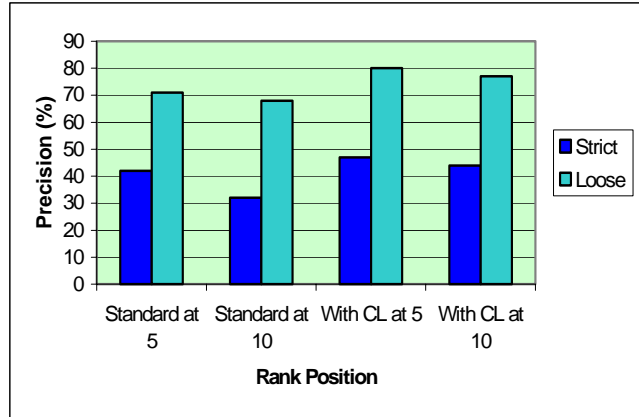
3. The relevance judgment information is used as input for the CL. In general terms, the CL consists of the following several stages/component:
  - ❑ identifying the terms from the query and relevant documents,
  - ❑ scoring the terms according to weights assigned by the probabilistic model and the incorporation of heuristics,
  - ❑ scoring the old context terms according to relevance to the current search, and merging the two sets of terms.
4. The CL will output a 'bag of words' describing the context of the user's information need. The user's context in term form is then used for query refinement. There are several ways the context terms can be used in query refinement. Firstly, they could be used to expand the query with the new set of (context) terms thereby retrieving a different set of references to the system's default. Secondly, they could be used to reorder the system's default output to adapt to that particular user. Thirdly, they could be used to solely break any ties in document scores — effectively reordering with subsets of the default output. The third method was implemented for the least interference with the user's query.<sup>3</sup>

#### Example of Context Learning

This example is based on a real user search episode. The very first query that this particular user inputted was “*active vision*”. The user scrolled through a few screens of returned documents, went backwards and forwards, and chose to view a few “full” documents in the process of judging relevance. In total three documents that were “fully” viewed were marked as relevant. This formed the basis for the input to the CL. The terms ‘learnt’, based on extracting from the pool of terms in the relevant document set, were: *motion, track visual, mobile, perception, egomotion, computer, vision, active, navigation*. After some time, the user made the second query “*sandini*” (presumably an author name). Thus, the terms just identified for the context were used to help improve the ranking for that particular individual (by breaking the ties in document scores with a representation

<sup>3</sup> Subsequent work on the probabilistic model indirectly addressed the problem of having to break ties. However, the use of Context provides an alternative solution to the same problem.

of the context of the need as input). Subsequently, from the list of documents presented, the user marked four as being relevant. Taking into account this new search together with the history of the previous episode the CL updated its list of context terms which then resulted in the



modified list of documents and so on.

Figure 2: The evaluation results of the Okapi user context learner

#### Experiments

We designed the experiment to be as natural for users to search as possible. At various points of the semesters, there were users (including faculty members and students) enrolled to participate experimenting the Okapi system, from which each user was given a unique user ID. In the subsequent months, we observed the users’ search behaviors, and also identified some of them as the potential users for our experiment. After several months when the users had done adequate searches, we started our experiment. Immediately after a user’s search session, we collected the data for the CL to generate a set of documents. We then mixed the output from the CL with that of the baseline system to produce a list of documents, found the user as soon as possible, and asked the user to judge the relevance of each documents on the list. With this relevance information, we then calculated the precision values for the CL and the baseline system. Via this way, we achieved obtaining relevance judgments from users, and at the same time not affect users’ natural behavior in their searches.

