

Web Search Personalization with Ontological User Profiles

Ahu Sieg
asieg@cti.depaul.edu

Bamshad Mobasher
mobasher@cti.depaul.edu

Robin Burke
rburke@cti.depaul.edu

Center for Web Intelligence
School of Computer Science, Telecommunication and Information Systems
DePaul University
243 South Wabash Avenue
Chicago, Illinois, 60604, USA

ABSTRACT

Every user has a distinct background and a specific goal when searching for information on the Web. The goal of Web search personalization is to tailor search results to a particular user based on that user's interests and preferences. Effective personalization of information access involves two important challenges: accurately identifying the user context and organizing the information in such a way that matches the particular context. We present an approach to personalized search that involves building models of user context as ontological profiles by assigning implicitly derived interest scores to existing concepts in a domain ontology. A spreading activation algorithm is used to maintain the interest scores based on the user's ongoing behavior. Our experiments show that re-ranking the search results based on the interest scores and the semantic evidence in an ontological user profile is effective in presenting the most relevant results to the user.

Categories and Subject Descriptors

H.3.3 [Information Systems]: Information Storage and Retrieval—*Information Search and Retrieval*

General Terms

Algorithms, Experimentation

Keywords

Search Personalization, Ontological User Profiles, User Context

1. INTRODUCTION

Web personalization alleviates the burden of information overload by tailoring the information presented based on an individual user's needs. Every user has a specific goal when searching for information through entering keyword

queries into a search engine. Keyword queries are inherently ambiguous but often formulated while the user is engaged in some larger task [15]. For example, a historian may enter the query *Madonna and child* while browsing Web pages about art history, while a music fan may issue the same query to look for updates on the famous pop star.

In recent years, personalized search has attracted interest in the research community as a means to decrease search ambiguity and return results that are more likely to be interesting to a particular user and thus providing more effective and efficient information access [29, 1, 4]. One of the key factors for accurate personalized information access is user context.

Researchers have long been interested in the role of context in a variety of fields including artificial intelligence, context-aware applications, and information retrieval. While there are many factors that may contribute to the delineation of the user context, here we consider three essential elements that collectively play a critical role in personalized Web information access. These three independent but related elements are the user's short-term information need, such as a query or localized context of current activity, semantic knowledge about the domain being investigated, and the user's profile that captures long-term interests. Each of these elements are considered to be critical sources of contextual evidence, a piece of knowledge that supports the disambiguation of the user's context for information access.

In this paper, we present a novel approach for building ontological user profiles by assigning interest scores to existing concepts in a domain ontology. These profiles are maintained and updated as annotated specializations of a pre-existing reference domain ontology. We propose a spreading activation algorithm for maintaining the interest scores in the user profile based on the user's ongoing behavior. Our experimental results show that re-ranking the search results based on the interest scores and the semantic evidence in an ontological user profile successfully provides the user with a personalized view of the search results by bringing results closer to the top when they are most relevant to the user.

The paper is organized as follows. In the next section, we discuss some related background work and our motivation. Our approach for building the ontological user profiles is described in Section 3. In Section 4, we discuss the application of our context model for re-ranking Web search results for *Search Personalization*. The experimental evaluation and results are presented in Section 5. In the final section, we share our conclusions and plan for future work.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CIKM'07, November 6–8, 2007, Lisboa, Portugal.

Copyright 2007 ACM 978-1-59593-803-9/07/0011 ...\$5.00.

2. BACKGROUND AND MOTIVATION

This section includes related work and introduces the terminology we will use throughout the remainder of this paper.

2.1 Related Work

Despite their popularity, users' interactions with Web search engines can be characterized as one size fits all [3]. The representation of user preferences, search context, or the task context is generally non-existent in most search engines [16]. Allan et al. [3] define the problem of *contextual retrieval* as follows: "Combine search technologies and knowledge about query and user context into a single framework in order to provide the most appropriate answer for a user's information needs."

Effective personalization of information access involves two important challenges: accurately identifying the user context and organizing the information in such a way that matches the particular context. Since the acquisition of user interests and preferences is an essential element in identifying the user context, most personalized search systems employ a *user modeling* component.

Recent studies show that users often start browsing through pages that are returned by less precise queries which are cognitively easy to construct. Since the users are reluctant to specify their underlying intent and search goals, personalization must pursue techniques that leverage implicit information about the user's interests [27, 33].

*Google Personalized Search*¹ builds a user profile by means of implicit feedback where the system adapts the results according to the search history of the user. Many systems employ search personalization on the client-side by re-ranking documents that are suggested by an external search engine [30, 18] such as Google. Since the analysis of the pages in the result list is a time consuming process, these systems often take into account only the top ranked results. Also, only the snippets associated with each page in the search results is considered as opposed to the entire page content.

One increasingly popular method to mediate information access is through the use of ontologies [12]. Researchers have attempted to utilize ontologies for improving navigation effectiveness as well as personalized Web search and browsing, specifically when combined with the notion of automatically generating semantically enriched ontology-based user profiles [34]. Our research follows recent ontology-based personalized search approaches [10, 22] in utilizing the *Open Directory Project (ODP)*² taxonomy as the Web topic ontology.

The ODP is the largest and most comprehensive Web directory, which is maintained by a global community of volunteer editors. The ODP taxonomy is used as the basis for various research projects in the area of Web personalization [7, 35].

