Personalized Social Query Expansion Using Social Annotations

Mohamed Reda Bouadjenek[1], Hakim Hacid[2(✉)], and Mokrane Bouzeghoub[3]

[1] Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, Canada
mrb@mie.utoronto.ca

[2] Zayed University, Dubai, United Arab Emirates
hakim.hacid@zu.ac.ae

[3] University of Versailles, Versailles, France
mokrane.bouzeghoub@uvsq.fr

Abstract. Query expansion is a query pre-processing technique that adds to a given query, terms that are likely to occur in relevant documents in order to improve information retrieval accuracy. A key problem to solve is "*how to identify the terms to be added to a query?*" While considering social tagging systems as a data source, we propose an approach that selects terms based on (i) the semantic similarity between tags composing a query, (ii) a social proximity between the query and the user for a personalized expansion, and (iii) a strategy for expanding, on the fly, user queries. We demonstrate the effectiveness of our approach by an intensive evaluation on three large public datasets crawled from delicious, Flickr, and CiteULike. We show that the expanded queries built by our method provide more accurate results as compared to the initial queries, by increasing the MAP in a range of 10 to 16% on the three datasets. We also compare our method to three state of the art baselines, and we show that our query expansion method allows significant improvement in the MAP, with a boost in a range between 5 to 18%.

Keywords: Personalization · Social Information Retrieval
Social networks · Query expansion

CR Subject Classification: H.3.3 [Information Systems]: Information Storage and Retrieval · Information Search and Retrieval

1 Introduction

Web 2.0 has strengthened end-users position in the Web through their integration in the heart of the content generation ecosystem. This has been made possible mainly through the availability of tools such as social networks, social bookmarking systems, social news sites, etc., impacting the way information is produced, processed, and consumed by both humans and machines. As a result,

A. Hameurlain et al. (Eds.): TLDKS XL, LNCS 11360, pp. 1–25, 2019.
https://doi.org/10.1007/978-3-662-58664-8_1

on the one hand, the user is no longer able to digest the large quantity of information he has access to and is generally overwhelmed by it. On the other hand, most of popular Information Retrieval (IR) systems lack in offering efficient personalization techniques, which provide users only with the necessary information that fulfill their needs. Two types of constraints make the situation more complex: *information-dependent* constraints and *user-dependent* constraints. The first class of constraints includes (i) the large scale due to the continuous activities of users and their ability to generate new content, (ii) information diversity or heterogeneity, since different types of media are used to communicate, e.g., text, image, video, etc. (iii) versatility, since information is dynamic and is continuously updated (confirmed, contradicted, etc.), (iv) its disparity, since it can be in different places, and as a result (v) the variation in the quality of information. The second class of constraints is mainly related to users' diversity and the high dynamics in their profiles.

To improve the IR process and reduce the amount of irrelevant documents, there are mainly three possible improvement tracks: (i) query reformulation using extra knowledge, i.e., expansion or refinement of the user query, (ii) post filtering or re-ranking of the retrieved documents (based on the user profile or context), and (iii) improvement of the IR model, i.e., reengineering of the IR process to integrate contextual information and relevant ranking functions. In this paper, we focus on query reformulation, especially on personalized query expansion for personalized search, i.e., personalizing the reformulation of queries.

Query expansion consists of enriching the user's initial query with additional information so that the IR system may propose suitable results that better satisfy user's needs [14,15,19]. We explore the possibility of using the data available in social networks, and more precisely data of social bookmarking systems, as a source of explicit feedback information. These latter enable users to freely add, annotate, edit, and share bookmarks of web resources, e.g., web pages. Basically, we propose an approach which reuses the users vocabulary (the terms used to annotate web pages) in order to expand their queries in a personalized way and thus, increase their satisfaction regarding the quality of search. Exploiting social knowledge for improving web search has a number of advantages:

- Feedback information in social networks is provided directly by the user, so users interests accurate information can be harvested as people actively express their opinions on social platforms. Thus, this user interest can be easily modeled to provide personalized services.
- A huge amount of social information is published and available with the agreement of the publishers. Exploiting these information should not violate user privacy, in particular social tagging information, which doesn't contain sensitive information about users.
- Finally, social resources are often publicly accessible, as most of social networks provide APIs to access their data (even if often, a contract must be established before any use of the data).

Our approach in this work[1] consists of three main steps: (i) determining similar and related tags to a given query term through their co-occurrence over resources and users, (ii) constructing a profile of the query issuer based on his tagging activities, which is maintained and used to compute expansions, and finally, (iii) expanding the query terms, where each term is enriched with the most interesting tags based on their similarities and their interest to the user.

The problem we are tackling in this paper is strongly related to personalization since we want to expand queries in a personalized way and consequently propose adapted search results. Personalization allows to differentiate between individuals by emphasizing on their specific domains of interest and their preferences. It is a key point in IR and its demand is constantly increasing by various users for adapting their results [3]. Several techniques exist to provide personalized services among which the user profiling. The user profile is a collection of personal information associated to a specific user that enables to capture his interests. Details of how we model user profiles are given in Sects. 2 and 3.1.4.

The main contributions of this work can be summarized as follows:

1. We propose an approach in which we use social knowledge as explicit feedback information for the expansion process. Reusing such a social knowledge aims at expanding user queries with their own vocabularies instead of using a public thesaurus, which is made by people who are not aware of the individual users needs and expectations.
2. We propose a Personalized Social Query Expansion framework called PSQE. This latter provides a user-dependent query expansion based on social knowledge, i.e., for the same query of two different users, PSQE will provide two different expanded queries, which will be processed by a search engine.
3. Using an evaluation on real data gathered from three different large bookmarking systems, we demonstrate the effectiveness of our framework for socially driven query expansion compared to many state of the art approaches.

The rest of this paper is organized as follows: in Sect. 2 we introduce all the concepts that we use throughout this paper. Section 3 introduces our method of query expansion using folksonomy. In Sect. 4, we discuss the different experiments that evaluate the performance of our approach. Related work is discussed in Sect. 5. Finally, we conclude and provide some future directions in Sect. 6.

2 Background and Notations

In this section, we formally define the basic concepts that we use throughout this paper namely, a bookmarks, a folksonomy, and a user profile. We also provide a formal definition of the problem we are intending to solve.

[1] This is an extended and revised version of a preliminary conference report that was presented in [12].

2.1 Background

Social bookmarking websites are based on the techniques of *social tagging or collaborative tagging*. The principle behind social bookmarking platforms is to provide the user with a means to annotate resources on the Web, e.g., URIs in *delicious*[2], videos in *youtube*[3], images in *flickr*[4], or academic papers in *CiteULike*[5]. These annotations (also called tags) can be shared with others. This unstructured (or better, free structured) approach to classification with users assigning their own labels is variously referred to as a *folksonomy* [21,28]. A folksonomy is based on the notion of bookmark, which is formally defined as follows:

Definition 1 (Bookmark). Let U, T, R be respectively the set of Users, Tags and Resources. A *bookmark* is a triplet (u, t, r) such as $u \in U$, $t \in T$, $r \in R$, which represents a user u who used a tag t to annotate a resource r.

Then, a group of bookmarks which forms a folksonomy is formally defined as follows:

Definition 2 (Folksonomy). Let U, T, R be respectively the set of Users, Tags and Resources. A folksonomy $\mathbb{F}(U, T, R)$ is a subset of the cartesian product $U \times T \times R$ such that each triple $(u, t, r) \in \mathbb{F}$ is a *bookmark*.

A folksonomy can then be naturally represented by a tripartite-graph where each ternary edge represents a bookmark. In particular, the graph representation of the folksonomy $\mathbb{F}$ is defined as a tripartite graph $\mathcal{G}(V, E)$ where $V = U \cup T \cup R$ and $E = \{(u, t, r)|(u, t, r) \in \mathbb{F}\}$. Figure 1 shows seven bookmarks provided by two users on three resources using three tags.

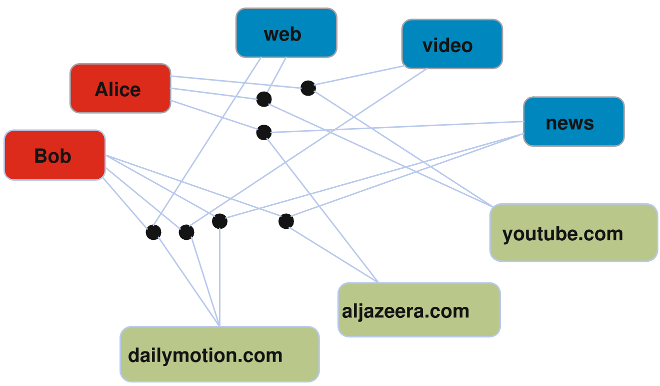

Fig. 1. Example of a folksonomy. The triples (u, t, r) are represented as ternary-edges connecting users, resources and tags.

