Improving Personalized Search on the Social Web Based on Similarities between Users

Zhenghua Xu, Thomas Lukasiewicz, and Oana Tifrea-Marciuska

Department of Computer Science, University of Oxford, UK
{zhenghua.xu,thomas.lukasiewicz,oana.tifrea}@cs.ox.ac.uk

Abstract. To characterize a user's preferences and the social summary of a document, the user profile and the general document profile are widely adopted in existing folksonomy-based personalization solutions. However, in many real-world situations, using only these two profiles cannot personalize well the search results on the Social Web, because (i) different people usually have different perceptions about the same document, and (ii) the information contained in the user profile is usually not comprehensive enough to characterize a user's preference. Therefore, in this work, in order to improve personalized search on the Social Web, we propose a dual personalized ranking (D-PR) function, which adopts two novel profiles: an extended user profile and a personalized document profile. For each document, instead of using a general document profile for all users, our method computes for each individual user a personalized document profile to better summarize his/her perception about this document. A solution is proposed to estimate this profile based on the perception similarities between users. Moreover, we define an extended user profile as the sum of all of the user's personalized document profiles to better characterize a user's preferences. Experimental results show that our D-PR ranking function achieves better personalized ranking on the Social Web than the state-of-the-art baseline method.

1 Introduction

Recently, with the rise of Web 2.0 applications, such as social bookmarking systems, electronic commerce websites, blogs, and social network sites, the Web has evolved towards the so-called *Social Web*, where users can freely provide social annotations to online documents (i.e., Web pages or resources on the Social Web) via bookmarking, tagging, rating, commenting, and so on. Social annotations are valuable resources for personalized search on the Social Web. On the one hand, annotations provided by different Web users from different perspectives are usually good summaries of the corresponding documents. On the other hand, social annotations are also ideal data for privacy-enhanced personalization: first, they are provided by a user directly, so these annotations can be treated as a user's individual opinion about a document; these interests and preferences of the user can be harvested by the aggregation of his/her social annotations; second, these social annotations are usually publicly available and contain little sensitive information about users, so they can be safely utilized without violating user privacy. In this paper, we refer to social annotations as social tags assigned to documents by users in bookmarking systems, but relevant techniques can be easily adapted to other social metadata (e.g., comments, blogs, etc.) as well.

U. Straccia and A. Calì (Eds.): SUM 2014, LNAI 8720, pp. 306–319, 2014.

Consequently, more and more research activities focus on personalizing the search on the Social Web using social tags [3,4,25,26]. Generally, given a query issued by a user, the existing methods rank the online documents by the corresponding *ranking scores*, which are normally comprised of two parts: a query-related part, measuring the textual similarity between the given query and each document, and a personalization part, measuring the similarity between the user's preferences (in the user profile) and the social summary of each document (in the general document profile).

A user profile is a weighted vector, whose dimensions are tags and whose values in each dimensions are the corresponding tag weights. In the user profile, the tag weight is influenced by the number of times that this user uses the tag for bookmarking. Similarly, the general document profile is also a weighted vector and its tag weight is influenced by the number of times that the document is bookmarked with the tag.

However, in many real-world situations, using these two profiles cannot personalize well the search results on the Social Web. On one hand, users usually have different perceptions about the same document, so, for a specific user, not all tags assigned by all the other users are equally helpful to summarize his real perception about the document (some of them are actually harmful). Therefore, the general document profile, which treats tags from all users with equal importance, cannot properly summarize a special user's personal perception about the document. On the other hand, in practice, there are tens of billions of documents on the Web and even a long-time Social Web user can only annotate a very small portion of them. Therefore, the user profile, based on only the tags assigned by the corresponding user, usually does not contain sufficient information to comprehensively characterize the user's preferences.

To solve these problems, we propose a *dual personalized ranking (D-PR)* function which utilizes two novel profiles, called *personalized document profile* and *extended user profile*, to better characterize a user's preferences and better summarize his/her personal perception about a document, respectively. Instead of using the same general document profile for all users, for each of the documents, our method computes for each individual user a personalized document profile to characterize his/her personal perception about this document. Furthermore, the extended user profile is defined as the sum of all of the user's personalized document profiles. As each user has a personalized document profile for each of the documents, the extended user profile contains more information to comprehensively characterize a user's preferences.

However, how to obtain the user's personalized document profiles for all online documents is a challenge. The tags assigned by a user to a document may be a good outline of his personal perception about this document. But, in fact, this is unpractical: on the one hand, a user normally uses only a few (typically 1 or 2) tags to annotate a document, so these tags contain too little information to comprehensively summarize the document; on the other hand, only a small portion of online documents are annotated by a user, but we need a personalized document profile for each document.

Therefore, we propose to estimate the personalized document profile of a user u by using the perception similarity between u and the other users as weights to sum up tags assigned to the relevant document by the users having high perception similarity with u. The underlying intuition is that users having similar perceptions about the existing documents are very likely to also share similar perceptions about future documents; so, for a user u, tags assigned by users having high perception similarity with u are more helpful to characterize u's personal perception about the document than tags assigned by users having low perception similarity with u. Intuitively, the higher perception similarity between two users, the higher their tags are weighted for each other.

In summary, we make the following contributions in this paper.

