

SonetRank: Leveraging Social Networks to Personalize Search

Abhijith Kashyap
Univ. of California at Riverside,
Riverside, CA. USA.
akash001@ucr.edu

Reza Amini
Florida International University
Miami, FL. USA
ramin001@fiu.edu

Vagelis Hristidis
Univ. of California at Riverside,
Riverside, CA. USA.
vagelis@cs.ucr.edu

ABSTRACT

Earlier works on personalized Web search focused on the click-through graphs, while recent works leverage social annotations, which are often unavailable. On the other hand, many users are members of the social networks and subscribe to social groups. Intuitively, users in the same group may have similar relevance judgments for queries related to these groups. SonetRank utilizes this observation to personalize the Web search results based on the aggregate relevance feedback of the users in similar groups. SonetRank builds and maintains a rich graph-based model, termed *Social Aware Search Graph*, consisting of groups, users, queries and results click-through information. SonetRank's personalization scheme learns in a principled way to leverage the following three signals, of decreasing strength: the personal document preferences of the user, of the users of her social groups relevant to the query, and of the other users in the network. SonetRank also uses a novel approach to measure the amount of personalization with respect to a user and a query, based on the query-specific richness of the user's social profile. We evaluate SonetRank with users on Amazon Mechanical Turk and show a significant improvement in ranking compared to state-of-the-art techniques.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval - *relevance feedback, information filtering.*

Keywords

Search Personalization, Social Search, Results Re-ranking.

1. INTRODUCTION

Many queries issued by users on a search engine are inherently ambiguous and have multiple interpretations [1]. As an example, consider the query "City of Hope", which returns the critically acclaimed 1991 film with the same name as the query and websites related to a city named City of Hope in Arkansas, in addition to websites for renowned cancer research institution, 'City of Hope'. Personalization is used to enhance the search experience of users based on a *profile* of the individual user and using this profile to *re-rank* her query results. For example, if the search engine knows that the user lives in Arkansas or has issued queries or clicked results relevant to the city 'City of Hope', then the search engine might rank results pertaining to this city higher.

The user profile may consist of information *explicitly* marked by

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, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CIKM'12, October 29–November 2, 2012, Maui, HI, USA.
Copyright 2012 ACM 978-1-4503-1156-4/12/10...\$15.00.

the user as interesting or *implicit* indicators of user preferences such as query history [2] and click logs [3]. Recently several works have proposed inferring these preferences by automatically grouping (clustering) users [7, 13] and using aggregated group preferences to enhance ranking of results for queries that are related to these groups. However, these approaches typically generate the same ranking for each group. Further, these approaches do not utilize user-specific preferences, e.g. Alice clicked on a specific cancer institute in the past.

The emergence of social networks like Facebook provides new opportunities to personalize search. Users spend a considerable amount of time on social networks and build their social profile explicitly by subscribing to groups that reflect to a large extent, their topics of interest. This provides a unique opportunity for Web search personalization that, to the best of our knowledge, has not been leveraged before. We refer to this setting, which unifies the Web search activity with social preferences, as **Social-Aware Search (SAS)**. That is, SAS leverages the social interactions of users to personalize their queries on the Web. This is in contrast to *social search*, which generally refers to searching the content published by your social connections.

Personalizing solely on the subscribed social groups would incur similar limitations to the above mentioned works, that automatically cluster users. Hence, a successful personalization scheme must combine the user-specific query (and click-through) history with aggregated behavior of users in query-related groups in her social profile. Furthermore, the relevance judgments provided by the user on this personalized query can now be used to enhance the personalization for other users who have groups in common with the said user. To illustrate this point, consider user Alice who has subscribed to groups for *Leukemia Awareness* and treatment facility *City of Hope*. If she executes query '*blood cancer*' which is related (e.g. as measured by text similarity to group description) to both aforementioned groups, then results of similar queries by other users of these groups should be ranked higher. While viewing the results, Alice could mark some of the results as relevant, for e.g. by clicking on them. One of the results she marks is a link to a page about research updates on blood cancer treatments on the *City of Hope* website. Now imagine another user Bob who is a member of *Leukemia* group but not of *City of Hope*. Bob also executes the query *blood cancer* and receives relevant results that other users in *Leukemia* group find relevant. In addition, Bob could also see the result for research updates at *City of Hope* marked by Alice, even though he is *not* a member of the group, but is related to a group that is *similar*, because it has similar description or because the groups have many common members.