Liu et al. [17] utilize the first three levels of the ODP for learning profiles as bags of words associated with each category. The user's query is mapped into a small set of categories as a means to disambiguate the words in the query. The Web search is then conducted based on the user's original query and the set of categories. As opposed to using a set of categories, Chirita et al. [6] utilize the documents stored locally on a desktop PC for personalized query expansion.

¹<http://www.google.com/psearch>

²<http://www.dmoz.org>

The query terms are selected for Web search by adapting summarization and natural language processing techniques to extract keywords from locally stored desktop documents.

Hyperlink-based approaches have also been explored as a means to personalize Web search. In Persona [32], the well-known *Hyperlink Induced Topic Selection (HITS)* algorithm [5] is enhanced with an interactive query scheme utilizing the Web taxonomy provided by the ODP to resolve the meaning of a user query.

Considerable amount of Web personalization research aims at enhancing the original PageRank algorithm. In *Personalized Page Rank* [14], a set of personalized hub pages with high PageRank is needed to drive the personalized rank values. In order to automate the hub selection in *Personalized Page Rank*, a set of user collected bookmarks is utilized in a ranking platform called *PROS* [8].

Instead of computing a single global PageRank value for every page, the *Topic-Sensitive PageRank* [13] approach tailors the PageRank values based on the 16 main topics listed in the Open Directory. Multiple *Topic-Sensitive PageRank* values are computed off-line. Using the similarity of the topics to the query, a linear combination of the topic-sensitive ranks are employed at run-time to determine more accurately which pages are truly the most important with respect to a particular query. This approach is effective only if the search engine can estimate the suitable topic for the query and the user. Thus, Qui and Cho [21] extend the topic-sensitive method to address the problem of automatic identification of user preferences and interests.

2.2 Terminology

The notion of *context* may refer to a diverse range of ideas depending on the nature of the work being performed. Previous work defines context by using a fixed set of attributes such as location, time or identities of nearby individuals or objects, as is commonly done in ubiquitous computing [26]. In this section, we define more precisely what we mean by *context* and other related terminology used in the paper.

Context: The representation of a user's intent for information seeking. We propose to model a user's information access context by seamlessly integrating knowledge from the immediate and past user activity as well as knowledge from a pre-existing ontology as an explicit representation of the domain of interest. In our framework, *context* is implicitly defined through the notion of ontological user profiles, which are updated over time to reflect changes in user interests. This representation distinguishes our approach from previous work which depends on the *context* information to be explicitly defined.

Ontology: An ontology is an explicit specification of concepts and relationships that can exist between them. When the knowledge of a domain is represented in a declarative formalism, the set of objects that can be represented is called the universe of discourse. This set of objects, and the describable relationships among them, are reflected in the representational vocabulary with which a knowledge-based program represents knowledge [11]. The set of relations such as subsumption *is-a* and meronymy *part-of* describe the semantics of the domain. Rather than creating our own ontology, we choose to base our reference ontology on an exist-




Figure 1: Ontological User Profile as the Context Model

ing hierarchical taxonomy; a tree-like structure that organizes Web content into pre-defined topics.

Query: A search query that comprises of one or more keywords.

3. ONTOLOGIES FOR WEB PERSONALIZATION

Our goal is to utilize the user context to personalize search results by re-ranking the results returned from a search engine for a given query. Our unified context model for a user is represented as an instance of a reference domain ontology in which concepts are annotated by *interest scores* derived and updated implicitly based on the user’s information access behavior. We call this representation an *ontological user profile*.

Since semantic knowledge is an essential part of the user context, we use a domain ontology as the fundamental source of semantic knowledge in our framework. An ontological approach to user profiling has proven to be successful in addressing the *cold-start problem* in recommender systems where no initial information is available early on upon which to base recommendations [19]. When initially learning user interests, systems perform poorly until enough information has been collected for user profiling. Using ontologies as the basis of the profile allows the initial user behavior to be matched with existing concepts in the domain ontology and relationships between these concepts.

Trajkova and Gauch [34] calculate the similarity between the Web pages visited by a user and the concepts in a domain ontology. After annotating each concept with a weight based on an accumulated similarity score, a user profile is created consisting of all concepts with non-zero weights.

In our approach, the purpose of using an ontology is to identify topics that might be of interest to a specific Web user. Therefore, we define our ontology as a hierarchy of topics, where the topics are utilized for the classification and categorization of Web pages. The hierarchical relationship among the concepts is taken into consideration for building the ontological user profile as we update the annotations for existing concepts using spreading activation.

3.1 Ontological User Profiles

The Web search personalization aspect of our research is built on the previous work in ARCH [28]. In ARCH, the initial query is modified based on the user’s interaction with a concept hierarchy which captures the domain knowledge. This domain knowledge is utilized to disambiguate the user context.

In the present framework, the *user context* is represented using an *ontological user profile*, which is an annotated in-

stance of a reference ontology. Figure 1 depicts a high-level picture of our proposed context model based on an *ontological user profile*. When disambiguating the context, the domain knowledge inherent in an existing reference ontology is called upon as a source of key domain concepts.

Each ontological user profile is initially an instance of the reference ontology. Each concept in the user profile is annotated with an *interest score* which has an initial value of one. As the user interacts with the system by selecting or viewing new documents, the ontological user profile is updated and the annotations for existing concepts are modified by spreading activation. Thus, the *user context* is maintained and updated incrementally based on user’s ongoing behavior.