- We propose a *dual personalized ranking (D-PR)* function to improve personalized search on the Social Web by introducing two novel profiles: the extended user profile and the personalized document profile, to better characterize a user's preferences and better summarize his/her personal perception about a document.
- We formally define the extended user profile as the sum of all the user's personalized document profiles; and we further propose to estimate a user's personalized document profile using the perception similarity between users. Finally, a method used to quantify the perception similarity is also presented.
- We conduct extensive experimental studies based on a public real-world large scale research dataset [15]. The results validate the effectiveness of our *D-PR* function: it outperforms the state-of-the-art *SoPRa* function [3].

The rest of this paper is organized as follows. In Section 2, we present some preliminaries. Section 3 formally defines two state-of-the-art personalized ranking solutions and illustrates their potential problems. In Section 4, we propose a novel D-PR function to solve these problems; while the approaches of estimating the user's personalized document profile and constructing the extended user profile are also presented in this section. Experiments are discussed in Section 5. Section 6 reviews some closely related works. Finally, Section 7 concludes this work and provides some future directions.

2 Preliminaries

Social bookmarking systems are based on the techniques of social tagging. The main idea behind them is to provide the user with a means to freely annotate resources on the Web (e.g., URIs in delicious[1] or images in Flickr[2]) with tags. Since the annotations can be shared with others, this practice of collaboratively creating and translating tags to annotate and categorize online content is usually called *collaborative tagging* or *social tagging*, and the resulting tag-based classification is called a *folksonomy*.

Definition 1. Let U, T, and D be the sets of *users*, *tags*, and *documents*. A *bookmark* is a triple $(u, t, d) \in U \times T \times D$, which represents the fact that the user u has annotated the document d with the tag t. A *folksonomy* $\mathcal{F}(U, T, D)$ is a subset of $U \times T \times D$.

The following example illustrates the above concepts, including folksonomies and bookmarks; it will be used in the sequel as a running example.

Example 1. Consider the set of users $U = \{Alice, Bob, Carl\}$, their set of tags $T = \{English, Chinese, Comedy, Action, Interesting, Boring\}$, and a set of documents $D = \{d_1, d_2, d_3\}$. A folksonomy $\mathcal{F}$ may then express the following knowledge:

[1] `https://delicious.com/`
[2] `https://www.flickr.com/`

Table 1. Tags used by users to annotate documents

	tags in d_1	**tags in** d_2	**tags in** d_3
Alice	English, Comedy, Interesting	Boring	Chinese, Comedy, Interesting
Bob	Boring	Chinese, Action, Interesting	Boring
Carl	English, Comedy, Interesting	Boring	(Null)

(i) $Alice$ and $Carl$ are interested in all comedies and dislike action movies, while Bob has the right opposite preferences; and (ii) d_1 is an introduction page of an English comedy movie, d_2 is an introduction page of a Chinese action movie, and d_3 is an online video of a Chinese comedy movie. The specific tags used by each of these users to annotate each of these online documents are shown in Table 1. □

The *personalized ranking problem* [3,25] can be fomalized as follows: given a folksonomy $\mathcal{F}(U, T, D)$ and a query q submitted by a user $u \in U$ to a search engine, it re-ranks the set of documents $d_q \in D$ that match q, in such a way that relevant documents for u are highlighted and pushed to the top for maximizing this user's satisfaction and personalizing the search results. The ranking follows an ordering $\tau = [d_1 \geq d_2 \geq \cdots \geq d_k]$ in which (i) $d_k \in D$ and (ii) $d_i \geq d_j$ iff $Rank(d_i, q, u) \geq Rank(d_j, q, u)$, where $Rank(d, q, u)$ is the result of a ranking function that quantifies similarity between q and the document d relative to u .

The vector space model (VSM) [19] is a general model used in information retrieval where the profile of a user (resp., a document) is mapped to a weighted vector in a universal term space. The terms can be tags or words. We use words when we deal with the text of the document and of the query, while tags are used when we deal with the tags of the document and of the query (each query word is considered a tag).

To calculate the similarity between two vectors, we use the well-known cosine similarity. Given two vectors $\boldsymbol{A} = (A_1, \ldots, A_n)$ and $\boldsymbol{B} = (B_1, \ldots, B_n)$, its cosine similarity $Sim(\boldsymbol{A}, \boldsymbol{B})$ is formally defined as follows, where $Sim(\boldsymbol{A}, \boldsymbol{B})$ ranges from 0 (independence) to 1 (identity):

$$Sim(\boldsymbol{A}, \boldsymbol{B}) = \frac{\boldsymbol{A} \cdot \boldsymbol{B}}{|\boldsymbol{A}||\boldsymbol{B}|} = \frac{\sum_{i=1}^{n} A_i \cdot B_i}{\sqrt{\sum_{i=1}^{n} A_i^2} \cdot \sqrt{\sum_{i=1}^{n} B_i^2}}. \tag{1}$$

We use the textual matching score, $Score(q, d)$, to indicate how similar a query q is to the textual content of a document d using words as terms. This score is not folksonomy-based and has been widely adopted in most commercial search engines. Thus, it can be obtained directly from these search engines, when it is incorporated into a personalized ranking function.

3 Personalized Ranking Functions

In this section, we recall two state-of-the-art personalized ranking methods and illustrate their potential ordering problems in the running example. All ranking functions presented in this paper follow the widely used VSM, where the weights of tags are based on tag frequencies (*tf*), and the extension to *tf-idf* (tag frequency-inverse document frequency) [10] is trivial.