Contributions: We propose SonetRank, which combines the factors identified above: (a) the individual user preferences, (b) preferences in the user's groups or similar groups related to query, (c) preferences of similar queries from other groups. To achieve

this, we propose a rich graph-based model, called *Social-Aware Search (SAS)* graph model (Section 2) that captures users, queries, groups, documents, and their associations. We propose an authority-flow based algorithm (Section 3.1) on the SAS graph to estimate the best search results based on the user’s profile. A key challenge here is to decide the extent of personalization for a given user query. For that, SonetRank estimates the query- and user-specific *confidence factor* (Section 3.2), which measures how rich and accurate the SAS graph is with respect to the user and the query. We evaluated SonetRank with real users from Amazon Mechanical Turk (Section 4) and achieve a significant improvement over state-of-the-art techniques.

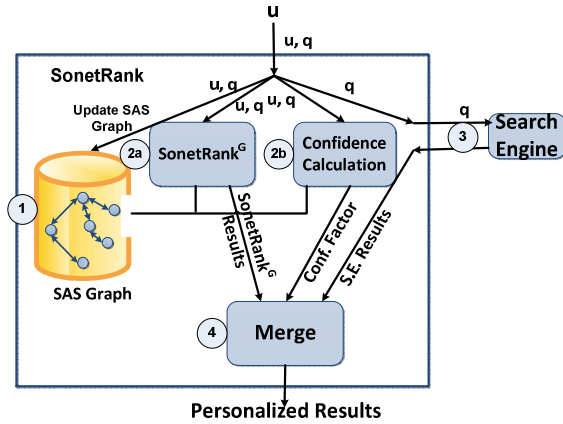


Figure 1. SonetRank Solution

2. SOCIAL AWARE SEARCH MODEL

In SAS settings, the following preferences are available:

1. **Individual Preferences:** A user is likely to prefer results of a query that she found relevant in the past for the same query or for *similar* queries.
2. **Group Preferences:** Results that other users belonging to same or similar groups found relevant for the same or similar queries might also be relevant to the user.
3. **Query Preferences:** Results judged relevant by users of different groups for the same or similar queries. This case can arise if the social profile information is partial and does not include all the interests of the user.

Intuitively, the first type of preference is stronger than the second and the second is stronger than the third. By combining all the three in a principled way, SonetRank improves the personalization experience. This is in contrast to previous works which consider only a single type of such diverse preferences.

Solution Overview: SonetRank maintains a rich model of user associations in social groups and their search preferences in a graph-based model termed the Social-Aware Search (SAS) graph (described next) which captures various entities in the social aware search settings, such as users, groups, queries and documents, and relationships (edges) between these entities. The SAS graph is incrementally built as users join in, subscribe to groups, execute queries and select (click) results.

Figure 1 shows in numbered circles the key steps in personalizing query q issued by user u :

1. The SAS graph is updated to record this interaction.
2. The user-query (u, q) is evaluated over the SAS graph.

- a. A personalized ranking is generated based on the SAS graph, using the authority-flow based [5, 6] SonetRank^G algorithm.
- b. For some queries, the SAS graph may not have enough information to personalize the results for the user u . SonetRank judges the *capability* of the SAS graph to personalize query q for user u , by calculating a *confidence factor*.
3. Query q is evaluated on a search engine to receive fresh results. This step can occur in parallel to Steps 1 or 2.
4. Finally, the result rankings of q evaluated on SAS graph (using SonetRank^G) are merged with the results of q obtained from a search engine, using the confidence factor.