Accurate information about the user’s interests must be collected and represented with minimal user intervention. This can be done by passively observing the user’s browsing behavior over time and collecting Web pages in which the user has shown interest. Several factors, including the frequency of visits to a page, the amount of time spent on the page, and other user actions such as bookmarking a page can be used as bases for heuristics to automatically collect these documents [9].

3.2 Representation of Reference Ontology

Our current implementation uses the *Open Directory Project*, which is organized into a hierarchy of topics and Web pages that belong to these topics. We utilize the Web pages as training data for the representation of the concepts in the reference ontology. The textual information that can get extracted from Web pages explain the semantics of the concepts and is learned as we build a term vector representation for the concepts.

We create an aggregate representation of the reference ontology by computing a term vector \vec{n} for each concept n in the concept hierarchy. Each concept vector represents, in aggregate form, all individual training documents indexed under that concept, as well as all of its subconcepts.

We begin by constructing a global dictionary of terms extracted from the training documents indexed under each concept. A stop list is used to remove high frequency, but semantically non-relevant terms from the content. Porter stemming [20] is utilized to reduce words to their stems. Each document d in the training data is represented as a term vector $\vec{d} = \langle w_1, w_2, \dots, w_k \rangle$, where each term weight, w_i , is computed using term frequency and inverse document frequency [25]. Specifically, $w_i = tf_i * \log(N/n_i)$, where tf_i is the frequency of term i in document d , N is the total number of documents in the training set, and n_i is the number of documents that contain term i . We further normalize each document vector, so that \vec{d} represents a term vector with unit length.

The aggregate representation of the concept hierarchy can be described more formally as follows. Let $S(n)$ be the set of subconcepts under concept n as non-leaf nodes. Also, let $\{d_1^n, d_2^n, \dots, d_{k_n}^n\}$ be the individual documents indexed under concept n as leaf nodes. $Docs(n)$, which includes of all of the documents indexed under concept n along with all of the documents indexed under all of the subconcepts of n is defined as:

$$Docs(n) = [\bigcup_{n' \in S(n)} Docs(n')] \cup \{d_1^n, d_2^n, \dots, d_{k_n}^n\}$$

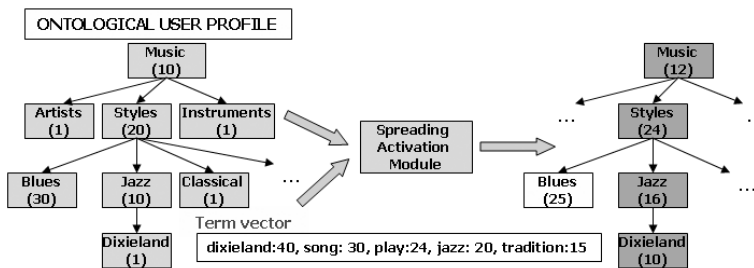


Figure 2: Portion of an Ontological User Profile where Interest Scores are updated based on Spreading Activation

The concept term vector \vec{n} is then computed as:

$$\vec{n} = \left[\sum_{d \in Docs(n)} \vec{d} \right] / |Docs(n)|$$

Thus, \vec{n} represents the centroid of the documents indexed under concept n along with the subconcepts of n . The resulting term vector is normalized into a unit term vector.

3.3 Context Model

Figure 2 depicts a portion an ontological user profile corresponding to the node *Music*. The interest scores for the concepts are updated with spreading activation using an input term vector.

Each node in the ontological user profile is a pair, $\langle C_j, IS(C_j) \rangle$, where C_j is a concept in the reference ontology and $IS(C_j)$ is the interest score annotation for that concept. The input term vector represents the active interaction of the user, such as a query or localized context of current activity.

Based on the user’s information access behavior, let’s assume the user has shown interest in *Dixieland Jazz*. Since the input term vector contains terms that appear in the term vector for the *Dixieland* concept, as a result of spreading activation, the interest scores for the *Dixieland*, *Jazz*, *Styles*, and *Music* concepts get incremented whereas the interest score for *Blues* gets decreased. The *Spreading Activation* algorithm and the process of updating the interest scores are discussed in detail in the next section.

3.4 Learning Profiles by Spreading Activation

We use *Spreading Activation* to incrementally update the *interest score* of the concepts in the user profiles. Therefore, the ontological user profile is treated as the semantic network and the interest scores are updated based on activation values.

Traditionally, the spreading activation methods used in information retrieval are based on the existence of maps specifying the existence of particular relations between terms or concepts [24]. Alani et al. [2] use spreading activation to search ontologies in Ontocopi, which attempts to identify communities of practice in a particular domain. Spreading activation has also been utilized to find related concepts in an ontology given an initial set of concepts and corresponding initial activation values [23].

In our approach, we use a very specific configuration of spreading activation, depicted in Algorithm 1, for the sole purpose of maintaining *interest scores* within a user profile. We assume a model of user behavior can be learned through the passive observation of user’s information access activity

and Web pages in which the user has shown interest in can automatically be collected for user profiling.

The algorithm has an initial set of concepts from the ontological user profile. These concepts are assigned an initial activation value. The main idea is to activate other concepts following a set of weighted relations during propagation and at the end obtain a set of concepts and their respective activations.