[2] http://www.delicious.com/.
[3] http://www.youtube.com/.
[4] http://www.flickr.com/.
[5] http://www.citeulike.org/.

Folksonomies have proven to be a valuable knowledge for user profiling [17, 35, 41, 43]. Especially, because users tag interesting and relevant information to them with keywords that may be a good summary of their interest. Hence, in this paper, and in the context of folksonomies, the profile includes all the terms used as tags along with their weights to capture user's tagging activities. It is formally defined as follows:

Definition 3 (User Profile). Let U, T, R be respectively the set of Users, Tags and Resources of a folksonomy $\mathbb{F}(U, T, R)$. A profile p_u assigned to a user $u \in U$, is modeled as a weighted vector $\mathbf{p_u}$ of m dimensions, where each dimension represents a tag the user employed in his tagging actions. More formally, $\mathbf{p_u} = \{w_{t_1}, w_{t_2}, ..., w_{t_m}\}$ such that w_{t_m} is the weight of t_m, such as $t_m \in T \wedge (\exists r \in R \mid (u, t_m, r) \in \mathbb{F})$.

Thus, the profile includes the most relevant terms for the user and not all his activities, i.e., the documents that he has tagged. A value is associated to each term of the profile expressing its strength and importance for the given user.

Later in Sect. 3.1.4, we propose a method to assign weights to each term in the user profile in order to better define his interests.

2.2 Problem Definition

As mentioned before, query expansion consists of enriching the initial query with additional information. This expansion is generally expected to provide better search results. However, providing merely a uniform expansion to all users is, from our point of view, not really suitable nor efficient since relevance of documents is relative for each user. Thus, a simple and uniform query expansion is not enough to provide satisfactory search results for each user. Hence, having a folksonomy $\mathbb{F}(U, T, R)$, the problem we are addressing can be formalized as follows:

For a given user $u \in U$ who issued a query $q = \{t_1, t_2, ..., t_n\}$, how to provide for each term $t_i \in q$ a ranked list of related terms $\mathcal{L} = \{t_{i1}, t_{i2}, ..., t_{ik}\}$, such that when expanding the term t_i with the top k of $\mathcal{L}$, the most relevant documents are put earlier in the ranking?

3 Social Query Expansion Approach

The approach we are proposing aims at expanding user's queries in a personalized way. It can be decomposed into two parts: (i) *an offline* and (ii) *an online* part. The offline part performs the heavy computation which consists of transforming the whole social graph of a folksonomy $\mathbb{F}$ into a graph of tags where two tags are related if they are semantically related. This part is also responsible for the construction and the update of the users' profiles, for serving the online part. The online part of the approach is responsible for computing the concrete expansion using the graph of tags and the user' profiles constructed in the offline part. In the following, we describe in more details each part and we explicitly highlight our contributions.

3.1 Offline Part

The offline part is also decomposed into two facets: (i) the transformation of the social graph of a folksonomy $\mathbb{F}$ into a graph of tags, representing similarities between tags that either occur on the same resources or are shared by the same users, and (ii) the computation of the users' profiles to highlight their interests for personalizing their queries.

The approach is based on the creation and the maintenance of a graph of tags that represents all the similarities that exist between the tags of $\mathbb{F}$. There exist two kinds of approaches that propose to achieve that: (i) an approach based on the co-occurrence of tags over resources, and (ii) an approach based on their co-occurrence over users.

3.1.1 Extracting Semantics from Resources

In the first category of approaches, [24,30,33] state that semantically related tags are expected to occur over the same resources. For example, tags that most occur for *google.com* on *delicious* are: *search, google, engine, web, internet*.

Thus, extracting semantically related tags can be carried out by computing similarities. There exist many similarity measures [30], but all of them need preprocessing that consists of reducing the dimensionality of the tripartite graph $\mathbb{F}$ into a bipartite graph. This reduction is generally performed through aggregation methods. From the study of existing aggregation methods proposed in [30], we have chosen the *projectional aggregation* along with the *Jaccard*, the *Dice*, and the *Overlap* similarity measures to compute the similarity between tags. We choose this aggregation method because its simplicity, and it is one which gives better results in semantic information extraction [30]. Hence, we follow the same process as [30] to extract a graph of related tags from $\mathbb{F}$ according to their co-occurrence over resources:

1. Using a function $\mathcal{F}$ on the whole *folksonomy* $\mathbb{F}$ performs a *projectional aggregation* over the user dimension, resulting in a bipartite graph Tag-Resource.
2. Then, using a function $\mathcal{G}$ on the resulting bipartite graph Tag-Resource provides a graph of tags $\mathcal{T}_R$, in which each link is weighted with the similarity between tags according using the Jaccard, the Dice or the Overlap metrics [30].

Therefore, we may obtain either a graph of tags $\mathcal{T}_R$ using the *Jaccard*, the *Dice*, or the *Overlap*. Note that we do not merge the similarity measures in a same graph of tags, meaning that a graph of tags is constructed using only one similarity measure.

We end-up with an undirected weighted graph in which nodes represent tags, and an edge between two tags represents the fact that these tags occur together at least on one resource. The weights associated to edges are computed from similarities between tags as explained beforehand. This first step is illustrated in the left upper part of Fig. 2.