Due to different strategies of acquiring, scoring, and merging terms from queries and relevant documents as the representation of the contexts, we explored several variations of the CL in the experiments. In total, we conducted two rounds of experiments that involved 20 users. The first round of experiment, which involved 11 users, 108 queries and 63 sessions, was preliminary and

aimed at exploring different parameters of constructing a CL. Based on the results of the first round experiment, we selected four different versions of CLs, and conducted the second round experiment. The aim of this round was to compare the four CLs against the baseline system - the plain Okapi system. This round involved 9 users, 102 queries and 57 sessions. Figure 2 shows the comparison between precision values of the best performed CL against the baseline plain Okapi system in the second round. Here a three-way judgment (*relevant*, *partially-relevant*, and *not relevant*) was used to overcome users' difficulties in providing a binary judgment for document relevance. The figure shows two types of precision values. One refers to the precision value when only the documents judged as *relevant* are considered. In the graph, "Strict" indicates this case where we call it "strict relevance judgment". The other precision value refers to the case when documents assessed as either *relevant* or *partially-relevant* are included. In the graph, this is marked as "Loose", which stands for "Loose relevance judgment". We also considered two cut-off points (rank positions 5 and 10) for calculating precision values for both strict and loose relevance judgments to obtain more evaluation probes.

On average, the best CL gave a 10% increase in the precision values at both cut-off points positions. It also did uniformly better than the baseline system in both strict and loose relevance judgment conditions. In short, our approach did appear to improve the retrieval accuracy of the IRS. This indicates that it is worth pursuing context learning on a batch of consecutive queries for a user.

### *Lessons Learnt from Personalization in a Traditional IRS*

Many lessons can be drawn from the above experiments. However, here the discussion will focus on those that are related to the challenges in developing effective information access/retrieval systems.

After exploring a variety of context learning algorithms, it appears that efforts to fine-tune the context learner to a greater granularity do not necessarily pay off. Generally, the simpler the context approximation, the more likely it will be useful in document ordering.

The way context terms are merged (along with any heuristics or thresholds) with new contexts appear to be critical—even with the most modest form of personalizing the retrieved set of documents by reordering subsets within it.

More detailed investigation, involving specific questions to users about some of the poorly performing

terms, might provide fuller explanations to the questions like why certain CL variations did not work well.

There is a difference in experiment methodology between a batch learning procedure and a live learning one. In batch mode, there is the benefit of knowing when the search ends and therefore when context learning should take place. In live mode, the system has to rely on the user for some indication as to the end of the query such as the action of performing a new query, editing the old one or exiting the system.

In order to ensure high and accurate user participation, it is advisable to monitor the system regularly and keep track of frequent users' searches in order to target the user for evaluation. When a targeted user performs a search that should be evaluated, he/she should be contacted to obtain the relevance judgments within a short period (such as 20 minutes) after the last search.

### **Towards Personalization via Collaboration in Web Retrieval Systems**

Current searching on the Web is in some ways similar to searching in the traditional IR environment. There is a search engine where users input their queries, the retrieval mechanism then aims at matching queries to the indexed documents or pages, and produces a list of these results. There are also, however, some differences in terms of the documents, the users, their tasks and variety of information needs.

Firstly, the content and the structure of the documents are different. Besides containing hyperlinks, web pages could contain heterogeneous data - which can be unstructured, volatile or redundant, and differ in quality - whereas bibliographic documents in traditional IRSs are usually homogeneous and high quality. Secondly, the volume and extent of distribution of the Web data is much larger than the counter part in traditional IRSs. Thirdly, there are a greater variety of Web users, and the tasks they perform are much more varied in comparison to those in traditional IRSs. For example, users now use the Web not only to find bibliographic information about a particular subject, which are the most performed tasks in traditional (online) information retrieval systems, but also to find out the latest weather, travel information, shopping details, and so on.