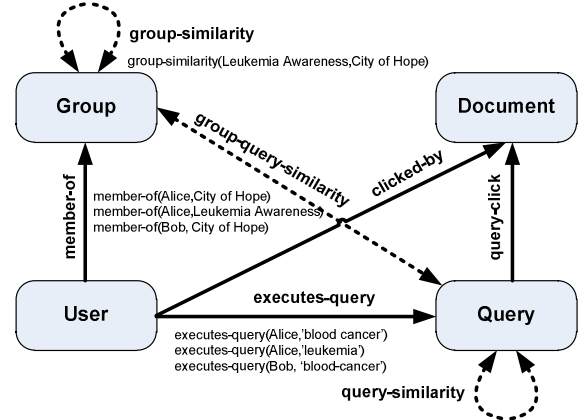


Figure 2. Social Aware Search (SAS) Graph Model

SAS Graph Model: The SAS Graph is illustrated in Figure 2 and shows the various entity types in a SAS setting, namely users, groups, queries and documents (results). The figure also shows the relationship types between entities, e.g., *member-of* relationship captures association of a user to a group. These relationships can be of two types:

Interaction relationships: These relationships are depicted by solid lines and capture the actions and preferences of users.

1. **Group Membership** (*member-of*(user,group)): Users join an online group. The relationship *member-of* captures this association.
2. **User Query Execution** (*executes-query*(user,query)): Query executions by users are captured by the *executes-query* relationship. This relationship instance is updated each time a user u issues a query q .
3. **User Clicks** (*clicked-by*(document, user)): Users click on documents in query results that they find interesting. These implicit relevance judgments are indicators of the preference of a user u for a document d .
4. **Query Click** (*query-click*(document, query)): Analogous to preferences over result documents by users, *query-click* captures the preferences for a document d for query q .

Semantic similarity relationships: The semantic similarity edges between entities, represented by dashed-edges in Figure 2, are used to leverage approximate matches of entities in the SAS graph. For example, if a user subscribed to group “European soccer” submits query “most expensive players”, previously clicked results for the same or similar queries from users subscribed to similar groups could be leveraged. We consider three types of semantic similarities in this work (a) between groups (*group-similarity*) (b) between queries (*query-similarity*), and (c) between a group and a query (*group-query-similarity*).

SAS Graph: The SAS Graph $G^D(V^D, E^D)$ is a labeled directed graph where each node $v \in V^D$ has an associated $type(v) \in \{user, group, query, document\}$ and the edge set E^D captures and quantifies the semantic relationships between nodes in the SAS Graph Model. This graph is created in order to enable us to rank documents using authority-flow algorithms. For each relationship between two entities (u, v) in the SAS Graph Model, we create two edges $e^f(u \rightarrow v)$ and $e^b(v \rightarrow u)$ and each edge $e \in E^D$ has an associated weight $w(e) \rightarrow [0,1]$ to capture the strength of association between the entities. Due to space limitations, we omit the discussion on setting these weights, which are set based on intuitive measures of the strength of association between nodes in the SAS Graph.

3. SONETRANK RANKING FUNCTIONS

Section 3.1 presents SonetRank^G, which computes a personalized ranking for the user-query on the SAS Graph. Section 3.2 describes how the ranking produced by SonetRank^G is combined with the ranking produced by the search engine.

3.1 SonetRank^G: Ranking on SAS Graph

SonetRank^G ranks documents on the SAS graph $G^D(V^D, E^D)$ using an adaptation of the personalized PageRank [5]. The personalized PageRank generates a ranking of documents using the following equation:

$$\mathbf{r} = c\mathbf{A}\mathbf{r} + (1 - c)\mathbf{s}/|S| \quad (1)$$

where, $\mathbf{r} = \{r(v_1), \dots, r(v_n)\}$ is the rank vector and \mathbf{A} is the $n \times n$ transition matrix of transition probabilities with $A_{ij} = 1/O(v_i)$, where $O(v_i)$ is the out-degree of node v_i . The algorithm to compute ranking in Equation 1, works by biasing the ranking towards a *base set* S of nodes in the graph. Typically, this base-set consists of a small set of documents (web-pages) that represent the personalized preference of a given user (or query).