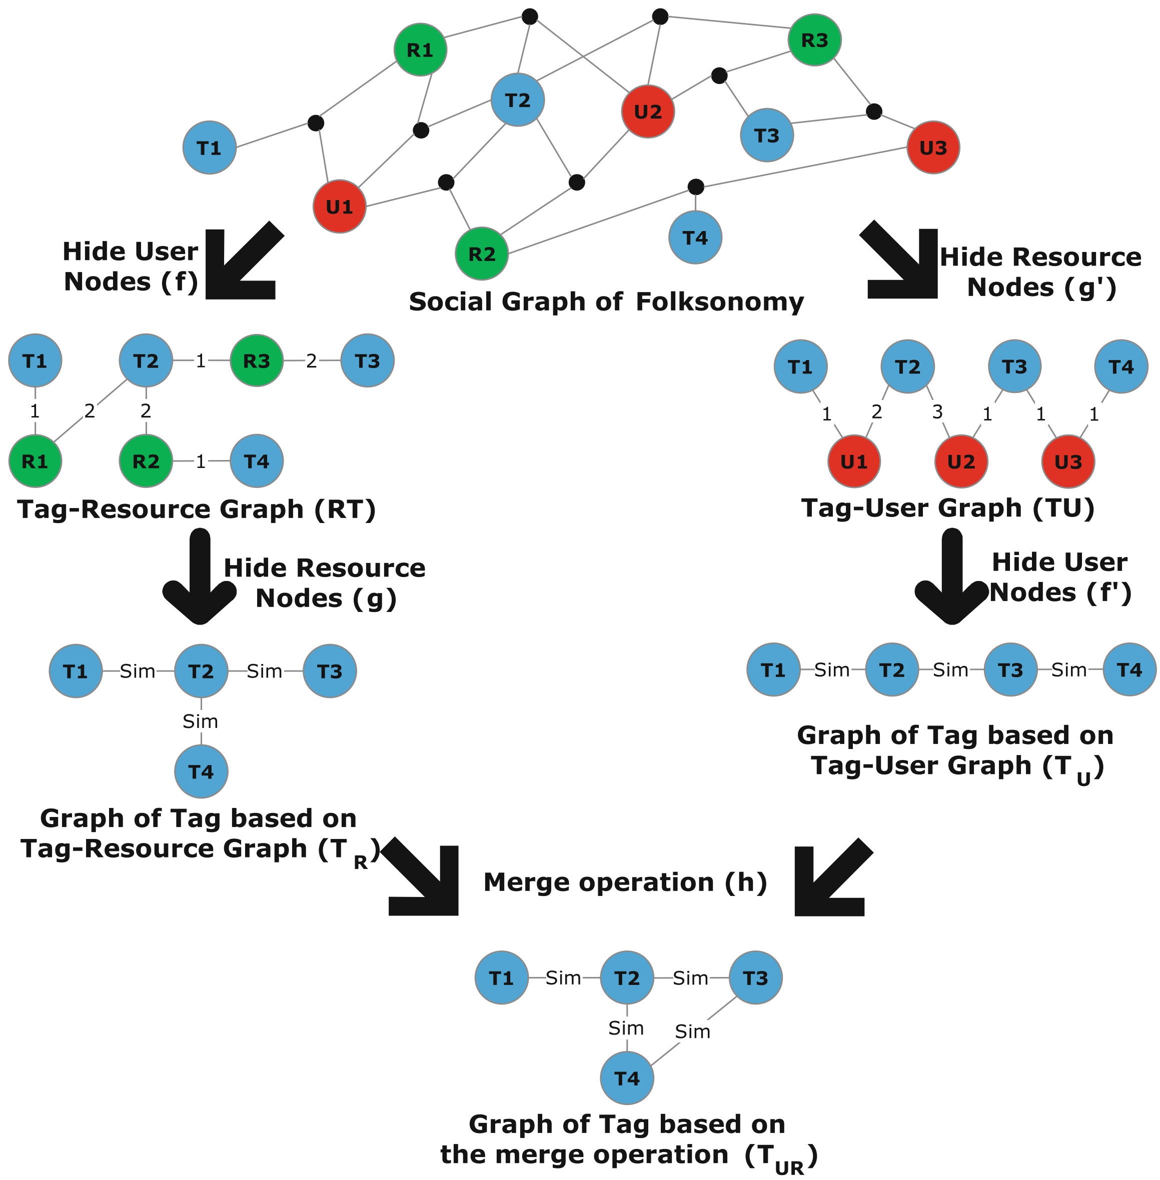

Fig. 2. Summary of the graph reduction process, which transform the whole folksonomy $\mathbb{F}$ into a graph of tags $\mathcal{T}_{\mathcal{UR}}$. The similarity values on the Figure are computed using the Jaccard measure on both graphs $\mathcal{T}_{\mathcal{R}}$ and $\mathcal{T}_{\mathcal{U}}$, and using $\alpha = 0.5$ on the graph $\mathcal{T}_{\mathcal{UR}}$.

3.1.2 Extracting Semantics from Users

In the second category, [4,33] state that correlated tags are also used by the same users to annotate resources. For example, the tags *Collaborative* and *Blog* have been used 13,557 times together by users in our *delicious* dataset.

This observation is more expected to happen in certain folksonomies, where users are encouraged to upload their personal resources which leads to generate private bookmarks, e.g., a folksonomy such as *CiteULike*, *Flickr*, or *YouTube* where users are expected to upload respectively their research papers, images, and videos. Therefore, similarly to the previous approach, [33] proposes to extract semantically related tags using the following process:

1. Using a function $\mathcal{G}'$ on the folksonomy $\mathbb{F}$ performs a *projectional aggregation* over the resource dimension for obtaining a bipartite graph Tag-User.
2. Then the function $\mathcal{F}'$ is used to get another graph of tags $\mathcal{T}_{\mathcal{U}}$ where similarities between tags are computed using one of the three previous similarity measures.

This process is illustrated in the right upper part of Fig. 2. Notice that the structure of the graph of tags $\mathcal{T}_R$ is different from the one of the graph of tags $\mathcal{T}_{\mathcal{U}}$.

3.1.3 Construction of the Graph of Tag Similarities

Using only one of the two previous methods to construct a graph representing similarities between tags leads to a loss of information on one side or the other. For example, if we choose to extract related tags according to their co-occurrence over resources, we neglect the fact that there are some tags which are expected to be shared by the same users and vice versa.

Therefore, we propose to use a function $\mathcal{M}$ which is applied on the graphs of tags $\mathcal{T}_R$ and $\mathcal{T}_{\mathcal{U}}$ to merge them and to get a unique graph of tags $\mathcal{T}_{\mathcal{UR}}$ where the new similarity values are computed by merging the values using the *Weighted Borda Fuse (WBF)* [18]. This merge is summarized in Eq. 1, where $0 \leq \alpha \leq 1$:

$$Sim_{\mathcal{T}_{\mathcal{UR}}}(t_i, t_j) = \alpha \times Sim_{\mathcal{T}_R}(t_i, t_j) + (1 - \alpha) \times Sim_{\mathcal{T}_{\mathcal{U}}}(t_i, t_j) \qquad (1)$$

Where, $Sim_{\mathcal{T}_{\mathcal{UR}}}(t_i, t_j)$ calculates the similarity between two tags relying on the two other types of nodes, i.e., users and resources. The parameter α represents the importance one wants to give to the two types of graphs, i.e., resources or users, in the consideration of the similarity calculation. In fact, depending on the context, when computing the similarity between two tags, one may want to give a higher importance to users sharing these two tags than documents having these tags as a common tags. Another user may want to give more importance to their co-occurrence over resources than to the users sharing these tags. Depending on the nature of the folksonomy, we set α to its optimal value in order to maximize the tags semantics extraction. Finally, it should be noted that the merge is performed between graphs generated with the same similarity measure.

This step of the offline part extracts semantics from the whole social graph of $\mathbb{F}$ without a loss of information, i.e., by exploiting the co-occurrences of tags over resources and users. This step leads to the creation of a graph of tags, where edges represent semantic relations between tags. This graph will be further used to extract terms that are semantically related to a given term of a query to perform the query expansion. The contribution at this stage is the combination of the graphs resulting from resources and users to construct a better graph of tag similarities without loss of information. This is different from the existing approaches where only one graph is used.

In the following, we introduce our method of constructing and weighting the user profiles in order to personalize the expansions.

3.1.4 Construction of the User Profile

To achieve a personalized expansion, we also propose to build a user profile that consists of capturing information regarding real user interests. There are different ways to build user profiles [23,40,41]. For example, a person may be modeled as a vector of attributes of his online personal profiles including the name, affiliation, and interests. Such simple factual data provides an inadequate description of the individual, as they are often incomplete, mostly subjective and do not reflect dynamic changes [23].

Since we focus on folksonomies, the user feedback is expected to be mostly explicit (because of the tagging action, where the user explicitly assigns tags to resources).

Thus, in a folksonomy, users are expected to tag and annotate resources that are interesting to them using tags that summarize their understanding of resources. In other words, these tags are in turn expected to be a good summary of the user's topics of interests as also discussed in [2,17,23,35,37,43]. Hence, each user can be modeled as a set of tags and their weights.

The definition of a user profile is given in Definition 3. The main challenge here is *how to define the weight of each tag in the user profile?* We propose to use an adaptation of the well known *tf-idf* measure to estimate this weight. Hence, we define the weight w_{t_i} of the term t_i in the user profile as the *user term frequency, inverse user frequency (utf-iuf)*, which is computed as follows:

$$utf-iuf_{t_i,u_j} = \frac{n_{t_i,u_j}}{\sum\limits_{t_k \in \mathbf{p_u^m}} n_{t_k,u_j}} \times log\left(\frac{|U|}{|U_{t_i}|}\right) \tag{2}$$

where n_{t_i,u_j} is the number of time the user u_j used the tag t_i.

A high value of *utf-iuf* is reached by a high user term frequency and a low user frequency of the term in the whole set of users. Note that we perform a stemming on tags before computing the profiles, to eliminate the differences between terms having the same root to better estimate the weight of each term.

User profiles are created offline and maintained incrementally. This is motivated by the fact that profiles and tagging actions are not evolving as quickly as query formulation on the system. As an analogy, it is well known that 90% of users in the social Web consume the content (i.e., query formulation), 9% update content, and 1% generate new content (profile updates) [34]. Thus, we have decided to handle the profile construction as an offline task while providing a maintenance process for keeping it up to date.

In summary, at the end of the offline part, we build two assets: (i) a graph of tags similarities which is used to represent semantically relatedness of terms, and (ii) user profiles which are leveraged in the personalization step.

3.2 Online Part

The online part of the approach is responsible for computing the concrete expansion using the graph $\mathcal{T}_{\mathcal{UR}}$ and the profiles constructed in the offline part. Before presenting our algorithm of query expansion, we propose a method to compute, on the fly, the interest of a user to a given tag.

3.2.1 Interest Measure to Tag

Having computed the similarity graph between tags and built users' profiles containing the degree to which a set of tags are representative of a user, it becomes possible to compute a degree of interest a user may have to other tags, e.g., query tags. This is useful in our approach to compute, in real time, the suitable expansions of a tag w.r.t. a given user. In our approach, this interest is seen as a similarity between the user profile $\mathbf{p_u}$ and a tag t_i. Intuitively, the computed similarity captures the interest of the user u in the query term t_i denoted $\mathcal{I}^u_{t_i}$:

$$\mathcal{I}_u(t_i) = \sum_{t_j \in \mathbf{p_u}} (Sim_{\mathcal{T}_{\mathcal{UR}}}(t_i, t_j) \times w_j) \tag{3}$$

where $Sim(t_i, t_j)$ is the similarity between the term t_i and t_j, the j^{th} term of the user profile, and w_j is the weight of the term t_j in the profile computed during the previous process. Notice that any similarity measure can be used for computing $Sim(t_i, t_j)$, as discussed in [30]. In this work, we consider the *Jaccard*, the *Overlap*, and the *Dice* similarity measures, as discussed in the previous sections.