As any given concept propagates its activation to its neighbors, the weight of the relation between the origin concept and the destination concept plays an important role in the amount of activation that is passed through the network. Thus, a one-time computation of the weights for the relations in the network is needed. Since the nodes are organized into a concept hierarchy derived from the domain ontology, we compute the weights for the relations between each concept and all of its subconcepts using a measure of containment. The containment weight produces a range of values between zero and one such that a value of zero indicates no overlap between the two nodes whereas a value of one indicates complete overlap.

The weight of the relation w_{is} for concept i and one of its subconcepts s is computed as $w_{is} = \frac{\vec{n}_i \cdot \vec{n}_s}{\vec{n}_i \cdot \vec{n}_i}$, where \vec{n}_i is the term vector for concept i and \vec{n}_s is the term vector for subconcept s . Once the weights are computed, we process the weights again to ensure the total sum of the weights of the relations between a concept and all of its subconcepts equals to 1.

The algorithm considers in turn each of the documents assumed to represent the current context. For each iteration of the algorithm, the initial activation value for each concept in the user profile is reset to zero. We compute a term vector for each document d_i and compare the term vector for d_i with the term vectors for each concept C_j in the user profile using a cosine similarity measure. Those concepts with a similarity score, $sim(d_i, C_j)$, greater than zero are added in a priority queue, which is in a non-increasing order with respect to the concepts’ activation values.

The activation value for concept C_j is assigned to $IS(C_j) * sim(d_i, C_j)$, where $IS(C_j)$ is the existing interest score for the specific concept. The concept with the highest activation value is then removed from the queue and processed. If the current concept passes through restrictions, it propagates its activation to its neighbors. The amount of activation that is propagated to each neighbor is proportional to the weight of the relation. The neighboring concepts which are activated and are not currently in the priority queue are added to queue, which is then reordered.

Input: Ontological user profile with interest scores and a set of documents
Output: Ontological user profile concepts with updated activation values

$CON = \{C_1, \dots, C_n\}$, concepts with interest scores
 $IS(C_j)$, interest score
 $IS(C_j) = 1$, no interest information available
 $I = \{d_1, \dots, d_n\}$, user is interested in these documents

```

foreach  $d_i \in I$  do
  Initialize priorityQueue;
  foreach  $C_j \in CON$  do
     $C_j.Activation = 0$ ; // Reset activation value
  end
  foreach  $C_j \in CON$  do
    Calculate  $sim(d_i, C_j)$ ;
    if  $sim(d_i, C_j) > 0$  then
       $C_j.Activation = IS(C_j) * sim(d_i, C_j)$ ;
      priorityQueue.Add( $C_j$ );
    else
       $C_j.Activation = 0$ ;
    end
  end
  while priorityQueue.Count > 0 do
    Sort priorityQueue; // activation values(descending)
     $C_s = priorityQueue[0]$ ; // first item(spreading concept)
    priorityQueue.Dequeue( $C_s$ ); // remove item
    if passRestrictions( $C_s$ ) then
      linkedConcepts = GetLinkedConcepts( $C_s$ );
      foreach  $C_l$  in linkedConcepts do
         $C_l.Activation += C_s.Activation * C_l.Weight$ ;
        priorityQueue.Add( $C_l$ );
      end
    end
  end
end

```

Algorithm 1: Spreading Activation Algorithm

The process repeats itself until there are no further concepts to be processed in the priority queue. The neighbors for the spreading concept are considered to be the linked concepts. For a given spreading concept, we can ensure the algorithm processes each edge only once by iterating over the linked concepts only one time. The order of the iteration over the linked concepts does not affect the results of activation. The linked concepts that are activated are added to the existing priority queue, which is then sorted with respect to activation values.

Input: Ontological user profile concepts with updated activation values
Output: Ontological user profile concepts with updated interest scores

$CON = \{C_1, \dots, C_n\}$, concepts with interest scores
 $IS(C_j)$, interest score
 $C_j.Activation$, activation value resulting from Spreading Activation
 k , constant
 $n = 0$;

```

foreach  $C_j \in CON$  do
   $IS(C_j) = IS(C_j) + C_j.Activation$ ;
   $n = n + (IS(C_j))^2$ ; // sum of squared interest scores
   $n = \sqrt{n}$ ; // square root of sum of squared interest scores
end
foreach  $C_j \in CON$  do
   $IS(C_j) = (IS(C_j) * k) / n$ ; // normalize to constant length
end

```

Algorithm 2: Algorithm for the Normalization and Updating of Interest Scores in the Ontological User Profile




Figure 3: Personalized Web Search based on Ontological User Profiles

The interest score for each concept in the ontological user profile is then updated using Algorithm 2. First the resulting activation value is added to the existing interest score. The interest scores for all concepts are then treated as a vector, which is normalized to a unit length using a pre-defined constant, k , as the length of the vector. Rather than gradually increasing the interest scores, we utilize normalization so that the interest scores can get decremented as well as getting incremented. The concepts in the ontological user profile are updated with the normalized interest scores.

4. SEARCH PERSONALIZATION

Our goal is to utilize the user context to personalize search results by re-ranking the results returned from a search engine for a given query. Figure 3 displays our approach for search personalization based on ontological user profiles.

Assuming an ontological user profile with interest scores exists and we have a set of search results, Algorithm 3 is utilized to re-rank the search results based on the interest scores and the semantic evidence in the user profile.