3.1 User Profile Personalized Ranking Function

Xu et al. [26] propose a ranking function to compute the ranking score $Rank(d, q, u)$ of a document d relative to a given query q issued by a user u from two aspects: (i) the textual matching score $Score(q, d)$, measuring the statistical textual quality of d relative to q; and (ii) a profile matching score $Sim(\boldsymbol{p_u}, \boldsymbol{p_d})$, which estimates the interest of the user u in the document d, and which is measured by the similarity between the user profile and the general document profile. As this method uses the user's preferences that are implicitly contained in the user profile to personalize the ranking result, we call it *user profile personalized ranking (UP-PR) function*, formally defined as follows:

$$Rank(d, q, u) = \alpha \cdot Sim(\boldsymbol{p_u}, \boldsymbol{p_d}) + (1 - \alpha) \cdot Score(q, d), \tag{2}$$

where $\boldsymbol{p_u}$ is the user profile indicating this user's personal preferences, and $\boldsymbol{p_d}$ is the general document profile measuring the understandings and perceptions of all users about this document. Following the VSM, $\boldsymbol{p_u}$ (resp., $\boldsymbol{p_d}$) is a weighted vector with tags as dimensions and tag weights as values, where a tag's weight is the number of times that this tag is used by the user (resp., is used to annotate the document) for bookmarking. The following example illustrates the UP-PR function.

Example 2. Recalling Example 1, $Carl$ would like to find an interesting Chinese comedy film, so he issues a query "Interesting Chinese film" to a non-personalized search engine. Obviously, based on the knowledge in Example 1, $Carl$ would expect the ordering of the search result to be $\tau_0 = [d_3 \geq d_1 \geq d_2]$. However, the search engine computes $Score(q, d_1) = 0.6$, $Score(q, d_2) = 0.52$, and $Score(q, d_3) = 0.5$, i.e., the resulting ordering on the search results is $\tau_1 = [d_1 \geq d_2 \geq d_3]$, which is an unexpected ordering, as the desired document d_3 is ranked at the bottom.

On the other hand, if we use *UP-PR* to personalize the ranking result, then we first compute the weighted vectors of the query (denoted $\boldsymbol{q}$), the profile of $Carl$ (denoted $\boldsymbol{p_{Carl}}$), and the profiles of the documents (denoted $\boldsymbol{p_{d_1}}$, $\boldsymbol{p_{d_2}}$, and $\boldsymbol{p_{d_3}}$) as shown in Table 2. Then, the personalized *UP-PR* ranking scores of d_1, d_2, and d_3 for $Carl$ relative to this query can be computed as shown in Equation 3 with $\alpha = 0.5$. Therefore, the personalized ranking of these search results is $\tau_2 = [d_1 \geq d_3 \geq d_2]$.

$$\begin{aligned}
Rank(d_1, q, Carl) &= \frac{1}{2}(Sim(\boldsymbol{p_{Carl}}, \boldsymbol{p_{d_1}}) + Score(q, d_1)) \\
&= \frac{1}{2}(\frac{7}{\sqrt{4} \cdot \sqrt{13}} + 0.68) = \frac{1}{2}(0.97 + 0.6) = 0.79, \\
Rank(d_2, q, Carl) &= \frac{1}{2}(Sim(\boldsymbol{p_{Carl}}, \boldsymbol{p_{d_2}}) + Score(q, d_2)) \\
&= \frac{1}{2}(0.57 + 0.52) = 0.55, \\
Rank(d_3, q, Carl) &= \frac{1}{2}(Sim(\boldsymbol{p_{Carl}}, \boldsymbol{p_{d_3}}) + Score(q, d_3)) \\
&= \frac{1}{2}(0.75 + 0.5) = 0.63.
\end{aligned} \tag{3}$$

However, although τ_2 obtained via *UP-PR* is better than τ_1 (promoting d_3 from the bottom to the middle), τ_2 is still not the best ordering, as d_3 is ranked lower than d_1. Specifically, we note that in Equation 3, d_1 has a higher ranking score than d_3, which is intuitively inaccurate, because (based on the knowledge in Example 1) $Sim(\boldsymbol{p_{Carl}}, \boldsymbol{p_{d_3}})$ should have similar value to $Sim(\boldsymbol{p_{Carl}}, \boldsymbol{p_{d_1}})$ (as $Carl$ prefers all comedies), and $Score(q, d_3)$ should be the highest text matching score (as d_3 is a Chinese comedy film perfectly matching the query). In the following (sub)sections, we will analyze in detail the reasons for such an inaccurate ordering. □

3.2 Social Personalized Ranking Function

We obtain a low textual matching score $Score(q, d_3)$ in Example 2, because d_3 is an online video that has little textual content to compute a proper textual matching score. This problem is common on the Social Web, and, to solve it, Bouadjenek et al. [3] propose a *social personalized ranking (SoPRa) function*, which extends the *UP-PR* function in [26] by considering a new non-personalized matching score: the social matching score $Sim(\boldsymbol{q}, \boldsymbol{p_d})$ between the given query q and the social summary of document $\boldsymbol{p_d}$. This score indicates how relevant the social summary of a document d is to q. The intuition is to use social tags to better summarize the content of a document and add further information for social resources with very little textual content (e.g., videos and images). Therefore, we have two query-related scores in *SoPRa*, which are defined as follows:

$$Rank(d, q, u) = \alpha \cdot Sim(\boldsymbol{p_u}, \boldsymbol{p_d}) + (1 - \alpha) \cdot [\beta \cdot Sim(\boldsymbol{q}, \boldsymbol{p_d}) + (1 - \beta) \cdot Score(q, d)] \,. \quad (4)$$