3.2.2 Effective Query Expansion

In this step of query expansion, we consider that the similarity between two terms t_i (a query term) and t_j (a potential candidate for the expansion of t_i), to be influenced by two main features: (i) the semantic similarity between t_i and t_j (the semantic strength between the two terms), and (ii) the extent to which the tag t_j is likely to be interesting to the considered user.

Once these two similarities are computed, a merge operation is necessary to obtain a final ranking value that indicates the similarity of t_j with t_i w.r.t. the user u. For this, several aggregation methods and algorithms exist. We choose the *Weighted Borda Fuse (WBF)* as summarized in Eq. 4, where $0 \leq \gamma \leq 1$ is a parameter that controls the strength of the semantic and social parts of our approach. Using Eq. 4, we can rank a list of terms $\mathcal{L}$, which are semantically related to a given term t_i from a user perspective.

$$Rank^u_t(t_j) = \overbrace{\gamma \times Sim_{\mathcal{T}_{\mathcal{UR}}}(t, t_j)}^{\texttt{Semantic Part}} + \underbrace{(1-\gamma) \times \mathcal{I}^u_{t_j}}_{\texttt{Social Part}} \tag{4}$$

The effective social query expansion is summarized in Algorithm 1. Hence, for a query $q = t_1 \wedge t_2 \wedge ... \wedge t_m$ issued by a user u, we first get the user's profile, which is computed as explained above (Sect. 3.1.4 and Line 1 in Algorithm 1). At this stage, the purpose is to enrich each term t_i of q with related terms (line 2). Then, the objective is to get all the neighboring tags t_j of t_i in the tag graph $\mathcal{T}_{\mathcal{UR}}$ (line 3). After that (in line 4), we compute for each t_j, the ranking value that indicates its similarity with t_i w.r.t. the user u using formula 4 (line 5). Next, the neighbor list has to be sorted according to the computed values and we keep only the k top tags (line 7). Finally, t_i and its remaining neighbors must be linked with the OR ($\vee$) logical connector (line 8) and updated in q'.

Algorithm 1. Effective Social Query Expansion

Require: A folksonomy $\mathbb{F}$
u : a User. $q = \{t_1, t_2, ..., t_n\}$: a Query.

1: $p_u[m] \leftarrow$ extract profile of u from $\mathbb{F}$
2: **for all** $t_i \in q$ **do**
3: $\mathcal{L} \leftarrow$ list of neighbor of t_i in tag graph $\mathcal{T}_{\mathcal{UR}}$
4: **for all** $t_j \in l$ **do**
5: $t_j.Value \leftarrow$ Compute the ranking score $Rank^u_{t_i}(t_j)$
6: **end for**
7: Sort $\mathcal{L}$ according to $t_j.Value$ and keep only the top k terms in $\mathcal{L}$
8: Make a logical OR ($\vee$) connection between t_i and all terms of $\mathcal{L}$
9: Set the weight of the new terms t_j as the $t_j.Value$ or the TF-IDF value, depending on the choosed strategy (See Section 3.2.3)
10: Insert $\mathcal{L}$ in q'
11: **end for**
12: **return** q'

Example 1. If a user issues a query $q = t_1 \wedge t_2 \wedge ... \wedge t_m$, it will be expanded to $q' = \{(t_1 \vee t_{11} \vee ... \vee t_{1l}) \wedge (t_2 \vee t_{21} \vee ... \vee t_{2k}) \wedge ... \wedge (t_m \vee t_{m1} \vee ... \vee t_{mr})\}$, where t_{ij} is a term that is semantically related to $t_i \in q$ and socially to u.

It should be noted that in this paper, we consider that the selection of each query term is determined independently, without considering latent term relations. Most past work on modeling term dependencies has analyzed three different underlying dependency assumptions: full independence, sequential dependence [39], and full dependence [32]. Taking into account terms dependency is part of our future works.

3.2.3 Terms Weighting

Term weighting in query expansion is challenging since there is no formal method for assigning weights to new terms. Indeed, appropriately weighting terms should result in better retrieval performance. Thus, we experiment the following two strategies for weighting new terms:

– Using the ranking values of Formula 4 as the weight of the new expanded terms. This strategy provides personalized term weight assignment while considering both semantic strength and user interest.

– Using the *Term Frequency-Inverse Document Frequency (TF-IDF)* [1] as the weight of the new expanded terms as follows:

$$tf - idf_{t_i,q} = tf_{t_i} \times log\left(\frac{|D|}{|D_{t_i}|}\right) \quad (5)$$

where tf_{t_i} denotes the term frequency of t_i in the query q. This strategy provides a uniform term weight to the query while keeping the personalizing aspect in choosing terms. Notice that weights are assigned to terms in the line 9 of Algorithm 1.

4 Evaluations

In this section, we describe the two types of evaluations we performed on our approach: (i) an estimation of the parameters of our approach to provide insights regarding their potential impact on the system, and (ii) a comparison study, where our approach is compared to the closest state of the art approaches to provide insights about the obtained results and position the proposal.

4.1 Datasets

A number of social bookmarking systems exist [21]. We have selected three datasets to perform an offline evaluation: *delicious, flickr and CiteULike*. These datasets are available and public. The interest of using such data instead of crawled data is to work on widely accepted data sets, reduce the risk of noise, and an ability to reproduce the evaluations by others as well as the ability to compare our approach to other approaches on "standardized datasets". Hereafter is the description of the different datasets.

– **Delicious:** a social bookmarking web service for storing, sharing, and discovering web bookmarks. We have used a dataset which is described and analyzed in [42][6].
– **Flickr:** an image hosting, tagging and sharing website. The *Flickr* dataset is the one used and studied in [38][7].
– **CiteULike:** an online bookmarking service that allows users to bookmark academic articles. This dataset is the one provided by the *CiteULike* website[8].

Before the experiments, we performed three data preprocessing tasks: (1) Several annotations are too personal or meaningless, such as "toread", "Imported IE Fa-vorites", "system:imported", etc. We remove some of them manually. (2) Although the annotations from delicious are easy for users to read and understand, they are not designed for machine use. For example, some users may concatenate several words to form an annotation such as "java.programming"

[6] http://data.dai-labor.de/corpus/delicious/.
[7] http://www.tagora-project.eu/data/#flickrphotos.
[8] http://static.citeulike.org/data/2007-05-30.bz2.

or “java/programming”. We split this kind of annotations before using them in the experiments. (3) The list of terms undergoes a stemming by means of the Porter's algorithm [36] in such a way to eliminate the differences between terms having the same root. In the same time, the system records the relations between stemmed terms and original terms. As for the *delicious* dataset, we add two other data preprocessing tasks: (i) we downloaded all the available web pages while removing those which are no longer available, and (ii) we removed all the non-english web pages. This operation was performed using *Apache Tika* toolkit. Table 1 gives a description of these datasets.

Table 1. Corpus details

	Bookmarks	Users	Resources	Tags
Delicious	9,675,294	318,769	425,183	1,321,039
Flickr	22,140,211	112,033	327,188	912,102
CiteULike	16,164,802	107,066	3,508,847	712,912

4.2 Evaluation Methodology

Making evaluations for personalized search is a challenge in itself since relevance judgements can only be assessed by end-users themselves [17]. This is difficult to achieve at a large scale. Different contributions [5,8,25,31] state that the tagging behavior of a user of folksonomies closely reflects his behavior of search on the Web. In other words, if a user u tags a resource r with a tag t, he will choose to access the resource r if it appears in the result obtained by submitting t as a query to the search engine. Thus, we can easily state that any bookmark (u, t, r) can be used as a test query for evaluations. The main idea of the experiments is based on the following assumption:

Proposition 1. *For a personalized query $q = \{t\}$ issued by a user u with a query term t, the relevant documents are those tagged by u with t.*

Hence, in the off-line study, for each evaluation, we randomly select $2,000$ pairs (u, t), which are considered to form a personalized query set. For each corresponding pair (u, t), we remove all the bookmarks $(u, t, r) \in \mathbb{F}, \forall r \in R$ in order to not promote the resource r in the obtained results. For each pair, the user u sends the query $q = \{t\}$ to the system. Then, the query q is enriched and transformed into q' following our approach. For the *delicious* dataset, documents that match q' are retrieved, ranked and sorted using the *Apache Lucene*. For the *Flickr* and *CiteULike* datasets, we retrieve all resources that are annotated with tags of q' while representing them according to the *Vector Space Model (VSM)*. Then, the cosine similarity is used to compute similarity between a query q' and a resource r_j.