The first key difference from PageRank, is that SonetRank^G operates on the heterogeneous SAS Graph consisting of varied entities and relationships between them. In contrast, PageRank operates on a homogenous graph of Web pages and their inter-links. In such a heterogeneous scenario, it is necessary to distribute authority fairly to the neighbors of various types, to avoid biasing PageRank computation to favor types of nodes that have many connections to nodes of same type [9]. For example, consider a SAS graph where a query q is connected to 100 other queries (based on similarity relationships) and to 2 users who executed this query. By using personalized PageRank equation directly, the transition probability to each of the 102 nodes is $1/102$ and therefore the computation will transfer most authority to other queries at the expense of user nodes.

To ensure *fair* distribution of authority, we define the transition probability on the edge $e: v_i \rightarrow v_j$ based on the type of the target node as $1/O(v_i, type(v_j))$ instead of $1/O(v_i)$ in PageRank, where $type(v_j)$ is the type (user, group, query or document) of node v_j , and $O(v_i, type(v_j))$ is the number of outgoing edges from v_i of type $type(v_j)$. To ensure that the transition probability distribution is valid, we multiply these transition probabilities by an *authority transfer factor* $\alpha(v_i, v_j)$ to normalize the transition probabilities, so that they sum to at most 1.0 for a given node. Therefore the transition probability A_{ij} is given by.

$$A_{ij} = \begin{cases} \left(\frac{\alpha(v_i, v_j) \times w(v_i, v_j)}{O(v_i, type(v_j))} \right) & \text{if } O(v_i, type(v_j)) > 0 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

In this work, we set $\alpha(u, v)$ to $1/4$ for all node pairs, given that we have 4 types of entities.

Another key difference from PageRank is the *base set*. We are interested in ranking the results of a particular query q issued by a user u . That is, we are looking for documents that are related to both the q and u through the SAS graph. Therefore, we set the base set S to $\{q, u\}$. The choice of this base set biases SonetRank^G towards nodes in SAS graph G^D that are in close proximity to the user and the query in G^D . These closely associated nodes include previous queries executed by the user and the documents judged relevant for those queries.

3.2 Search Engine Ranking

Relying solely on the personalized ranking generated by SonetRank^G, we run the risk of returning an empty result or results that include many irrelevant documents, since the SAS graph has relevance information for a limited number of documents and queries. For example, if the SAS Graph contains groups and queries related to *City of Hope* and *Leukemia* as in the example from Section 1, and a user executes a query related to movies (e.g. *Godfather*), then this query cannot be personalized using this SAS graph and would return an empty answer to the query. Furthermore, if the graph contains data for a small number of queries and one or more documents is preferred (clicked) by a large number of them, then this document would be ranked high for any marginally related query.

To avoid these scenarios, we base our ranking on a combination of global relevance of documents returned by executing the query on a search engine and the ranking based on the social-aware interactions, SonetRank^G. Therefore, in addition to computing the SonetRank^G ranking, we also compute a *confidence factor* to judge the quality of the SonetRank^G ranking. Using this confidence factor, we appropriately merge the rankings obtained from SonetRank^G and those obtained by executing the query on the search engine.

Merge Rankings: To merge the list of documents (d^{SE}) returned by the Search Engine and the list of documents returned by the SonetRank^G (d^{SR}), we first get the union of these two lists, and then compute the final score of each document d_i by a linear combination of its corresponding scores in each list (Equation 3).

$$d_i = C_f \cdot score(d_i^{SR}) + (1 - C_f) \cdot score(d_i^{SE}) \quad (3)$$

where $score(d)$ is the score of the document as returned by SonetRank^G or the Search Engine and C_f is the confidence factor, as defined next.