Input: Ontological user profile with interest scores and a set of search results
Output: Re-ranked search results

$CON = \{C_1, \dots, C_n\}$, concepts with interest scores
 $IS(C_j)$, interest score
 $R = \{d_1, \dots, d_n\}$, search results from query q

```

foreach  $d_i \in R$  do
  Calculate  $sim(d_i, q)$ ;
   $maxSim = 0$ ;
  foreach  $C_j \in CON$  do
    Calculate  $sim(d_i, C_j)$ ;
    if  $sim(d_i, C_j) \geq maxSim$  then
      (Concept) $c = C_j$ ;
       $maxSim = sim(d_i, C_j)$ ;
    end
  end
  Calculate  $sim(q, c)$ ;
  if  $IS(c) > 1$  then
     $rankScore(d_i) = IS(c) * \alpha * sim(d_i, q) * sim(q, c)$ ;
  else
     $rankScore(d_i) = IS(c) * sim(d_i, q) * sim(q, c)$ ;
  end
end
Sort  $R$  based on rankScore;

```

Algorithm 3: Re-ranking Algorithm

A term vector \vec{r} is computed for each document $r \in R$, where R is the set of search results for a given query. The term weights are obtained using the *tf.idf* formula described earlier. To calculate the rank score for each document, first the similarity of the document and the query is computed using a cosine similarity measure. Then, we compute the

similarity of the document with each concept in the user profile to identify the best matching concept.

Once the best matching concept is identified, a rank score is assigned to the document by multiplying the interest score for the concept, the similarity of the document to the query, and the similarity of the specific concept to the query. If the interest score for the best matching concept is greater than one, it is further boosted by a tuning parameter α . Once all documents have been processed, the search results are sorted in descending order with respect to this new rank score.

5. EXPERIMENTAL EVALUATION

Our experimental evaluation is designed to address three particular questions:

- Do the interest scores for individual concepts in the ontological profile converge?
- Do the changes in interest scores accurately reflect user interest in specific topics?
- Can the semantic evidence provided by the ontological profiles be used to effectively re-rank Web search results to present the user with a personalized view?

Since the queries of average Web users tend to be short and ambiguous [31], our goal is to demonstrate that re-ranking based on ontological user profiles can help in disambiguating the user’s intent particularly when such queries are used.

5.1 Experimental Metrics

For the user profile convergence experiments, we employ two statistical measures; the arithmetic mean (average) and variance. We compute the average interest scores so that we can demonstrate the average rate of increase converges as a result of updating the ontological user profiles over time. Also, we utilize variance in order to measure how the interest scores are spread around the mean as a result of incremental updates. Our results are discussed in Section 5.3.

For the personalized search experiments, we measure the effectiveness of re-ranking in terms of *Top-n Recall* and *Top-n Precision*. For example, at $n = 100$, the top 100 search results are included in the computation of recall and precision, whereas at $n = 90$, only the top 90 results are taken into consideration.

Starting with the top one hundred results and going down to top ten search results, the values for n include $n = \{100, 90, 80, 70, \dots, 10\}$. The *Top-n Recall* is computed by dividing the number of relevant documents that appear within the top n search results at each interval with the total number of relevant documents for the given concept.

$$\text{Top-n Recall} = \frac{\# \text{ of relevant retrieved within } n}{\text{total } \# \text{ of relevant documents}}$$

We also compute the *Top-n Precision* at each interval by dividing the number of relevant documents that appear within the top n results with n .

$$\text{Top-n Precision} = \frac{\# \text{ of relevant retrieved within } n}{n}$$

5.2 Experimental Data Sets

As of December 2006, the *Open Directory* contained more than 590,000 concepts. For experimental purposes, we use a branching factor of four with a depth of six levels in the hierarchy. Our experimental data set contained 563 concepts in the hierarchy and a total of 10,226 documents that were indexed under various concepts.

The indexed documents were pre-processed and divided into three separate sets including a *training set*, a *test set*, and a *profile set*. For all of the data sets, we kept track of which concepts these documents were originally indexed under in the hierarchy. The *training set* was utilized for the representation of the reference ontology, the *profile set* was used for spreading activation, and the *test set* was utilized as the document collection for searching.

The *training set* consisted of 5041 documents which were used for the one-time learning of the reference ontology. The concept terms and corresponding term weights were computed using the formula described in the Representation of Reference Ontology section.

A total of 3067 documents were included in the *test set*, which were used as the document collection for performing our search experiments. Depending on the search query, each document in our collection can be treated as a signal or a noise document. The signal documents are those documents relevant to a particular concept that should be ranked high in the search results for queries related to that concept. The noise documents are those documents that should be ranked low or excluded from the search results.

The *test set* documents that were originally indexed under a specific concept and all of its subconcepts were treated as signal documents for that concept whereas all other test set documents were treated as noise. In order to create an index for the signal and noise documents, a *tf.idf* weight was computed for each term in the document collection using the global dictionary of the reference ontology.

The *profile set* consisted of 2118 documents, which were treated as a representation of specific user interest for a given concept to simulate ontological user profiles. As we performed the automated experiments for each concept/query, only the profile documents that were originally indexed under that specific concept were utilized to build an ontological user profile by updating the interest scores with the spreading activation algorithm.