For the *Flickr* and *CiteULike* datasets, we rank all the retrieved resources using values of the cosine similarity and we consider that relevant resources are those tagged by u using tags of q' to assess the obtained results. The random selection was carried out 10 times independently, and we report the average results.

A query expansion is expected to provide more resources as an answer to a query because of its enrichment, which generally causes an increase in the total recall. In our evaluation, we are more interested in studying the ability of the method to push relevant documents to the top of the ranking. Thus, we use the *Mean Average Precision (MAP)* and the *Mean Reciprocal Rank (MRR)*, two performance measures that take into account the ranking of relevant resources.

4.3 Study of the Parameters

We intend here to observe the parameters of our approach and estimate their optimal values. These parameters are:

- γ, which controls the semantic part and the social part in the ranking of tags for an expansion (see Eq. 4). The higher its value is, the stronger is the semantic part in tag similarity ranking, and vice versa.
- The number of tags which are suitable for the expansion.
- α, which gives either a higher importance to resources or to users, when computing the graph of tags $\mathcal{T}_{\mathcal{UR}}$. We set this parameter such that: the higher its value is, the stronger are the resources' links, and thus weaker the users links are, and vice versa (see Eq. 1).
- We evaluate two strategies for weighting the expanded terms (see Sect. 3.2.3).
- Finally, we observe the impact of the similarity measures over the search results.

We refer to our approach in Figs. 3, 4, 5, and 6 as *Personalized Social Query Expansion (PSQE)*. Also, all the Figures contain the results according to each similarity measure, and for each similarity measure, the results of the two weighting strategies are shown (this results in six curves per graph).

4.3.1 Impact of the Social Interest (γ)

The results showing the impact of the user interest w.r.t. the semantic similarity is given in Fig. 3. This latter shows the evolution of the MAP and the MRR for different values of γ, while fixing $\alpha = 0.5$ and query size to 4 for our three datasets, and using the three similarity measures. We note that the smaller the value of γ is, the better is the performance. This can be explained by the fact that the higher the value of the user interest part, the more resources that the user tags are highlighted (probably other users tag them with the same tags), and the higher is the value of the MAP and the MRR. However, we consider that neglecting the semantic part of Eq. 4 is not suitable for the following reasons: (i) First, if we fix γ to 0, we are going to neglect the semantic part, and perhaps lose the query sense (even if the potential terms to expand the query are those related to the query terms); (ii) Second, if we fix γ to 0 we are going to face cold start problems, since new users don't have an initial profile that allows us to rank terms. Thus, we choose to fix γ to 0.5 for the rest of the evaluations.

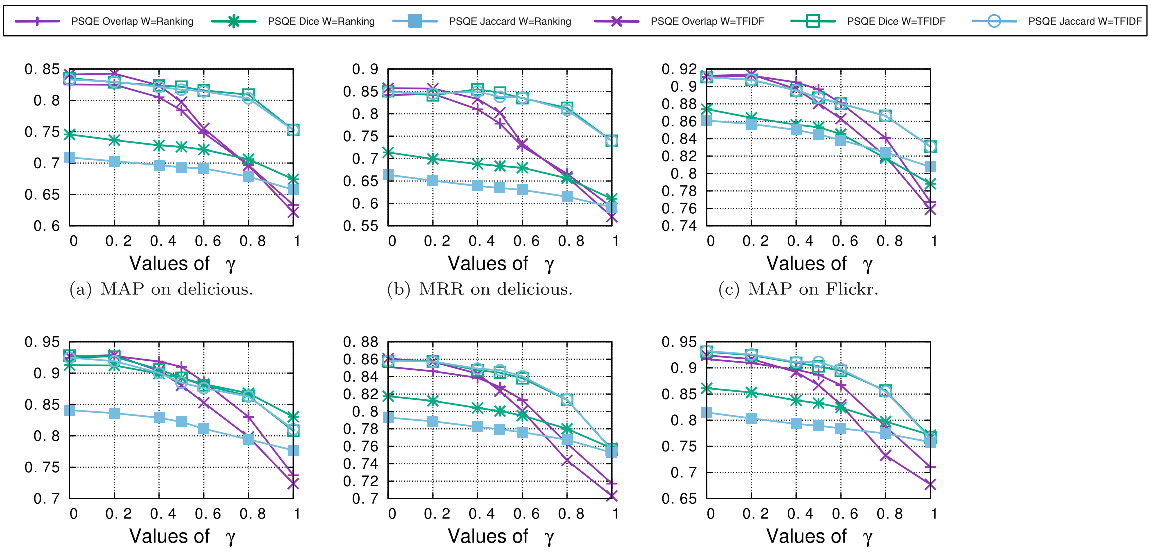

(a) MAP on delicious. (b) MRR on delicious. (c) MAP on Flickr. (d) MRR on Flickr. (e) MAP on CiteULike. (f) MRR on CiteULike.

Fig. 3. Measuring the impact of the social interest (γ). For different values of γ, we fix $\alpha = 0.5$, query size $= 4$ and we use the three similarity measures and the two weighting strategies for new terms averaged over 1000 queries, using the VSM.

4.3.2 Impact of the Query Size

The objective here is to check if the length of a query impacts the obtained results. The results are illustrated in Fig. 4. Through all the experiments we have performed, it comes out that the maximum performance is achieved while adding 4 to 6 related terms to the query. Adding more than 6 related terms has no impact on the quality of the results when using values of Eq. 4 as weight

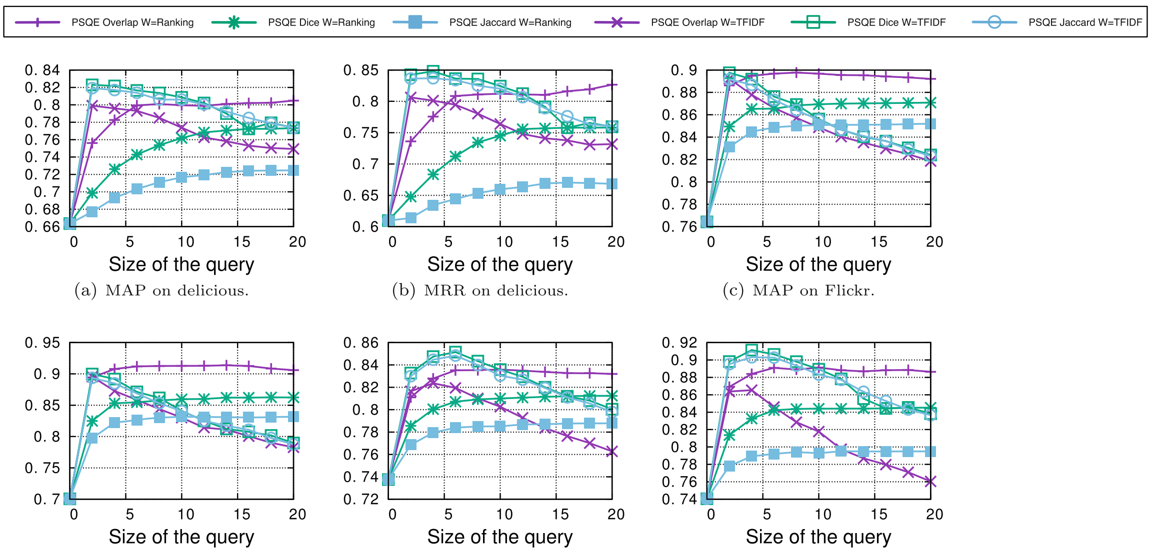

(a) MAP on delicious. (b) MRR on delicious. (c) MAP on Flickr. (d) MRR on Flickr. (e) MAP on CiteULike. (f) MRR on CiteULike.

Fig. 4. Evaluating the impact of the query size on the expansion. For different values of the query size, we use $\gamma = 0.5$, $\alpha = 0.5$ and our two strategies of weighting new terms.

for new term. This has even a negative impact when using TF-IDF values for term weighting as Fig. 4 shows. For the first case, this is due to the fact that the weight of the added terms is close to 0 (we remind that the weight of the added terms is the value of Eq. 4). Hence, this makes it natural and intuitive to pick a value in the provided interval, between 4 and 6.