Confidence factor: A high value of C_f would give more importance to SonetRank^G scores and documents that have high SonetRank^G scores would be ranked higher in the final ranking. Some documents that are not returned by the search engine might also be included in the query since they are relevant based on user interactions. Intuitively, the confidence factor is a measure of the suitability of the SAS Graph in personalizing the results of a query q for a given user u . In the SAS graph, association between nodes is formed by user interaction or by semantic relationship between nodes. This user-query association is *strong* when the same documents that were marked relevant for query q were also previously preferred by the user u . Therefore, to compute the confidence factor C_f , we need to measure the strength of association between user and query nodes in the SAS Graph. SimRank [10] measures the strength of association between pairs of nodes in a graph G based on the recursive notion of structural similarity which states that two nodes are *similar* if they are

referenced by similar nodes. Therefore, to calculate the confidence factor C_f , we use SimRank [10] and compute C_f as the SimRank score the pair of user and query in the SAS Graph. We evaluate the effectiveness of the confidence factor in Section 4.

4. EXPERIMENTAL EVALUATION

Our aim in these experiments is to evaluate the effectiveness of the SonetRank, which enhances the search results of a query issued by a particular user in a *social aware setting*, with aggregate relevance judgments of users who belong to similar groups as the user or on similar previously executed queries. We evaluate SonetRank by means of a *controlled* user-study in which we predefine groups that users can belong to and also select the queries they would execute.

Table 1. Social Groups used for Evaluation

#	Group	Sample Queries
1	Comedy Movies	<ul style="list-style-type: none"> best drama movies in 1936 comedy drama movies of 20's
2	Charlie Chaplin Fans	<ul style="list-style-type: none"> city lights actor comedy movie Harry Crocker
3	Jim Carey Fans	<ul style="list-style-type: none"> cast of The Mask cast of Dumb and Dumber
4	Mafia Movies	<ul style="list-style-type: none"> classic crime movie in 1972 Joe Pesci classic thriller movie
5	Robert De Niro Fans	<ul style="list-style-type: none"> crime movie with Diane Keaton Taxi Driver cast
6	Al Pacino Fans	<ul style="list-style-type: none"> cast of Godfather 3 cast crime movie with James Caan

We defined six social groups (Table 1) around the widely-familiar subject of movies. The groups in Table 1 are constructed around two movie genres: Comedy (1-3) and Mafia or Gangster (4-6). For each group, we selected 14 queries for our evaluation including the sample queries shown in Table 1. The queries were chosen such that they return a few results related to the corresponding group, e.g., the query ‘cast of Godfather 3’ returns the IMDB bio page of ‘Al Pacino’.

Baselines: Since our approach personalizes results based on a number of aspects namely group, query and document relevance judgments, we compare with baselines that offer varying degrees of personalization on each of the aspects:

- Baseline 1 (No Personalization):** The results returned by executing the query on a search engine. We use Google Custom Search API with disabled web-history personalization.
- Baseline 2 (Click-through Random Walk):** Based on [3] a query-dependent ranking is generated using the past relevance judgments (clicks). The relevance of a document is computed using a backward-random walk from document to query on a click-graph, which contains users and documents.
- Baseline 3 (SonetRank-NoGroups):** Group nodes and all associated edges are removed from the SAS graph to study the effect of groups in improving personalization.

Methodology: We invited users via Amazon Mechanical Turk. Users were presented with the textual description of the groups and assigned up to 3 groups. The evaluation was performed in two stages (a) Training and (b) Evaluation.

- Training:** Users execute 3 (out of 10) queries at random. These queries are executed against the search engine and the results are presented to the user, who marks all relevant

results, based on group choices. These relevance indications were added to the SAS graph as preferences of the user.

- Evaluation:** We re-invited users from the training phase and these users executed 4 *additional* queries for each selected group. For each query, the results were personalized independently using SonetRank and the three baselines and the results from all algorithms were merged and presented to the user after randomizing the order. As in the training phase, we asked users to mark all relevant results.

The effect of personalization using SonetRank increases with increase in *richness* of the SAS graph as more users join in, associate with groups, execute queries and mark relevant results. To study this effect, we performed evaluation in three consecutive phases, designated as Phases 1, 2 and 3. In each phase, the SAS Graph is augmented with relevance judgments from additional users. A total of 74 users participated in the evaluation, with an average of 24 users in each phase.