Example 3. Continuing Example 2, if we personalize the results of the search engine by *SoPRa*, the ranking scores of d_1, d_2, and d_3 for $Carl$ are computed as shown in Equation 5 (with $\alpha = 0.5$ and $\beta = 0.5$).

$$\begin{aligned}
Rank(d_1, q, Carl) &= \frac{1}{2}(Sim(\boldsymbol{p_{Carl}}, \boldsymbol{p_{d_1}}) + \frac{1}{2}(Sim(\boldsymbol{q}, \boldsymbol{p_{d_1}}) + Score(q, d_1))) \\
&= \frac{1}{2}(0.97 + \frac{1}{2}(\frac{4}{\sqrt{3} \cdot \sqrt{13}} + 0.6)) = \frac{1}{2}(0.97 + \frac{1}{2}(0.64 + 0.6)) = 0.8, \\
Rank(d_2, q, Carl) &= \frac{1}{2}(Sim(\boldsymbol{p_{Carl}}, \boldsymbol{p_{d_2}}) + \frac{1}{2}(Sim(\boldsymbol{q}, \boldsymbol{p_{d_2}}) + Score(q, d_2))) \\
&= \frac{1}{2}(0.56 + \frac{1}{2}(0.44 + 0.52)) = 0.52, \qquad (5) \\
Rank(d_3, q, Carl) &= \frac{1}{2}(Sim(\boldsymbol{p_{Carl}}, \boldsymbol{p_{d_3}}) + \frac{1}{2}(Sim(\boldsymbol{q}, \boldsymbol{p_{d_3}}) + Score(q, d_3))) \\
&= \frac{1}{2}(0.75 + \frac{1}{2}(0.87 + 0.5)) = 0.72.
\end{aligned}$$

As we can see, the resulting ordering is the same as the one of *UP-PR*, which is not desired. Specifically, by using a social matching score, *SoPRa* narrows the gap between the ranking scores of d_1 and d_3, but the improvement is still not big enough to change the ordering of three documents. □

4 Dual Personalized Ranking Function

The reasons for having a low profile matching score $Sim(\boldsymbol{p_{Carl}}, \boldsymbol{p_{d_3}})$ in the previous examples are twofold: on the one hand, the general document profile $\boldsymbol{p_{d_3}}$ does not correctly characterize $Carl$'s real perception about d_3, since tags from all users are treated equally, and the tag from Bob brings a bias; on the other hand, the user profile $\boldsymbol{p_{Carl}}$ does not properly model $Carl$'s preference, because $\boldsymbol{p_{Carl}}$ does not tag d_3, so the information used for preference modeling is not comprehensive.

Generally, the widely used general document profile, which treats tags from all users with equal importance, may not be able to summarize a special user's personal perception about a document. Similarly, the information contained in the user profile (i.e., the tags assigned by the user) is usually insufficient to comprehensively characterize the preferences of the user.

Therefore, to solve these problems, we propose a new ranking function, which will be able to better personalize search results by introducing two novel profiles: the extended user profile and the personalized document profile, to better characterize a user's preferences and better summarize his/her personal perception about a document, respectively. Specifically, instead of using the same general document profile for all users, for each of the documents, each individual user has a personalized document profile to characterize his/her perception about this document. Furthermore, we define an extended profile of user u as $\boldsymbol{p'_u}$, which sums up all personalized document profiles of u to more comprehensively characterize u's preference. This ranking function is called *dual personalized ranking (D-PR)* function and formally defined as follows:

$$Rank(d, q, u) = \alpha \cdot Sim(\boldsymbol{p'_u}, \boldsymbol{p_{u,d}}) + (1-\alpha) \cdot [\beta \cdot Sim(\boldsymbol{q}, \boldsymbol{p_d}) + (1-\beta) \cdot Score(q, d)] \,, \quad (6)$$

where the personalized profile of a document d for a user u, $\boldsymbol{p_{u,d}}$, is a weighted vector of tags characterizing u's perception about d; while $\boldsymbol{p'_u}$ is an extended profile of u, obtained by summing up all personalized document profiles of u and defined as follows:

$$\boldsymbol{p'_u} = \sum_{i=1}^{|D|} \boldsymbol{p_{u,d_i}} \,. \quad (7)$$

Note that in Equation 6, we still use the general document profile $\boldsymbol{p_d}$ to compute the query-related social matching score $Sim(\boldsymbol{q}, \boldsymbol{p_d})$. As defined in Section 3.2, $Sim(\boldsymbol{q}, \boldsymbol{p_d})$ is a non-personalized matching score, measuring the textual similarity between q and the social summary of d, and it aims at using social tags assigned by all users to better summarize the content of a document, so here it is unreasonable to replace $\boldsymbol{p_d}$ by $\boldsymbol{p_{u,d}}$.