4.3.3 Impact of the Users and Resources (α)

The importance of users and resources on the way the expansion is performed can be tuned by the parameter α of Eq. 1. Fixing $\alpha = 0$ considers only links between tags based on common users while fixing $\alpha = 1$ considers only links between tags based on common resources. The results regarding this parameter are illustrated in Fig. 5, where the MAP and the MRR's behaviors are quite different on the three datasets.

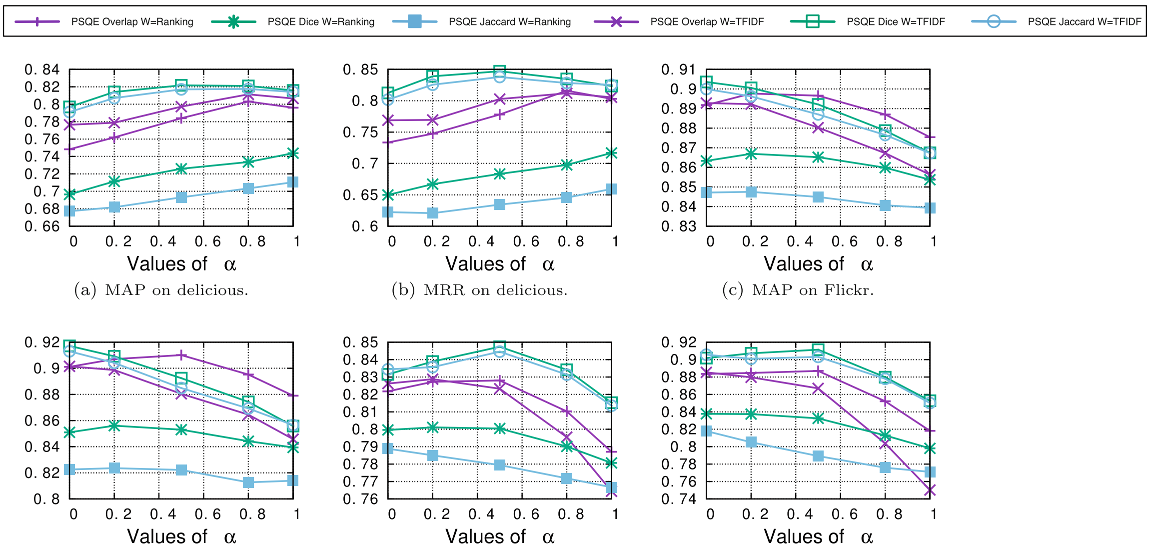

(a) MAP on delicious. (b) MRR on delicious. (c) MAP on Flickr.
(d) MRR on the Flickr. (e) MAP on CiteULike. (f) MRR on CiteULike.

Fig. 5. Evaluating the impact of the users/resources on the expansion. For values of α, using the three similarity measures, $\gamma = 0.5$, query size = 4 and for our two strategies of weighting new terms.

Indeed, in the *delicious* dataset, the values of the MAP and MRR increase by increasing the value of α using both the *Jaccard* and the *Dice* similarities achieving an optimal performance at $\alpha = 1$. As for *Flickr* and *CiteULike*, the optimal performance is achieved for $\alpha = 0.2$ and $\alpha = 0.5$ respectively. We believe that this is due to the fact that in social bookmarking systems like *delicious*, users are expected to share and annotate the same resources (URLs in *delicious*) to give rise to less private resources. Therefore, annotations are expected to occur more on resources than on users. However, in social bookmarking systems like *Flickr* and *CiteULike*, users are expected to upload their own resources (images and papers) resulting in more private resources. Thus, annotations are expected to occur more on users than on resources, a property which has been also observed and reported in [16].

4.3.4 Impact of the Weight of Terms

In Sect. 3.2.3, we explain that we experiment two strategies for weighting the new expanded terms by either (i) using value of Formula 4, or (ii) the *TF-IDF* value using Formula 5. We note that the performances follow almost the same distribution while varying γ and α in Fig. 3 and 5, and for our three similarity measures over our three datasets. However, we report that each time, the *TF-IDF* weighting strategy provides better performance. Hence, we conclude that personalizing the term weighting is less advantageous and less efficient comparing to a uniform weighting approach as used in the second strategy.

4.3.5 Impact of the Similarity Measures

The behavior of the performance seem to be the same for the three similarity measures with each time a small advantage to the *Dice* measure. Hence, taking into account the ratio between all the entities to which two tags are associated together versus the union of these entities leads to a better estimation of the similarity in folksonomies.

4.4 Comparison with Existing Approaches

Our objective here is to estimate how well our approach meets the users' information needs and compare its retrieval quality to that of other approaches, objectively. Our approach is evaluated using the optimal values computed in the previous section and using our two strategies of term weighting as explained in Sect. 3.2.3. The results are illustrated in Fig. 6 as *"PSQE-W = Ranking"* for the first strategy and *"PSQE-W = TFIDF"* for the second strategy, where we select four baselines for comparison as described in the following. Note that we choose the parameters that give the optimal performance for each of these baselines.

4.4.1 PSQE vs NoQE

The first approach for comparison is that with no query expansion or personalization. Documents that match queries are retrieved, and ranked as explained above. We report the following improvements:

***Delicious* dataset:** we obtain an improvement of almost 13% of the MAP and 18% of the MRR for our first strategy of term weighting using the Overlap similarity measure, and an improvement of almost 16% of the MAP and 24% of the MRR for our second strategy of term weighting using the Dice similarity measure.

***Flickr* dataset:** we obtain an improvement of almost 13% of the MAP and 21% of the MRR for our first strategy of term weighting using the Overlap similarity measure, and an improvement of almost 14% of the MAP and 21%

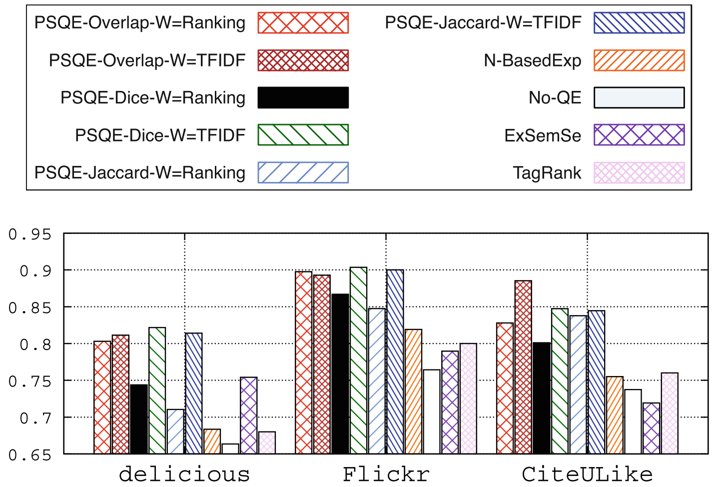

(a) Mean Average Precision (MAP).

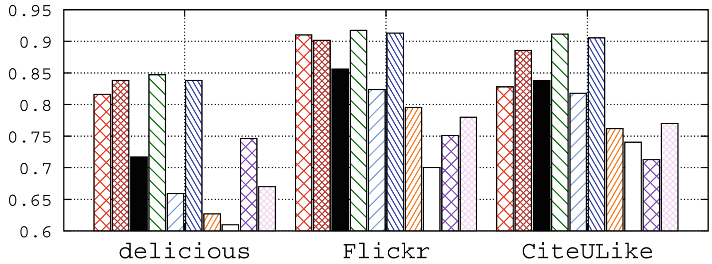

(b) Mean Reciprocal Rank (MRR).

Fig. 6. Comparison with the different baselines of the MAP and MRR, while fixing $\gamma = 0.5$ and query size$=4$, using the delicious, Flickr, and CiteULike datasets. We choose the optimal value of α for each similarity measure.

of the MRR for our second strategy of term weighting using the Dice similarity measure.

***CiteULike* dataset:** we obtain an improvement of almost 10% of the MAP and 7% of the MRR for our first strategy of term weighting using the Jaccard similarity measure, and an improvement of almost 15% of the MAP and 14% of the MRR for our second strategy of term weighting using the Overlap similarity measure.

Thus, it is clear that the query expansion has an evident advantage compared to a strategy with no expansion. We refer to this approach as **NoQE** in Fig. 6.