5.3 Experimental Methodology and Results

In this section, we provide our methodology and results for two independent but related aspects of our experimental evaluation. One aspect is to demonstrate user profile convergence. The second aspect of our evaluation is to design experiments to demonstrate the effectiveness of our approach for search personalization.

5.3.1 User Profile Convergence

With the user profile convergence experiments, our goal is to demonstrate that the rate of increase in interest scores stabilizes over incremental updates. Every time a new Web page, which the user has shown interest in, is processed via spreading activation, the interest scores for the concepts in the ontological user profile are updated.

Initially, the interest scores for the concepts in the profile will continue to change. However, once enough information has been processed for profiling, the amount of change in in-

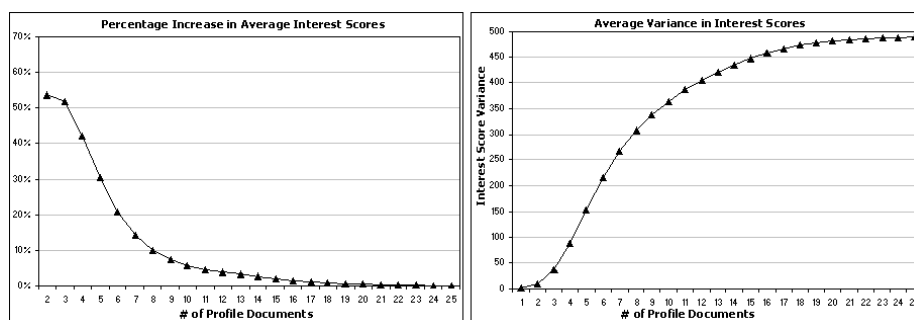


Figure 4: The average rate of increase and average variance in *Interest Scores* as a result of incremental updates.

terest scores should decrease. Our expectation is that eventually the concepts with the highest interest scores should become relatively stable. Therefore, these concepts will reflect the user’s primary interests.

To evaluate the user profile convergence, we used a single profile document for each concept and utilized that document as the input for the spreading activation algorithm for 25 rounds. We utilized the documents in the *profile set* for this experiment. For each concept, we used a profile document that was originally indexed under that specific concept, which we refer to as the signal concept.

Our methodology was as follows. We started with a given signal document and used a profile document to spread activation. As described in Section 3.4, after the propagation through the entire network is completed, the interest scores are normalized and updated. We recorded the interest scores for all concepts as well as the average interest score and variance across all concepts. This was considered round 1. For the same signal concept, we repeated the process for 25 rounds which is equivalent to updating the ontological user profile using 25 profile documents.

We ran the above experiment for 50 distinct signal concepts. The interest scores in the user profile were reset to one prior to processing each signal concept. Our goal was to measure the change in interest scores for the signal concept as well as the other concepts in the user profile.

As depicted in Figure 4, the average rate of increase for the interest scores for the signal concepts did converge. However, monitoring the interest scores for the signal concepts was not sufficient by itself. We needed to guarantee that the interest scores for all of the other concepts were not increasing at the same rate as the signal concept. Therefore, we computed the variance in interest scores after each round for a given signal concept.

Our expectation was that additional evidence in favor of a signal concept should result in discrimination of the signal concept from other concepts in the user profile. Figure 4 displays the average variance as a result of incremental updates. Our results on the average rate of increase and the average variance in interest scores together demonstrate that the interest scores in the user profiles are getting updated accurately.

5.3.2 Re-ranking Web Search Results

We constructed keyword queries to run our automated experiments. We decided to extract the query terms from the concept term vectors in the ontology. Each concept term

vector was sorted in descending order with respect to term weights.

Query	# of Terms	Criteria
Set 1	1	highest weighing term in concept term vector
Set 2	2	two highest weighing terms in concept term vector
Set 3	3	three highest weighing terms in concept term vector
Set 4	2 or more	overlapping terms within highest weighing 10 terms

Table 1: Set of Keyword Queries

Table 1 depicts the four query sets that were automatically generated for evaluation purposes. Our keyword queries were used to run a number of automated search scenarios for each concept in our reference ontology. The first set of keyword queries contained only one term and included the highest weighing term for each concept. In order to evaluate the search results when a single keyword was typed by the user as the search query, the assumption was that the user was interested in the given concept.

The second set of queries contained two terms including the two highest weighing terms for each concept. The third set of queries were generated using the three highest weighing terms for each concept. As the number of keywords in a query increase, the search query becomes less ambiguous.

Even though one to two keyword queries tend to be vague, we intentionally came up with a fourth query set to focus specifically on ambiguous queries. Each concept term vector was sorted with respect to term weights. We compared the highest weighing ten terms in each concept with all other concepts in the ontology. A given concept was considered to be overlapping with another concept if a specific term appeared in the term vectors of both concepts. The parents, children, and siblings of the concept were excluded when identifying the overlapping concepts for a given concept. Only the overlapping concepts were included in the experimental set with each query consisting of two or more overlapping terms within these concepts.

Our evaluation methodology was as follows. We used the system to perform a standard search for each query. As mentioned above, each query was designed for running our experiments for a specific concept. In the case of standard search, a term vector was built using the original keyword(s) in the query text. Removal of stop words and stemming




Figure 5: Average *Top-n Recall* and *Top-n Precision* comparisons between the personalized search and standard search using “overlap queries”.