Users in traditional IR search usually have 2-3 topics in respect to their search. It is this that motivates us to use information from previous searches to help the current one. Web users might be different in terms of the number of topics, but it is reasonable to assume that the users

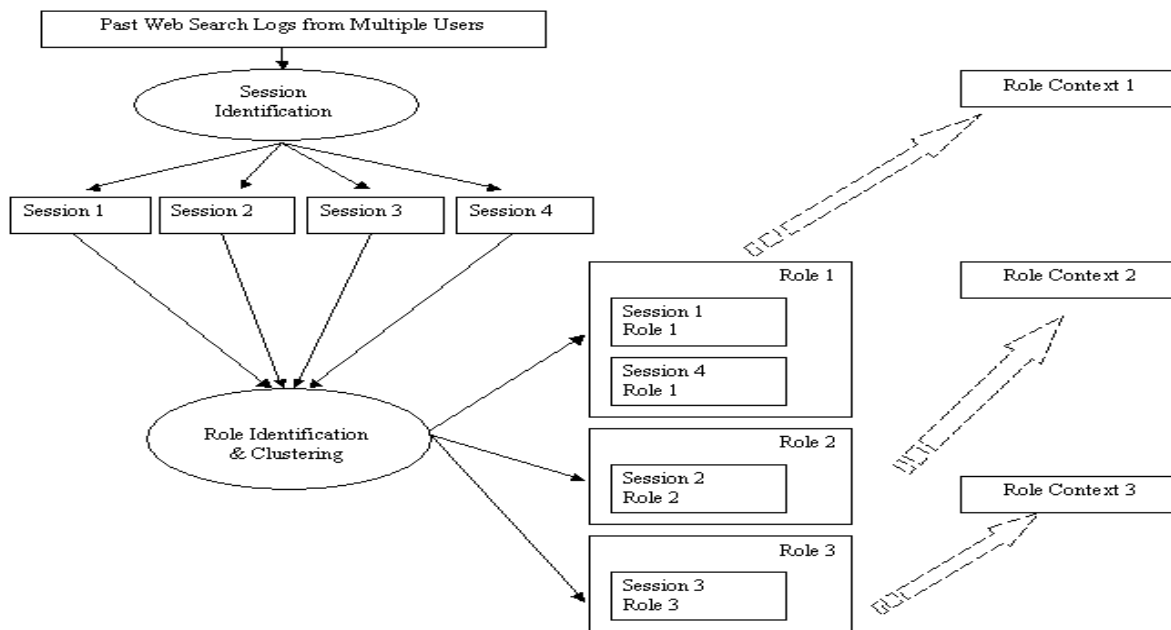


Figure 3: An approach of generating role context information from multiple users' search sessions

would cluster their interests among several topics. Therefore, our approach for the traditional environment is valid in personalization in Web retrieval too. However, due to the difference in aspects of users and retrieval systems between traditional IRs and Web search engines, extensions are essential.

#### *Collaboration among Peers: Role-based Context model for Personalization*

Two assumptions were made in designing the context-based personalization in traditional IRSs. Firstly, we assumed that it was easier to identify a user in a traditional IRS (due to often more explicit logon/off procedures), so the architecture of the personalization relies its learning on the identification of an individual user (see Fig. 1). Secondly, we assumed that the user in such environments tended to have a narrower spectrum of search interests, so there was no need for sub-division of the user's context.

However, the personalization in Web retrieval environments cannot rely on a specific identity of the user for its personalization. This is because, firstly, identifying a particular user in Web searches is difficult for the reasons of a lack of user background information and a reliable mechanism to identify a user. There is not a certain identifier of a user unless he/she has registered, which is seldom the case in most search engines.

Secondly, it is a general consensus that much less information about users' information needs is available in Web retrieval environments than in traditional situations. For example, there are less terms in users' queries and less activities in a search session (Han *et al*, 2001; Jansen *et al*, 1998). In addition, relevance feedback appears to be an additional option to Web retrieval users. There is no explicit "forced" relevance button in search engines for users to click, and the Web searching can be resumed without provision of feedback. Both of these features were incorporated differently in the Okapi IRS. Therefore, there could be much less information about an individual user and his/her searches for learning.

Thirdly, we actually do not want to make the personalization in Web retrieval environments rely solely on a specific identity of users either. There is no collaboration among different users in the personalization architecture in traditional environments, but the Web situation makes collaborative user modeling much more appealing. Comparing to tens of millions of Web users, the number of search topics conducted on the Web are much smaller. There must be large number of people who could share the same interests. Therefore, there may be a lack of information about one user and his/her information need on the Web, but there is useful information from others to be inferred from if a collaboration mechanism can be integrated. Even in the situation where there is enough information for an

individual user, the collaboration could provide much more additional assistance.