4.4.2 PSQE vs N-BasedExp

The second approach is the neighborhood based approach, which is based on the co-occurrence of terms over resources. This approach consists of enriching the query q with the most related terms without considering the user profile. Thus, queries are enriched similarly for each user. Our approach significantly outperform the neighborhood based approach as follows:

Delicious **dataset:** we obtain an improvement of almost 12% of the MAP and 19% of the MRR for our first strategy of term weighting using the Overlap similarity measure, and an improvement of almost 14% of the MAP and 22% of the MRR for our second strategy of term weighting using the Dice similarity measure.

Flickr **dataset:** we obtain an improvement of almost 8% of the MAP and 12% of the MRR for our first strategy of term weighting using the Overlap similarity measure, and an improvement of almost 9% of the MAP and 12% of the MRR for our second strategy of term weighting using the Dice similarity measure.

CiteULike **dataset:** we obtain an improvement of almost 8% of the MAP and 5% of the MRR for our first strategy of term weighting using the Jaccard similarity measure, and an improvement of almost 13% of the MAP and 12% of the MRR for our second strategy of term weighting using the Overlap similarity measure.

Therefore, we conclude that our personalized query expansion efforts bring a considerable contribution according to an approach based on the most related terms. We refer to this approach as **N-BasedExp** in Fig. 6.

4.4.3 PSQE vs ExSemSe

The third approach is an approach proposed in [4], which is a strategy that uses semantic search with query expansion named *Expanded Semantic Search*. In summary, this strategy consists of adding to the query q, k possible expansion tags with the largest similarity to the original tags in order to enrich its results. For each query, the query initiator u, ranks results using BM25 and tag similarity scores. We implemented this strategy and evaluated it over our datasets. We refer to this approach as **ExSemSe** in Fig. 6. We report the following improvements:

Delicious **dataset:** we obtain an improvement of almost 5% of the MAP and 7% of the MRR for our first strategy of term weighting using the Overlap similarity measure, and an improvement of almost 7% of the MAP and 10% of the MRR for our second strategy of term weighting using the Dice similarity measure.

Flickr **dataset:** we obtain an improvement of almost 11% of the MAP and 16% of the MRR for our first strategy of term weighting using the Overlap similarity measure, and an improvement of almost 12% of the MAP and 16% of the MRR for our second strategy of term weighting using the Dice similarity measure.

CiteULike **dataset:** we obtain an improvement of almost 12% of the MAP and 10% of the MRR for our first strategy of term weighting using the Jaccard similarity measure, and an improvement of almost 17% of the MAP and 17% of the MRR for our second strategy of term weighting using the Overlap similarity measure.

4.4.4 PSQE vs TagRank

The fourth approach is an approach proposed in [6], which is an algorithm called *TagRank* that automatically determines which tags best expand a list of tags in a given query. We implemented this strategy and evaluated it over our datasets. We refer to this approach as **TagRank** in Fig. 6. We report the following improvements:

***Delicious* dataset:** we obtain an improvement of almost 18.10% of the MAP and 21,79% of the MRR for our first strategy of term weighting using the Overlap similarity measure, and an improvement of almost 20.83% of the MAP and 26.42% of the MRR for our second strategy of term weighting using the Dice similarity measure.

***Flickr* dataset:** we obtain an improvement of almost 12.20% of the MAP and 16,67% of the MRR for our first strategy of term weighting using the Overlap similarity measure, and an improvement of almost 12.94% of the MAP and 17.58% of the MRR for our second strategy of term weighting using the Dice similarity measure.

***CiteULike* dataset:** we obtain an improvement of almost 10.23% of the MAP and 8,79% of the MRR for our first strategy of term weighting using the Jaccard similarity measure, and an improvement of almost 16.49% of the MAP and 18.35% of the MRR for our second strategy of term weighting using the Overlap similarity measure.

In summary, the obtained results show that our approach of personalization in query expansion using social knowledge may significantly improve web search. By comparing the PSQE framework to the closest state of the art approaches, we show that it is a very competitive approach that mays provide high quality results whatever the dataset used. Finally, we notice that the better performance are obtained with the *Dice* similarity measure and using TF-IDF for term weighting over our three datasets.

5 Related Work

Current models of information retrieval are blind to the social context that surrounds information resources, e.g., the authorship and usage of information sources, and the social context of the user that issues the query, i.e., his social activities of commenting, rating and sharing resources in social platforms. Therefore, recently, the fields of Information Retrieval (IR) and Social Networks Analysis (SNA) have been bridged resulting in Social Information Retrieval (SIR) models [20]. These models are expected to extend conventional IR models to incorporate social information [11].

In this paper, we are mainly interested in how to use social information to improve classic web search, in particular the query expansion process. Hence, we cite in the following, the main works that deal with social query expansion:

Biancalana et al. [7] proposed *Nereau*, a Query expansion strategy where the co-occurrence matrix of terms in documents is enhanced with meta-data retrieved from social bookmarking services. The system can record and interpret users' behavior, in order to provide personalized search results, according to their interests in such a way that allows the selection of terms that are candidates of the expansion based on original terms inserted by the user.

Bender et al. [4] consider SIR from both the query expansion and results ranking and propose a model that deals more with ranking results than query expansion. Lioma et al. [27] provide Social-QE by considering the query expansion (QE) as a logical inference and by considering the addition of tags as an extra deduction to this process. In the same spirit, Jin et al. [24] propose a method in which the used expansion terms are selected from a large amount of social tags in folksonomy. A tag co-occurrence method for similar terms selection is used to choose good expansion terms from the candidate tags directly according to their potential impact on the retrieval effectiveness. The work in [29] proposes a unified framework to address complex queries on multi-modal "social" collections. The approach they proposed includes a query expansion strategy that incorporates both textual and social elements. Finally, Lin et al. [26] propose this to enrich the source of terms expansion initially composed of relevant feedback data with social annotations. In particular, they propose a learning term ranking approach based on this source in order to enhance and boost the IR performances. Note that in these works, there is no personalization of the expansion process.

Bertier et al. [6] propose *TagRank* algorithm, an adaptation of the celebrated *PageRank* algorithm, which automatically determines which tags best expand a list of tags in a given query. This is achieved by creating and maintaining a *TagMap* matrix, a central abstraction that captures the personalized relationships between tags, which is constructed by dynamically computing the estimation of a distance between taggers, based on cosine similarity between tags and items. From our point of view, the proposed solution is not really suitable, since it needs the creation and the maintenance of a *TagMap* matrix for each user and the execution of an algorithm for determining close users with a high complexity.

Finally, a more recent work by Zhou et al. [44] proposes first a model to construct user profiles using tags and annotations together with documents retrieved from an external corpus. The model integrates the word embeddings text representation, with topic models in two groups of pseudo-aligned documents. Based on user profiles, the authors built two query expansion techniques based on: (i) topical weights-enhanced word embeddings, and (ii) the topical relevance between the query and the terms inside a user profile.

6 Conclusion and Future Work

This paper discusses a contribution to the area of query expansion leveraging the social context of the Web. We proposed a new approach based on social personalization to transform an initial query q to another query q' enriched with close terms that are mostly used by not only a given user but also by his social relatives. Given a social graph (folksonomy), the proposed approach starts by creating and maintaining a similarity graph of tags, that represents semantic strength between tags. The steps required to generate this graph of tags is operated offline, before the system is ready to process any query. Once this graph is created, a user profile is also created offline and maintained online for each user. These structures are used to compute personalized expansions on the fly thanks to the combination of the semantic and social dimensions. We demonstrated the effectiveness of our approach by an intensive evaluation on three large public datasets crawled from delicious, Flickr, and CiteULike. We showed that the expanded queries built by our method provide more accurate results as compared to the initial queries, by increasing the MAP in a range of 10 to 16% on the three datasets. We also compared our method to three state of the art baselines, and we showed that our query expansion method allows significant improvement in MAP, with a boost in a range between 5 to 18%. Finally, the proposed approach is being integrated into a system called *LAICOS* [9,13], which can be easily plugged into existing social bookmarking platforms.

Even with the interest of the proposed method, there are still possible improvements that we can bring. We believe that our approach is complementary to some existing approaches in the area of SIR. Thus, we are convinced that a combination with social ranking functions such as those proposed in [10,17,22,35,43] can be of a great interest.

Conflict of Interest. The author(s) declare(s) that there is no conflict of interest regarding the publication of this paper.