4.1 Personalized Document Profile

It is a challenge how to obtain the personalized document profiles of a user for all online documents. The tags assigned by a user to a document may be a good outline of this user's personal perception about this document. However, it is, in fact, not practical: on the one hand, a user normally uses only a few (typically, 1 to 3) tags to annotate a document, so these tags contain too little information to comprehensively summarize

Table 2. Weighted vectors of query and profiles

	English	Chinese	Comedy	Action	Interesting	Boring
$\boldsymbol{p}_{Alice}$	1	1	2	0	2	1
$\boldsymbol{p}_{Bob}$	0	1	0	1	1	2
$\boldsymbol{p}_{Carl}$	1	0	1	0	1	1
$\boldsymbol{p}_{d_1}$	2	0	2	0	2	1
$\boldsymbol{p}_{d_2}$	0	1	0	1	1	2
$\boldsymbol{p}_{d_3}$	0	1	1	0	1	1
$\boldsymbol{q}$	0	1	1	0	1	0

the document; on the other hand, only a small portion of online documents are annotated by a user, but we need a personalized document profile for each document.

Therefore, we propose to estimate the personalized document profile of a user u via using the perception similarities between u and other users as weights to sum up tags assigned to the relevant document by the users having high perception similarities with u. The underlying intuition is that users having similar perceptions about existing documents will very likely also share similar perceptions about future documents, so, for a user u, tags assigned by users having high perception similarity with u are more helpful to characterize u's personal perception about the document than tags assigned by users having low perception similarity with u. Intuitively, the higher the perception similarity between two users, the higher their tags are weighted for each other.

In this section, we first propose a method to quantify the perception similarities between users. Then, we present how to use perception similarities as weights of tags to estimate the personalized document profile. Finally, Example 4 illustrates how to apply *D-PR* in the running example.

Profile-Based Perception Similarity. Since the tags assigned by a user to a document can be treated as an outline of this user's perception about this document, it is natural to measure a user's overall perception by the weighted vector based on all the tags used by this user, i.e., his/her user profile. Thus, a perception similarity of two users can be measured by the similarity of their profiles, called *profile-based perception similarity*:

$$PerSim(u', u) = Sim(\boldsymbol{p}_{\boldsymbol{u'}}, \boldsymbol{p}_{\boldsymbol{u}}). \tag{8}$$

Estimate of Personalized Document Profile. For a given user, after obtaining the perception similarities between u and all other users, we first select a set of users $U_T \subseteq U$, whose perception similarity with u are higher than a predefined threshold T. Then, for a given document d, we estimate u's personalized document profile relative to d (denoted $\boldsymbol{p}_{\boldsymbol{u},\boldsymbol{d}}$) by using perception similarities as weights to sum up the tags assigned to d by the users belonging to U_T. Formally,

$$\boldsymbol{p}_{\boldsymbol{u},\boldsymbol{d}} = \sum_{i=1}^{|U_d \cap U_T|} (\boldsymbol{v}_{\boldsymbol{u_i},\boldsymbol{d}} \cdot PerSim(u_i, u)), \tag{9}$$

Table 3. Weighted vectors of personalized document profiles and extended user profile

	English	Chinese	Comedy	Action	Interesting	Boring
$\boldsymbol{p_{Carl,d_1}}$	1.9	0	1.9	0	1.9	0.56
$\boldsymbol{p_{Carl,d_1}}$	0	0.56	0	0.56	0.56	1.9
$\boldsymbol{p_{Carl,d_3}}$	0	0.9	0.9	0	0.9	0.56
$\boldsymbol{p'_{Carl}}$	1.9	1.46	2.8	0.56	3.36	2.16

where $\boldsymbol{v_{u_i,d}}$ is also a weighted vector of tags, whose weight of a tag is the number of times that the tag is assigned by u_i to d; while $U_d \subseteq U$ is the set of users who annotate document d, and $|U_d \cap U_T|$ is the cardinality of the intersection of U_d and U_T.

Example 4. Continuing the running example, based on Equation 8, we first use $\boldsymbol{p_{Alice}}$, $\boldsymbol{p_{Bob}}$, and $\boldsymbol{p_{Carl}}$ as shown in Table 2 to compute the perception similarities between $Carl$ and two other users as follows:

$$\begin{aligned}
PerSim(Carl, Alice) &= Sim(\boldsymbol{p_{Carl}}, \boldsymbol{p_{Alice}}) = \frac{6}{\sqrt{4}\cdot\sqrt{11}} = 0.9,\\
PerSim(Carl, Bob) &= Sim(\boldsymbol{p_{Carl}}, \boldsymbol{p_{Bob}}) = \frac{3}{\sqrt{4}\cdot\sqrt{7}} = 0.56,\\
PerSim(Carl, Carl) &= Sim(\boldsymbol{p_{Carl}}, \boldsymbol{p_{Carl}}) = 1.
\end{aligned} \tag{10}$$

Then, based on Equation 9, we estimate $Carl$'s personalized document profile of d_1, d_2, and d_3 (denoted $\boldsymbol{p_{Carl,d_1}}$, $\boldsymbol{p_{Carl,d_2}}$, and $\boldsymbol{p_{Carl,d_3}}$, respectively) as shown in Table 3, where the threshold T is set to 0.5, so $U_T = U$. Consequently, we further use Equation 7 to obtain the extended profile of $Carl$ (denoted $\boldsymbol{p'_{Carl}}$) as shown in Table 3. Finally, the personalized ranking scores of d_1, d_2, and d_3 relative to $Carl$ based on the D-PR function (Equation 6) can be computed as shown in Equation 11 (with $\alpha = 0.5$ and $\beta = 0.5$), and the resulted ordering is $\tau_3 = [d_3 \geq d_1 \geq d_2]$.