Figure 6: Percentage of improvement in *Top-n Recall* and *Top-n Precision* achieved by personalized search relative to standard search with various query sizes.

was utilized. Each term in the original query was assigned a weight of 1.0.

The search results were retrieved from the test set, the signal and noise document collection, by using a cosine similarity measure for matching. Using an interval of ten, we calculated the *Top-n Recall* and *Top-n Precision* for the search results.

Next, documents from the profile set were utilized to simulate user interest for the specific concept. For each query, we started with a new instance of the ontological user profile with all interest scores initialized to one. Such a user profile represents a situation where no initial user interest information is available. We performed our spreading activation algorithm to update interest scores in the ontological user profile.

After building the ontological user profile, we sorted the original search results based on our re-ranking algorithm and computed the *Top-n Recall* and *Top-n Precision* with the personalized results. In order to compare the standard search results with the personalized search results, we computed the average *Top-n Recall* and *Top-n Precision*, depicted in Figure 5. We have also computed the percentage of improvement between standard and personalized search for *Top-n Recall* and *Top-n Precision*, depicted in Figure 6.

5.4 Discussion of Experimental Results

When using a search engine, users typically formulate ambiguous queries which contain between one to three key-

words. The search results that are returned from the search engine may satisfy the search criteria but often fail to meet the user’s search intention.

Personalized search provides the user with results that accurately satisfy their specific goal and intent for the search. The queries used in our experiments were intentionally designed to be short to demonstrate the effectiveness of our Web search personalization approach, especially in the typical case of Web users who tend to use very short queries.

Simulating user behavior allowed us to run automated experiments with a larger data set. In the worst case scenario, the user would enter only a single keyword. The evaluation results show significant improvement in recall and precision for single keyword queries as well as gradual enhancement for two-term and three-term queries. As the number of keywords in a query increase, the search query becomes more clear.

In addition to the one, two, and three keyword queries, we ran experiments with the overlap query set to focus on ambiguous queries. Two users may use the exact same keyword to express their search interest even though each user has a completely distinct intent for the search. For example, the keyword *Python* may refer to *python as a snake* as well as the *Python programming language* sense.

The purpose of the overlap queries is to simulate real user behavior where the user enters a vague keyword query as the search criteria. Our evaluation results verify that using the ontological user profiles for personalizing search results is an

effective approach. Especially with the overlap queries, our evaluation results confirm that the ambiguous query terms are disambiguated by the semantic evidence in the ontological user profiles.

6. CONCLUSIONS AND OUTLOOK

We have presented a framework for contextual information access using ontologies and demonstrated that the semantic knowledge embedded in an ontology combined with long-term user profiles can be used to effectively tailor search results based on users' interests and preferences.

In our future work, we plan to continue evaluating the stability and convergence properties of the ontological profiles as interest scores are updated over consecutive interactions with the system. Since we focus on implicit methods for constructing the user profiles, the profiles need to adapt over time. Our future work will involve designing experiments that will allow us to monitor user profiles over time to ensure the incremental updates to the interest scores accurately reflect changes in user interests.