We introduced the concept of *context* when we discussed the personalization in traditional IRS. With a more task-oriented view of current Web search processes, we focus on *role contexts*. A user with an interest in a particular topic, when seeking to satisfy that information need, acts in a particular *role*. The role contains the information from the user or other users who had the same interest in the same retrieval task. Therefore, we can say that a role is related to an information need, and it is related to the topic area that covers the need. A role can also be recognised by a set of anticipated actions performed by users in a search task (Coon, 1992; Dix *et al.* 1993). The good aspect of a role is that it can relate to a group of users who share the same role, not only a particular user. In short, queries are an expression of user interests and user interests are related to tasks which can arise because users take on particular roles. A role can comprise of a set of tasks. Hence, a role context is a description of the user's role that can be expressed by queries.

The concept of session defined in (He *et al.*, 2002) helps us to collect the role information. A *session* consists of chronologically ordered queries related to the same topic, and thus these queries are the results of a search related to one role (He & Goker, 2001). Through identifying the boundaries of a session, only the truly related information is cumulated through the user's iteration of queries. In addition, through the role behind the search sessions, relevant information cross sessions can also be cumulated. That is, the role information can be used as a thread to pull relevant information cross sessions together, and to be used in the current session.

Figure 3 shows an example of using our session identification techniques to construct role context information. We started with a Web logs that contains searches issued by many users on the Web. Our session identification techniques can identify a set of sessions, each of which contains a sequence of queries that supposedly share the same role context. We then used role identification and clustering module, which currently is based on term clustering algorithm, but it could be more elaborate process if needed, to group all the sessions which share the same role context together. Based on these clusters, we derived our role context representations. These role representations are stored in the system (e.g., in a place called "role pool" in Fig. 4) for using in later search sessions. Notice that, we collected our data from the searches conducted by many users, and we did not require the identification of users in the process. The only thing we cared was that the information we merged together shares the same role context. Therefore, the introduction of role helps our

personalization approach to avoid closely relying on identifying each individual user, which makes collaboration among different users possible as long as they act in the same role during their searches.

#### *An Architecture and Functionality for Personalization in a Web Retrieval System*

Our approach of personalization in Web IRS depends on role identification to select appropriate relevant information to expand the description of a user's information need. All the roles are stored in a knowledge base called the *role pool*. The initial content of the pool can be obtained through a supervised learning process or by manual construction for small scale applications.

Therefore, the procedure of personalized Web retrieval with this architecture is (see Fig 4):

1. The *Session Identifier* receives the query terms with the time mark and uses them to identify whether or not the current query is in the same session as the previous ones. The detail of the algorithm is presented in (He *et al.*, 02). The role for current session is used for query refinement when this is a continuity of the same session, whereas a new role would be retrieved if the current query is the start of a different session.
2. In the case of a new session, the *Role Identifier*, which is currently based on Bayesian classifier, is used to retrieve the appropriate role from the *Role Pool*. The retrieval process involves partial matching between the input query terms and the existing roles in the pool. However, a default role may be used when the retrieval process cannot find a similar existing role. The retrieved role substitutes the role of previous session.
3. An expanded query is generated by incorporating terms from the role into the original search terms. The added terms would provide more detailed description of the user's information need. This means that more relevant documents could be retrieved or relevant documents could appear in more salient positions.
4. Like traditional IR systems, our system also uses relevance feedback as an essential means to grasp user's evaluation of the returned documents. The feedback can be in explicit form by a user selecting relevant documents or in implicit form by assuming the top returned documents to be relevant. The information collected from relevance feedback is fed to a module called *Role Refinement* to fine tune the content of the role for current session so that the role becomes a more accurate representation to the user's current need. At the end of a session, the corresponding role in the pool is updated.