Metrics: In a personalized search setting, the notion of global ground-truth to judge the relevance of query results does not exist. Therefore, we propose evaluation measures based on the *user-specific ground truth*, in which we consider the relevant documents to be those that were marked by the user for *any* query. We report the results of evaluation using the following metrics:

- Precision at K $Pre(K)$:** For a given user-query $Pre(K)$ is the fraction of the top K results that are labeled as relevant, as determined by the user-specific ground truth.
- Normalized Discounted Cumulative Gain $NDCG(K)$:** NDCG for a given query q is computed as follows:

$$NDCG(K) = \frac{1}{Z} \sum_{i=1}^K \frac{2^{r(i)} - 1}{\log(i + 1)} \quad (4)$$

where $r(i)$ is the binary relevance label of the result at rank position i and Z is the normalization factor.

Results: The results of the evaluation are shown in Table 2. The values reported in the table are averages over all queries executed by users across all phases. As seen in the table, SonetRank outperforms and achieves better personalization as compared to all other baselines. Improvement of SonetRank over Baseline 1 is significant with 62% for $Pre(5)$. Comparing to Baseline 2, SonetRank shows improvements in terms of both precision and significantly, NDCG scores (36% for $NDCG(5)$). The improvement in precision at 10 is lower, in all cases, because the results relevant in many evaluation queries appear within the first ten results. However, our Baseline3 and SonetRank are able to rank them higher, resulting in a higher NDCG.

Since Baseline 2 generates a query-dependent ranking based on past relevant judgments for a given query, it is able to include more relevant results (as compared to Baseline 1) in top 5 (and 10) and also rank them higher. Therefore, personalization performance is better than Baseline 1, and SonetRank outperforms due to the net effect of using aggregate group and query preferences.

Table 2: Precision and NDCG scores

	Baseline1	Baseline2	Baseline3	SonetRank
$Pre(5)$	0.053	0.073	0.080	0.086
$Pre(10)$	0.059	0.064	0.067	0.070
$NDCG(5)$	0.125	0.201	0.249	0.275
$NDCG(10)$	0.234	0.292	0.331	0.345

The effect of using aggregate user behavior across groups to personalize search can be seen by comparing the scores of

SonetRank with those of Baseline 3. The NDCG at 5 is 10% higher for SonetRank which is significant given that personalization based on groups carries a lesser significance than user's own relevance judgments on which the user specific ground truth is based. This shows that personalization can be improved by a shared notion of relevance among users across groups.

Effect of Confidence Factor: This ability of a SAS graph to personalize results for a particular user-query is determined by the confidence factor. Figure 3 shows the average confidence factor for SonetRank and SonetRank-NoGroups in each of the three phases. This factor grows with increased user activity as shown in Figure 3, from 0.11 for SonetRank in Phase 1 to 0.21 in Phase 3. Recall that the graph used by SonetRank-NoGroups is essentially the same as SAS Graph, but with all group nodes and associated edges removed. This removal results in a decreased value of confidence factor as the user and query become less *similar* as per SimRank semantics. The increased confidence factor in SonetRank leads to net increase in personalization as evidenced by increased average NDCG(10) for all these groups.

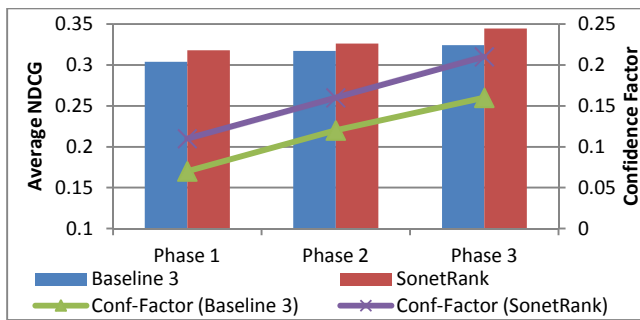


Figure 3. Conf. Factor and NDCG(10) scores for 3 Phases

5. RELATED WORK

Personalized Web-Search: Personalization entails modeling of preferences and interests of users using user-clicks [3, 11], data available on user's desktop [4] or user's browsing history and bookmarks [2, 12]. In [3], the results of a query are ranked based on a random walk on the click-graph of queries and clicks on documents while [11] uses a combination of topic-sensitive PageRank [8] and past user clicks to personalize search.