$$\begin{aligned}
Rank(d_1, q, Carl) &= \frac{1}{2}(Sim(\boldsymbol{p'_{Carl}}, \boldsymbol{p_{Carl,d_1}}) + \frac{1}{2}(Sim(\boldsymbol{q}, \boldsymbol{p_{d_1}}) + Score(q, d_1)))\\
&= \frac{1}{2}(\frac{17.0052}{\sqrt{11.1436}\cdot\sqrt{34.3052}} + \frac{1}{2}(0.64 + 0.6))\\
&= \frac{1}{2}(0.87 + \frac{1}{2}(0.64 + 0.6)) = 0.75,\\
Rank(d_2, q, Carl) &= \frac{1}{2}(Sim(\boldsymbol{p'_{Carl}}, \boldsymbol{p_{Carl,d_2}}) + \frac{1}{2}(Sim(\boldsymbol{q}, \boldsymbol{p_{d_2}}) + Score(q, d_2)))\\
&= \frac{1}{2}(0.7 + \frac{1}{2}(0.44 + 0.52)) = 0.59,
\end{aligned} \tag{11}$$

$$Rank(d_3, q, Carl) = \frac{1}{2}(Sim(\boldsymbol{p'_{Carl}}, \boldsymbol{p_{Carl,d_3}}) + \frac{1}{2}(Sim(\boldsymbol{q}, \boldsymbol{p_{d_3}}) + Score(q, d_3)))$$
$$= \frac{1}{2}(0.88 + \frac{1}{2}(0.87 + 0.5)) = 0.78.$$

In summary, τ_3 ranked by *D-PR* is identical to the desired ordering τ_0. This is because *D-PR* solves profile modeling problems existing in the state-of-the-art approaches in the following two ways: (i) for a given user (e.g., $Carl$), *D-PR* utilizes the perception similarities to weaken the influences of tags assigned by users having different perceptions with this user (e.g., Bob) such that the resulting personalized document profiles can better capture this user's real perception about the documents; (ii) for a user u, *D-PR* obtains a personalized document profile for each document, so the extended user profile of u, computed by summing up all these personalized document profiles, contains more sufficient information to characterize u's preferences more comprehensively. □

5 Experimental Study

In this section, we evaluate the personalization performance of our *D-PR* function by comparing it with the *SoPRa* function, which is the closest work and considered as the state-of-the-art baseline. As this experiment aims at verifying the personalization effect of introducing two novel personalized profiles, we set $\beta = 1$ to eliminate the influence of the possible non-personalized textual matching problem in $Score(q, d)$.

We conduct experimental studies based on a public real-world large scale research dataset, which is described and analyzed in [15]. This dataset gathers more than $100\,000$ URLs of online documents and retrieves their social annotations from Delicious.com. After removing the documents without any social annotation, the general information of the resulting dataset is as shown in Table 4. Statistically, each user assigns an average of 9.4 tags; only 0.038% of users annotate more than 100 (0.17%) online documents; and the maximum number of online documents annotated by a single user is only 442 (0.75%). These statistical results show that, for any individual user, only a very small proportion of online documents are annotated by him/her, so we need to estimate the user's personalized document profile with the help of tags assigned by others, using their perception similarities as weights.

5.1 Evaluation Methodology

Although the relevance judgment of personalized search result subjectively depends on end users, several researches [1,2,12] have already proved that the tagging behavior of a user on the Social Web is closely correlated to his/her online search behavior, i.e., if a document is annotated by a user with some tags, this document is very likely to

Table 4. Dataset information

Users	Tags	Documents
388,963	3,647,266	59,126

be visited by the same user if it appears as a search result of using the same tags as the search query. This finding provides the theoretical base of our automatic evaluation framework: if a query is issued by a user with some terms, the relevant document is the one annotated by this user using the same terms as tags.

Therefore, to generate a set of synthetic user queries, we randomly select a set of bookmarks from the dataset. For each bookmark (u, t, d), we create a query $q = t$, which is issued by user u and aims at finding document d. In this paper, we limit the size of each query to be 2 to 4 keywords, which is a typical query size issued for online search as studied in work [8]. Finally, we remove all selected bookmarks to avoid promoting the annotated document with bias. Furthermore, to reduce the influence of removing bookmarks, we only randomly create 100 synthetic user queries each time and conduct 10 times of evaluations independently and then report the average results.

The performance of the *D-PR* function and the *SoPRa* function are evaluated based on a widely adopted metric [3,25], called *mean reciprocal rank (MRR)*. MRR measures the performance of a personalized function by assigning a value $1/r$ for each tested personalized query answering and then computing the mean value. Formally,

$$MRR = \sum_{i=1}^{n} 1/(r_i \cdot n), \tag{12}$$

where r_i is the ranking position of the i^{th} user query's relevant document in the personalized search result ordering, and n is the total number of tested queries.

5.2 Results

Since both the *D-PR* function and the *SoPRa* function use a parameter α to adjust the proportion of the profile matching score in the ranking score, we vary the value of the parameter α from 0 to 1.0 and report the result in each case. We set the threshold T of the perception similarity to 0.5. Recall from above that we set $\beta = 1$.