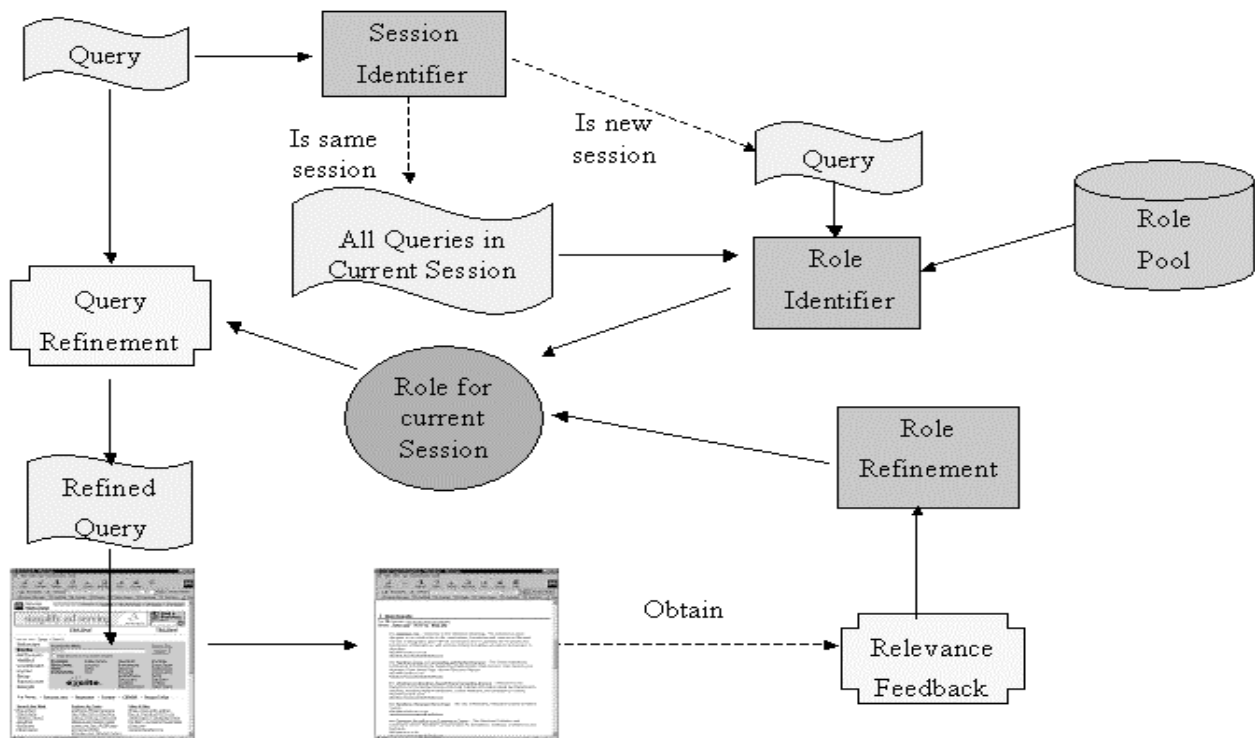


Figure 4: The new architecture for applying user adaptive learning in Web document searches.

## Proposed Evaluation

We consider evaluation to be an integral part of IRS improvement. In particular, where personalization is applied, user-oriented evaluation is the ideal means of assessing the retrieval effectiveness of the system. User-oriented evaluation will help highlight when and why personalization works well.

However, monitoring users and engaging them at various points for evaluation has its problems in the Web environment. These include identification of a specific user, tracking them for evaluation and ensuring that the user performing the evaluation is the same one for which the learning was performed. Thus, it may be necessary to use generic tasks encapsulating information needs that frequently occur amongst Web users and to ask users to perform searches around those needs whilst in the experiment. Below is a breakdown of the process:

- ❑ Prepare the description and scenario for the generic tasks based on profile of real Web users.
- ❑ Identify population (potential pool) of users for the experiment.
- ❑ Perform personalization of Web retrieval system.