Clustering-Based Personalization: Recently, there have been several works where users [13, 7] or queries [14] are grouped together using clustering techniques and aggregated preferences within each cluster are used to personalize search. Instead, our work considers social groups, and combines the effect of a richer set of factors.

Social Annotation-based Personalized Search: [15] uses the annotation data on *del.icio.us* and proposes a ranking method based on tag-relevance and match between user profile and document tags. The *UserRank* algorithm [16], is an adaptation of PageRank on a semantically weighted graph, similar to our SAS graph. Our work is related to [17], which models relationships between entities (Groups, Users etc.) and proposes a linear ranking function based on familiarity and similarity.

6. CONCLUSIONS AND FUTURE WORK

In this paper, we proposed the problem of *Social-Aware Search*, where we leverage social links and the query history to personalize Web search results. Intuitively, we combine the user's past query results' preferences, the preferences of users with similar social profile and aggregate preference of all other users. These disparate factors are combined in a principled way that

weighs each factor's importance. We evaluated the solution and showed its effectiveness based on a user study. As part of future work, we plan to study the effect of various parameters such as different authority flow rates on search personalization, i.e. assign different importance to various relationship types. We will also evaluate our approach on a real-life social network.

7. ACKNOWLEDGEMENTS

This project was supported in part by National Science Foundation grants IIS-1216032 and IIS-1216007.

8. REFERENCES

- [1] Radlinski, F. and Dumais, S. Improving personalized web search using result diversification. In Proceedings of SIGIR 2006.
- [2] Matthijs, N. and Radlinski, F. Personalizing web search using long term browsing history. In Proceedings of WSDM 2011.
- [3] Craswell, N. and Szummer, M. Random walks on the click graph. In Proceedings of CIDR 2007.
- [4] Teevan, J., Dumais, S. T. and Horvitz, E. Personalizing search via automated analysis of interests and activities. In Proceedings of SIGIR 2005.
- [5] Jeh, G. and Widom, J. Scaling personalized web search. In Proceedings of WWW 2003.
- [6] Page, L., Brin, S., Motwani, R. and Winograd, T. The PageRank Citation Ranking: Bringing Order to the Web. Technical Report, Stanford InfoLab, 1999.
- [7] Teevan, J., Morris, M. R. and Bush, S. Discovering and using groups to improve personalized search. In Proceedings of WSDM 2009.
- [8] Haveliwal, T. Topic-Sensitive PageRank. In Proceedings of WWW 2002.
- [9] Balmin, A., Hristidis, V. and Papakonstantinou, Y. Objectrank: authority-based keyword search in databases. In Proceedings of VLDB 2004.
- [10] Jeh, G. and Widom, J. SimRank: a measure of structural-context similarity. In Proceedings of SIGKDD 2002.
- [11] Qiu, F. and Cho, J. Automatic Identification of User Interest For Personalized Search. In Proceedings of the WWW 2006.
- [12] Sugiyama, K., Hatano, K. and Yoshikawa, M. Adaptive web search based on user profile constructed without any effort from users. In Proceedings of the WWW2004.
- [13] Teevan, J., Dumais, S. T. and Horvitz, E. Potential for personalization. ACM TCHI, 17, 1 2010, 1-31.
- [14] Liu, F., Yu, C. and Meng, W. Personalized web search by mapping user queries to categories. In Proceedings of CIKM 2002.
- [15] Xu, S., Bao, S., Fei, B., Su, Z. and Yu, Y. Exploring folksonomy for personalized search. In Proceedings of SIGIR 2008.
- [16] Bender, M., Crecelius, T., Kacimi, M., Michel, S., Neumann, T., Parreira, J. X., Schenkel, R. and Weikum, G. Exploiting social relations for query expansion and result ranking. 2008.
- [17] Carmel, D., Zwerdling, N., Guy, I., Ofek-Koifman, S., Har'el, N., Ronen, I., Uziel, E., Yogev, S. and Chernov, S. Personalized social search based on the user's social network. In Proceeding of CIKM 2009.