References

1. Baeza-Yates, R.A., Ribeiro-Neto, B.: Modern Information Retrieval: The Concepts and Technology Behind Search, 2nd edn. Addison-Wesley Longman Publishing Co. Inc., Boston (2011)
2. Bao, S., Xue, G., Wu, X., Yu, Y., Fei, B., Su, Z.: Optimizing web search using social annotations. In: Proceedings of the 16th International Conference on World Wide Web, WWW 2007, pp. 501–510. ACM, New York (2007)
3. Belkin, N.J.: Some(what) grand challenges for information retrieval. SIGIR Forum **42**(1), 47–54 (2008)
4. Bender, M., et al.: Exploiting social relations for query expansion and result ranking. In: 2008 IEEE 24th International Conference on Data Engineering Workshop (2008)
5. Benz, D., Hotho, A., Jaschke, R., Krause, B., Stumme, G.: Query logs as folksonomies. Datenbank-Spektrum **10**, 15–24 (2010)
6. Bertier, M., Guerraoui, R., Leroy, V., Kermarrec, A.-M.: Toward personalized query expansion. In: Proceedings of the Second ACM EuroSys Workshop on Social Network Systems, SNS 2009, pp. 7–12. ACM, New York (2009)
7. Biancalana, C., Micarelli, A., Squarcella, C.: Nereau: a social approach to query expansion. In: Proceedings of the 10th ACM Workshop on Web Information and Data Management, WIDM 2008, pp. 95–102. ACM, New York (2008)
8. Bischoff, K., Firan, C.S., Nejdl, W., Paiu, R.: Can all tags be used for search? In: Proceedings of the 17th ACM Conference on Information and Knowledge Management, CIKM 2008, pp. 193–202. ACM, New York (2008)

9. Bouadjenek, M.R., Hacid, H., Bouzeghoub, M.: LAICOS: an open source platform for personalized social web search. In: Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2013, pp. 1446–1449. ACM, New York (2013)
10. Bouadjenek, M.R., Hacid, H., Bouzeghoub, M.: SoPRa: a new social personalized ranking function for improving web search. In: Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2013, pp. 861–864. ACM, New York (2013)
11. Bouadjenek, M.R., Hacid, H., Bouzeghoub, M.: Social networks and information retrieval, how are they converging? A survey, a taxonomy and an analysis of social information retrieval approaches and platforms. Inf. Syst. **56**, 1–18 (2016)
12. Bouadjenek, M.R., Hacid, H., Bouzeghoub, M., Daigremont, J.: Personalized social query expansion using social bookmarking systems. In: Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2011, pp. 1113–1114. ACM, New York (2011)
13. Bouadjenek, M.R., Hacid, H., Bouzeghoub, M., Vakali, A.: PerSaDoR: personalized social document representation for improving web search. Inf. Sci. **369**, 614–633 (2016)
14. Bouadjenek, M.R., Sanner, S., Ferraro, G.: A study of query reformulation for patent prior art search with partial patent applications. In: Proceedings of the 15th International Conference on Artificial Intelligence and Law, ICAIL 2015, pp. 23–32. ACM, New York (2015)
15. Bouadjenek, M.R., Verspoor, K.: Multi-field query expansion is effective for biomedical dataset retrieval. Database **2017**, bax062 (2017)
16. Carmel, D., Roitman, H., Yom-Tov, E.: Social bookmark weighting for search and recommendation. VLDB J. **19**(6), 761–775 (2010)
17. Carmel, D., et al.: Personalized social search based on the user's social network. In: Proceedings of the 18th ACM Conference on Information and Knowledge Management, CIKM 2009, pp. 1227–1236. ACM, New York (2009)
18. De, A., Diaz, E.E., Raghavan, V.V.: On fuzzy result merging for metasearch. In: 2007 IEEE International Fuzzy Systems Conference, pp. 1–6, July 2007
19. Efthimiadis, E.N.: Query expansion. In: Annual Review of Information Systems and Technology (ARIST) (1996)
20. Goh, D., Foo, S.: Social Information Retrieval Systems: Emerging Technologies and Applications for Searching the Web Effectively. Information Science Reference - Imprint of: IGI Publishing (2007)
21. Hammond, T., Hannay, T., Lund, B., Scott, J.: Social bookmarking tools: a general review. D-Lib Mag. **11**(4) (2005). http://www.citeulike.org/group/684/article/80546
22. Hotho, A., Jäschke, R., Schmitz, C., Stumme, G.: Information retrieval in folksonomies: search and ranking. In: Sure, Y., Domingue, J. (eds.) ESWC 2006. LNCS, vol. 4011, pp. 411–426. Springer, Heidelberg (2006). https://doi.org/10.1007/11762256_31
23. Hung, C.-C., Huang, Y.-C., Hsu, J.Y., Wu, D.K.: Tag-based user profiling for social media recommendation. In: Workshop on Intelligent Techniques for Web Personalization and Recommender Systems at AAAI 2008, Chicago, Illinois (2008)
24. Jin, S., Lin, H., Su, S.: Query expansion based on folksonomy tag co-occurrence analysis. In: 2009 IEEE International Conference on Granular Computing, pp. 300–305, August 2009

25. 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). https://doi.org/10.1007/978-3-540-78646-7_12
26. Lin, Y., Lin, H., Jin, S., Ye, Z.: Social annotation in query expansion: a machine learning approach. In: Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2011, pp. 405–414. ACM, New York (2011)
27. Lioma, C., Blanco, R., Moens, M.-F.: A logical inference approach to query expansion with social tags. In: Azzopardi, L., et al. (eds.) ICTIR 2009. LNCS, vol. 5766, pp. 358–361. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-04417-5_39
28. Lund, B., Hammond, T., Hannay, T., Flack, M.: Social bookmarking tools (ii): a case study - connotea. D-Lib Mag. **11**(4) (2005). https://dblp.org/pers/hd/h/Hammond:Tony
29. Mantrach, A., Renders, J.-M.: A general framework for people retrieval in social media with multiple roles. In: Baeza-Yates, R., et al. (eds.) ECIR 2012. LNCS, vol. 7224, pp. 512–516. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-28997-2_53
30. Markines, B., Cattuto, C., Menczer, F., Benz, D., Hotho, A., Stumme, G.: Evaluating similarity measures for emergent semantics of social tagging. In: Proceedings of the 18th International Conference on World Wide Web, WWW 2009, pp. 641–650. ACM, New York (2009)
31. Mei, Q., Jiang, J., Su, H., Zhai, C.: Searching and tagging: two sides of the same coin? Technical report, University of Illinois at UrbanaChampaign (2007)
32. Metzler, D., Croft, W.B.: A Markov random field model for term dependencies. In: Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2005, pp. 472–479. ACM, New York (2005)
33. Mika, P.: Ontologies are us: a unified model of social networks and semantics. Web Semant. **5**(1), 5–15 (2007)
34. Nielsen, J.: Participation inequality: Encouraging more users to contribute (2006)
35. Noll, M.G., Meinel, C.: Web search personalization via social bookmarking and tagging. In: Aberer, K., et al. (eds.) ASWC/ISWC -2007. LNCS, vol. 4825, pp. 367–380. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-76298-0_27
36. Porter, M.F.: An Algorithm for Suffix Stripping, pp. 313–316. Morgan Kaufmann Publishers Inc., San Francisco (1997)
37. Schenkel, R., et al.: Efficient top-k querying over social-tagging networks. In: Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2008, pp. 523–530. ACM, New York (2008)
38. Schifanella, R., Barrat, A., Cattuto, C., Markines, B., Menczer, F.: Folks in folksonomies: social link prediction from shared metadata. In: Proceedings of the Third ACM International Conference on Web Search and Data Mining, WSDM 2010, pp. 271–280. ACM, New York(2010)
39. Srikanth, M., Srihari, R.: Biterm language models for document retrieval. In: Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2002, pp. 425–426. ACM, New York (2002)

40. Stoyanovich, J., Amer-Yahia, S., Marlow, C., Yu, C.: Leveraging tagging to model user interests in del.icio.us. In: AAAI Spring Symposium: Social Information Processing, pp. 104–109 (2008)
41. Vallet, D., Cantador, I., Jose, J.M.: Personalizing web search with folksonomy-based user and document profiles. In: Gurrin, C., et al. (eds.) ECIR 2010. LNCS, vol. 5993, pp. 420–431. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-12275-0_37
42. Wetzker, R., Zimmermann, C., Bauckhage, C.: Analyzing social bookmarking systems: a del.icio.us cookbook. In: Proceedings of the ECAI 2008 Mining Social Data Workshop, ECAI 2008 (2008)
43. Xu, S., Bao, S., Fei, B., Su, Z., Yu, Y.: Exploring folksonomy for personalized search. In: Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2008, pp. 155–162. ACM, New York (2008)
44. Zhou, D., Wu, X., Zhao, W., Lawless, S., Liu, J.: Query expansion with enriched user profiles for personalized search utilizing folksonomy data. IEEE Trans. Knowl. Data Eng. (2017)