- ❑ Obtain user relevance judgments and supplement with exploratory interviews.
- ❑ Compare precision values for strict and loose relevance judgments at different cut-off points (for the plain and personalized version of the system) - as per the evaluation in the traditional IRS. However, it might be appropriate to include a measure reflecting the pertinence or utility of the Web page.

The TREC framework can fit nicely with the above process. The TREC setup, for a decade now, has provided the information retrieval community with an evaluation framework and collection(s) for testing their algorithms and systems. Over the years, several specialized tracks have emerged to enable more focused experiments on particular aspects of retrieval (e.g. cross-lingual-track, filtering-track, interactive-track). The TREC Interactive Track had identified some common topics for Web based searching and defined several tasks for each of these (Hersh & Over, 00).

There were discussions about what kind of collection and topics to be used in this kind of experiment. Whether

searches should be performed on the 'live Web' or on a large but finite extracted portion of the Web? The former has the sense of authentic, but the latter is more manageable and easier to build test topics on. In terms of search topics, if we are too specific in describing a particular search or domain, it may not actually be realistic enough for the user. On the other hand, allowing users to search for any information need would make system comparisons more difficult. As a result, we focused our topic development effort on some major uses of the Web include questions about personal health, buying a given item, and planning travel to a specific place, and aimed for the specification of generic task descriptions around these kinds of information needs for the purposes of experiments. This work, largely based on the presently discontinued Interactive Track, together with the newly formed Hard Track could be used as a basis for the evaluation.

### Related Work

Context has been one of central concerns in studies of information retrieval (IR). With the widely use of information seeking and IR techniques to address new problems in new domains, investigation of context has been greatly expanded. The special issue of Information Processing and Management (Cool & Spink, 2002) represents one of the latest efforts on studying issues of context in IR.

Many previous studies on personalization on the Web have only focused on providing adaptation to link recommendations and browsing agents. Similar to ours, users' previous interests and behavior are treated as important clues in those studies. Syskill and Webert suggest links to a user, based on the user's ratings of Web pages (Pazzani & Billsus, 1997). Edwards and his colleagues have used both page profiles (type of pages the user finds interesting) and link profiles (type of links user explores) to provide interactive assistance during browsing of Web pages (Edwards et al, 1996). Their learning model is similar in structure to that described in this paper for the traditional IRS. Bloedorn et al. (1996) worked on automatically constructing user profiles by employing a generalization hierarchy based on a thesaurus. These profiles are derived from a subject in the hierarchy related to the information need of a user. Mobasher et al (2000) mine usage data for user's profiles to make recommendations during user browsing sessions and propose a corresponding architecture. In contrast, there are studies about Web users' search behavior and successive search phenomena from analyzing Web transaction logs (He & Goker, 2000; Jansen et al, 1998; Spink et al, 1998), but an architecture of personalization in Web retrieval is lacking.

### Conclusion and Further Work

In this paper, we have demonstrated how a user adaptive technique originated from a study in a traditional IRS can be extended to the different situations in Web searches, and described how the adaptive technique can be applied to help Web users in a collaborative way. We proposed an architecture for implementing user role learning on the Web. The main purpose of this framework is to enable learning from similar cluster of queries that are likely to pertain to the same role. This architecture, which enables collaborative learning, will help the IRS avoid the difficulty of identifying a particular user and address the problem of inadequate information to identify the user's need.

Further work on the Web IR system involves refining the Bayesian classifier based role learner to better suit the Web IR environment. Also, as was the case in evaluating the Okapi user context learner, the Web learner would need to be evaluated with user-oriented methods outlined.

Textual information in Web pages is different to the document records in standard bibliographic IRSs. There are differences in structure, content and quality. Likewise, information available via ubiquitous computing environments will result in differences in structure, usage, access methods, content and quality. We are interested in extending our personalization approach into ubiquitous computing environment as well.

### ACKNOWLEDGMENTS

This work is supported in part by UK-EPSC GR/R11742.