The experimental results of these two personalized functions are shown in Fig. 1. Generally, Fig. 1 shows that our *D-PR* function outperforms the *SoPRa* function in terms of MRR in almost all α cases, and its best ranking result at $\alpha = 0.3$ is about 11% better than the one of the *SoPRa* function at $\alpha = 0$. Specifically, we have the following observations in Fig. 1: (i) A continuous decline of the MRR of *SoPRa* is witnessed from $\alpha = 0$ (non-personalized) to $\alpha = 1$ (fully personalized); this observation verifies our argument that using user profiles and general document profiles cannot personalize well search results on the Social Web (here, they make ranking even worse). (ii) When the value of α rises from 0 to 0.3, the MRR of *D-PR* increases from around 0.147 to about 0.163. This indicates that the profile matching score in *D-PR* can better personalize the ranking of search results, which proves the effectiveness of the proposed personalized document profile and extended user profile. (iii) Afterwards, the MRR of *D-PR* continuously falls down to its bottom at $\alpha = 1$, which shows that excessive personalization will produce a bad ordering, because the topic matching between query and document is also critical for online search. Overall, these experimental results show that our *D-PR* function achieves better personalized search than the state-of-the-art *SoPRa* function.

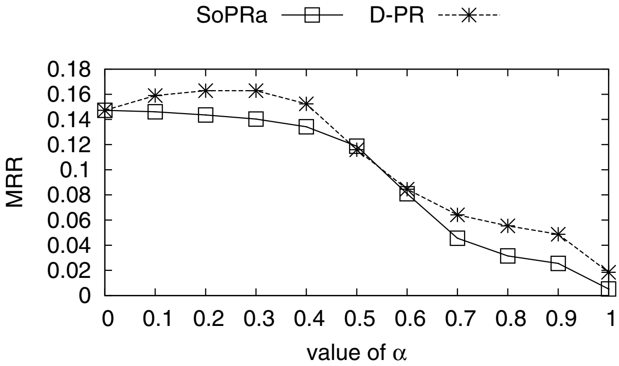

Fig. 1. *D-PR* vs *SoPRa*

6 Related Work

Personalized Web search by considering the searcher's personal attributes and preferences while evaluating a query is of great interest in information retrieval [18], since user queries are in general very short and provide an incomplete specification of the individual information need of a user. Some approaches have already been proposed to mine user preferences from both the user's explicit and implicit activities on the Web, such as query history [21], browsing history [22], the user's current task [13] or intent [23], and even eye-tracking during the search session [9]. Then, a user profile is built from the user's preferences and used for personalization by *query expansion* [5], i.e., a user's query is expanded based on the resulting profile to reflect the particular interest, or *re-ranking* [20], i.e., search results are re-ranked according to a user's profile such that personally relevant results appear higher in the search result list.

Specifically, the work in [9] shows that user preferences that are derived from click logs are reasonably reliable; so, in [16], the history of click data is used to estimate the user's hidden interests and to compute values of the *topic-sensitive PageRank* [7] for personalizing search results. Furthermore, Shen et al. [20] develop a method based on a decision-theoretic framework to convert user search histories into user profiles that are used to both expand queries and re-rank search results. As shown in [6], the benefits that can be achieved through personalization vary across queries; [13] and [23] thus propose solutions to discover the user's current tasks or intents by log analysis to help identify the queries that will benefit most from personalization.

However, mining user preferences by aggregating the user's online activities inevitably encounters a serious problem of privacy compromise [11]: due to the various online activities, Web logs usually contain some sensitive information of users, such as home address, medical record, bank account number, social security number, and so on. Therefore, as a privacy-enhanced personalization technique, folksonomy-based personalized Web search attracts more and more research efforts [3,14,25,26].

Xu et al. [26] propose to use the similarity of folksonomy-based user and document profiles to personalize search results. Then, Bouadjenek et al. [3] extend this work by introducing a social matching score to solve the textual matching problem. Instead of using *tf-idf*, Noll and Meinel [14] only use user tag frequency as the weighting of tag and normalize all document frequency to 1 to put more importance to the user profile. Vallet et al. [25] propose to use the probabilistic BM25 ranking model [17] to replace VSM. As these works weight tags from all users equally when modeling the document profile, they may encounter some personalization problems as discussed above.

In [24], Teevan et al. investigate how to use groups to improve personalized Web search and conclude that using group data collected across group members yields a significant improvement over individual personalization alone. Their work identifies the groups (or separate users) by either explicit properties (e.g., age, gender, job, location), interest groups, or desktop content; however, this information may result in privacy issues. Therefore, in our work, we propose to use the perception similarity computed from social annotations as groupization criteria to avoid such privacy problems.

7 Summary and Outlook

In this paper, we have proposed a dual personalized ranking (D-PR) function to improve personalized ranking of search on the Social Web via an extended user profile and a personalized document profile. We have formally defined the extended user profile of a user as the sum of all of his/her personalized document profiles; and we have further proposed to estimate the personalized document profile based on the perception similarities between users. Finally, a method used to quantify the perception similarity has also been presented. We have performed evaluations based on a public real-world large scale research dataset, and the results validate that our *D-PR* personalized function outperforms the state-of-the-art *SoPRa* function.

In future research, we will apply our *D-PR* ranking function to other Social Web datasets to evaluate its performance on various kinds of social resources. We will also investigate how to utilize categorical or ontology information of online documents to further enhance personalized search on the Social Web.

Acknowledgments. This work was supported by the EPSRC grant EP/J008346/1 “PrOQAW: Probabilistic Ontological Query Answering on the Web”, by the European Research Council (FP7/2007–2013/ERC) grant 246858 (“DIADEM”), by a Google European Doctoral Fellowship, and by a Yahoo! Research Fellowship. We also want to thank Mohamed Reda for providing us with the SoPRa implementation and dataset.