7. REFERENCES

- [1] M. Aktas, M. Nacar, and F. Menczer. Using hyperlink features to personalize web search. In *Advances in Web Mining and Web Usage Analysis, Proceedings of the 6th International Workshop on Knowledge Discovery from the Web, WebKDD 2004*, Seattle, WA, August 2004.
- [2] H. Alani, K. O'Hara, and N. Shadbolt. Ontocopi: Methods and tools for identifying communities of practice. In *Proceedings of the IFIP 17th World Computer Congress - TC12 Stream on Intelligent Information Processing*, pages 225–236, Deventer, The Netherlands, The Netherlands, 2002.
- [3] J. Allan, et al. Challenges in information retrieval and language modeling. *ACM SIGIR Forum*, 37(1):31–47, 2003.
- [4] O. Boydell and B. Smyth. Capturing community search expertise for personalized web search using snippet-indexes. In *Proceedings of the 15th ACM International Conference on Information and Knowledge Management, CIKM 2006*, pages 277–286, Arlington, VA, November 2006.
- [5] H. Chang, D. Cohn, and A. McCallum. Learning to create customized authority lists. In *Proceedings of the 7th International Conference on Machine Learning, ICML 2000*, pages 127–134, San Francisco, CA, July 2000.
- [6] P. Chirita, C. Firan, and W. Nejdl. Summarizing local context to personalize global web search. In *Proceedings of the 15th ACM International Conference on Information and Knowledge Management, CIKM 2006*, pages 287–296, Arlington, VA, November 2006.
- [7] P. A. Chirita, W. Nejdl, R. Paiu, and C. Kohlschutter. Using odp metadata to personalize search. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2005*, pages 178–185, Salvador, Brazil, August 2005.
- [8] P. A. Chirita, D. Olmedilla, and W. Nejdl. Pros: A personalized ranking platform for web search. In *Proceedings of the 3rd International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems, AH 2004*, Eindhoven, The Netherlands, August 2004.
- [9] S. Dumais, T. Joachims, K. Bharat, and A. Weigend. Implicit measures of user interests and preferences. *ACM SIGIR Forum*, 37(2), 2003.
- [10] S. Gauch, J. Chaffee, and A. Pretschner. Ontology-based personalized search and browsing. *Web Intelligence and Agent Systems*, 1(3-4), 2003.
- [11] T. R. Gruber. Towards principles for the design of ontologies used for knowledge sharing. In *Formal Ontology in Conceptual Analysis and Knowledge Representation*, Deventer, The Netherlands, 1993.
- [12] H. Haav and T. Lubi. A survey of concept-based information retrieval tools on the web. In *5th East-European Conference, ADBIS 2001*, pages 29–41, Vilnius, Lithuania, September 2001.
- [13] T. H. Haveliwala. Topic-sensitive pagerank. In *Proceedings of the 11th International World Wide Web Conference, WWW 2002*, Honolulu, Hawaii, May 2002.
- [14] G. Jeh and J. Widom. Scaling personalized web search. In *Proceedings of the 12th international conference on World Wide Web, WWW 2003*, pages 271–279, Budapest, Hungary, May 2003.
- [15] R. Kraft, F. Maghoul, and C. C. Chang. Y!q: contextual search at the point of inspiration. In *Proceedings of the 14th ACM International Conference on Information and Knowledge Management, CIKM 2005*, pages 816–823, Bremen, Germany, November 2005.
- [16] S. Lawrence. Context in web search. *IEEE Data Engineering Bulletin*, 23(3):25–32, 2000.
- [17] F. Liu, C. Yu, and W. Meng. Personalized web search for improving retrieval effectiveness. *IEEE Transactions on Knowledge and Data Engineering*, 16(1):28–40, 2004.
- [18] A. Micarelli and F. Sciarrone. Anatomy and empirical evaluation of an adaptive web-based information filtering system. *User Modeling and User-Adapted Interaction*, 14(2-3):159–200, 2004.
- [19] S. Middleton, N. Shadbolt, and D. D. Roure. Capturing interest through inference and visualization: Ontological user profiling in recommender systems. In *Proceedings of the International Conference on Knowledge Capture, K-CAP 2003*, pages 62–69, Sanibel Island, Florida, October 2003.
- [20] M. Porter. An algorithm for suffix stripping. *Program*, 14(3):130–137, 1980.
- [21] F. Qiu and J. Cho. Automatic identification of user interest for personalized search. In *Proceedings of the 15th International World Wide Web Conference, WWW 2006*, pages 727–736, Edinburgh, Scotland, May 2006.
- [22] D. Ravindran and S. Gauch. Exploring hierarchical relationships in conceptual search. In *Proceedings of the 13th International Conference on Information and Knowledge Management, ACM CIKM 2004*, Washington DC, November 2004.
- [23] C. Rocha, D. Schwabe, and M. P. de Aragao. A hybrid approach for searching in the semantic web. In *Proceedings of the 13th international conference on*

- World Wide Web, WWW 2004*, pages 374–383, New York, NY, USA, 2004.
- [24] G. Salton and C. Buckley. On the use of spreading activation methods in automatic information. In *Proceedings of the 11th annual international ACM SIGIR conference on Research and Development in Information Retrieval, SIGIR 1988*, pages 147–160, Grenoble, France, 1988.
- [25] G. Salton and M. McGill. *Introduction to Modern Information Retrieval*. McGraw-Hill, New York, NY, 1983.
- [26] B. Schilit and M. Theimer. Disseminating active map information to mobile hosts. *IEEE Network*, 8(5):22–32, 1994.
- [27] X. Shen, B. Tan, and C. Zhai. Ucair: Capturing and exploiting context for personalized search. In *Proceedings of the Information Retrieval in Context Workshop, SIGIR IRiX 2005*, Salvador, Brazil, August 2005.
- [28] A. Sieg, B. Mobasher, S. Lytinen, and R. Burke. Using concept hierarchies to enhance user queries in web-based information retrieval. In *Proceedings of the International Conference on Artificial Intelligence and Applications, IASTED 2004*, Innsbruck, Austria, February 2004.
- [29] A. Singh and K. Nakata. Hierarchical classification of web search results using personalized ontologies. In *Proceedings of the 3rd International Conference on Universal Access in Human-Computer Interaction, HCI International 2005*, Las Vegas, NV, July 2005.
- [30] M. Speretta and S. Gauch. Personalized search based on user search histories. In *Proceedings of the 2005 IEEE/WIC/ACM International Conference on Web Intelligence, WI 2005*, pages 622–628, Compigne, France, September 2005.
- [31] A. Spink, H. Ozmutlu, S. Ozmutlu, and B. Jansen. U.s. versus european web searching trends. *ACM SIGIR Forum*, 15(2), 2002.
- [32] F. Tanudjaja and L. Mui. Persona: A contextualized and personalized web search. In *Proceedings of the 35th Annual Hawaii International Conference on System Sciences, HICSS 2002*, page 67, Big Island, Hawaii, January 2002.
- [33] J. Teevan, S. Dumais, and E. Horvitz. Personalizing search via automated analysis of interests and activities. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2005*, pages 449–456, Salvador, Brazil, August 2005.
- [34] J. Trajkova and S. Gauch. Improving ontology-based user profiles. In *Proceedings of the Recherche d'Information Assistée par Ordinateur, RIAO 2004*, pages 380–389, University of Avignon (Vaucluse), France, April 2004.
- [35] C. Ziegler, K. Simon, and G. Lausen. Automatic computation of semantic proximity using taxonomic knowledge. In *Proceedings of the 15th ACM International Conference on Information and Knowledge Management, CIKM 2006*, pages 465–474, Arlington, VA, November 2006.