### REFERENCES

- Bloedorn, E., Mani, I. and MacMillan, T.(1996). Representational issues in machine learning of user profiles. In AAAI Spring Symposium on Machine Learning in Information Access, pages 25--27.
- Cool, C.,Spink, A. (2002). Issues of context in information retrieval, a special issue of Information Processing and Management, 38(5).
- Coon, D. (1992). Introduction to Psychology. West Publisher.
- Dix, A., Finlay, J., Abowd, G. and Beale, R. (1993). Human-Computer Interaction. Prentice-Hall
- Edwards, P., Bayer, D., Green, C.L. and Payne,T.R. (1996). Experience with learning agents which manage internet-based information. In M.A. Hearst and H. Hirsh, editors, AAAI Spring Symposium on Machine Learning in Information Access, number SS-96-05, pages 31--40. AAAI Press.
- Goker, A. and He, D. (2000). Analysing intranet logs to determine session boundaries for user-oriented learning. In AH2000: Proceedings of the International Conference on Adaptive Hypermedia and Adaptive Web-based Systems, pages 319--322, Trento, Italy. Springer-Verlag.
- Goker, A. and McCluskey, T.L. (1991). Towards an adaptive information retrieval system. In Z. W. Ras and M.

- Zemankova, editors, *Methodologies for Intelligent Systems*. 6th International Symposium, ISMIS'91, pages 348--357. Springer-Verlag.
- Goker, A. (1997). Context Learning in Okapi. *Journal of Documentation*, 53(1):80--83.
- Goker, A. (1999). Capturing information need by learning user context. In A. Rudstrom, editor, *Proceedings of IJCAI99 workshop "Learning about Users"*, pages 21--28.
- Han, S., Goker, A. and He, D. (2001). Web user search pattern analysis for modeling query topic changes. *Proceedings of User Modeling for Context-Aware Applications*, a workshop of the 8th International Conference on User Modeling.
- He, D. and Goker, A. (2000). Detecting session boundaries from Web user logs. In *Proceedings of the 22nd Annual Colloquium on IR Research IRSG 2000*, pages 57--66, Cambridge, UK.
- He, D., Goker, A. and Harper, D. (2002) Combining Evidence for Automatic Web Session Identification. *Information Processing and Management*, 38:727-742.
- Hersh, W. and Over, P. (2000). SIGIR workshop on interactive retrieval at TREC and beyond. *SIGIR Forum*, 34(1), Fall 2000. <http://www.acm.org/sigir/forum/F2000-TOC.html>.
- Jansen, Major B. J., Spink, A., Bateman, J. and Saracevic, T. (1998). Real Life Information Retrieval: A Study of User Queries on the Web. *SIGIR FORUM*, 32(1):5--17.
- Matwin, S. and Kubat, M. (1996). The Role of Context in Concept Learning. In *Proceedings of the ICML-96 workshop on Learning in Context-Sensitive Domains*, pages 1-5, Bari Italy.
- Mobasher, B., Cooley, R., and Srivastava, J. (2000). Automatic Personalization based on Web Usage Mining. *Communications of ACM*, 43(8):142--151.
- Pazzani, M. and Billsus, D. (1997). Learning and revising user profiles: The identification of interesting Web sites. *Machine Learning*, 27:313--331.
- Resnick, P. and Varian, H. (1997). *Communications of the ACM*, special issues on Recommender Systems. 40(3).
- Robertson, S. and Sparck Jones, K. (1976). Relevance weighting of search terms. *Journal of the American Society for Information Sciences*, 27(3):129--146.
- Robertson, S., Walker, S. and Beaulieu, M. (2000). Experimentation as a way of life: Okapi at TREC. *Information Processing and Management*, 36:95--108.
- Spink, A., Wilson, T., Ellis, D. and Ford, N. (1998). Modeling Users' Successive Searches in Digital Environments. *D-Lib Magazine*.
- Walker, S. and Hancock-Beaulieu, M. (1991). Okapi at City: An evaluation facility for interactive IR. Technical Report 6056, British Library R & D. Contributions by Goker A., McCluskey L., Palmer S.