References

1. Benz, D., Hotho, A., Jäschke, R., Krause, B., Stumme, G.: Query logs as folksonomies. Datenbank-Spektrum 10(1), 15–24 (2010)
2. Bischoff, K., Firan, C.S., Nejdl, W., Paiu, R.: Can all tags be used for search? In: Proceedings of CIKM, pp. 193–202 (2008)
3. Bouadjenek, M.R., Hacid, H., Bouzeghoub, M.: Sopra: A new social personalized ranking function for improving Web search. In: Proceedings of SIGIR, pp. 861–864 (2013)
4. Carmel, D., Zwerdling, N., Guy, I., Ofek-Koifman, S., Har’el, N., Ronen, I., Uziel, E., Yogev, S., Chernov, S.: Personalized social search based on the user’s social network. In: Proceedings of CIKM, pp. 1227–1236 (2009)

5. Chirita, P.A., Firan, C.S., Nejdl, W.: Personalized query expansion for the Web. In: Proceedings of SIGIR, pp. 7–14 (2007)
6. Dou, Z., Song, R., Wen, J.: A large-scale evaluation and analysis of personalized search strategies. In: Proceedings of WWW, pp. 581–590 (2007)
7. Haveliwala, T.H.: Topic-sensitive pagerank. In: Proceedings of WWW, pp. 517–526 (2002)
8. Jansen, B.J., Spink, A., Bateman, J., Saracevic, T.: Real life information retrieval: A study of user queries on the Web. In: SIGIR Forum, pp. 5–17 (1998)
9. Joachims, T., Granka, L., Pan, B., Hembrooke, H., Gay, G.: Accurately interpreting clickthrough data as implicit feedback. In: Proceedings of SIGIR, pp. 154–161 (2005)
10. Jones, K.S.: A statistical interpretation of term specificity and its application in retrieval. Journal of Documentation 28, 11–21 (1972)
11. Kobsa, A.: Privacy-enhanced personalization. Commun. ACM 50(8), 24–33 (2007)
12. Krause, B., Hotho, A., Stumme, G.: A comparison of social bookmarking with traditional search. In: Macdonald, C., Ounis, I., Plachouras, V., Ruthven, I., White, R.W. (eds.) ECIR 2008. LNCS, vol. 4956, pp. 101–113. Springer, Heidelberg (2008)
13. Luxenburger, J., Elbassuoni, S., Weikum, G.: Task-aware search personalization. In: Proceedings of SIGIR, pp. 721–722 (2008)
14. Noll, M.G., Meinel, C.: Web search personalization via social bookmarking and tagging. In: Aberer, K., Choi, K.-S., Noy, N., Allemang, D., Lee, K.-I., Nixon, L.J.B., Golbeck, J., Mika, P., Maynard, D., Mizoguchi, R., Schreiber, G., Cudré-Mauroux, P. (eds.) ISWC/ASWC 2007. LNCS, vol. 4825, pp. 367–380. Springer, Heidelberg (2007)
15. Noll, M.G., Meinel, C.: The metadata triumvirate: Social annotations, anchor texts and search queries. In: Proceedings of WI-IAT, pp. 640–647 (2008)
16. Qiu, F., Cho, J.: Automatic identification of user interest for personalized search. In: Proceedings of WWW, pp. 727–736 (2006)
17. Robertson, S.E., Walker, S.: Some simple effective approximations to the 2-Poisson model for probabilistic weighted retrieval. In: Proceedings of SIGIR, pp. 232–241 (1994)
18. Salton, G., McGill, M.J.: Introduction to Modern Information Retrieval. McGraw-Hill, Inc., New York (1986)
19. Salton, G., Wong, A., Yang, C.S.: A vector space model for automatic indexing. Commun. ACM 18(11), 613–620 (1975)
20. Shen, X., Tan, B., Zhai, C.: Implicit user modeling for personalized search. In: Proceedings of CIKM, pp. 824–831 (2005)
21. Shen, X., Zhai, C.X.: Exploiting query history for document ranking in interactive information retrieval. In: Proceedings of SIGIR, pp. 377–378 (2003)
22. Sugiyama, K., Hatano, K., Yoshikawa, M.: Adaptive Web search based on user profile constructed without any effort from users. In: Proceedings of WWW, pp. 675–684 (2004)
23. Teevan, J., Dumais, S.T., Liebling, D.J.: To personalize or not to personalize: modeling queries with variation in user intent. In: Proceedings of SIGIR, pp. 163–170 (2008)
24. Teevan, J., Morris, M.R., Bush, S.: Discovering and using groups to improve personalized search. In: Proceedings of WSDM, pp. 15–24. ACM (2009)
25. Vallet, D., Cantador, I., Jose, J.M.: Personalizing Web search with folksonomy-based user and document profiles. In: Gurrin, C., He, Y., Kazai, G., Kruschwitz, U., Little, S., Roelleke, T., Rüger, S., van Rijsbergen, K. (eds.) ECIR 2010. LNCS, vol. 5993, pp. 420–431. Springer, Heidelberg (2010)
26. Xu, S., Bao, S., Fei, B., Su, Z., Yu, Y.: Exploring folksonomy for personalized search. In: Proceedings of SIGIR, pp. 155–162 (2008)