Web Search Personalization Using Social Data

Dong Zhou[1,2], Séamus Lawless[2], and Vincent Wade[2]

[1] School of Computer Science and Engineering, Hunan University of Science and Technology, Taoyuan Road, Xiangtan, Hunan, 411201, China
[2] Center for Next Generation Localisation, Knowledge and Date Engineering Group, School of Computer Science and Statistics, Trinity College Dublin, Dublin 2, Ireland
dongzhou1979@hotmail.com,
{Seamus.Lawless,Vincent.Wade}@scss.tcd.ie

Abstract. Web search that utilizes social tagging data suffers from an extreme example of the vocabulary mismatch problem encountered in traditional Information Retrieval (IR). This is due to the personalized, unrestricted vocabulary that users choose to describe and tag each resource. Previous research has proposed the utilization of query expansion to deal with search in this rather complicated space. However, non-personalized approaches based on relevance feedback and personalized approaches based on co-occurrence statistics have only demonstrated limited improvements. This paper proposes an **I**terative **P**ersonalized Query Expansion **A**lgorithm for **W**eb Search (***iPAW***), which is based on individual user profiles mined from the annotations and resources the user has marked. The method also incorporates a user model constructed from a co-occurrence matrix and from a Tag-Topic model where annotations and web documents are connected in a latent graph. The experimental results suggest that the proposed personalized query expansion method can produce better results than both the classical non-personalized search approach and other personalized query expansion methods. An “adaptivity factor” was further investigated to adjust the level of personalization.

Keywords: Personalized Web Search, Query Expansion, Social Data, Tag-Topic Model, Graph Algorithm.

1 Introduction

Over the past decade, the area of personalized web search has gained much attention in the literature [8, 12]. An important concern in personalized search systems is how to store and represent the gathered usage information. Some systems store this information in an individualized user model [18], while other systems maintain an aggregate view of usage information [1].

Personalization in search systems can be achieved by query adaptation, result adaptation, or both. Query adaptation attempts to expand (augment) the terms of the user’s query with other terms, with the aim of retrieving more relevant results [8]. In terms of personalized query expansion, additional terms often come from individual user profiles to assist the user in formulating a better query.

P. Zaphiris et al. (Eds.): TPDL 2012, LNCS 7489, pp. 298–310, 2012.

Recent years have also witnessed the explosive growth of information on the World Wide Web (WWW). In social tagging systems such as del.icio.us[1], users are able to annotate each web resource with any number of free-form tags of their own choice. This type of system provides an ideal test bed for personalized search [17]. A user profile can be easily derived from their feedback, providing a good indication of the user's interests.

However, the uncontrolled manner of social tagging results in the use of an unrestricted vocabulary. This makes searching through the collection difficult and generally less accurate. In current social media systems, search algorithms also tend to be rather simplistic in nature, relying upon term matching methods, which often fail to deal with the vocabulary mismatch problem and result in poor ranking results.

Query expansion can partially solve the above mentioned problem. A classic technique is Pseudo relevance feedback (PRF) or local analysis [3]. This approach has been previously proven to work well. However, in the context of personalized search, the selected terms may be different from the users' true interests, so that the retrieved documents may not be relevant to a particular user. There have been few attempts at selecting the appropriate expansion terms from a user profile [4, 5, 6]. Past research appears to favor tag-tag relationships, by selecting the most related tags from a user's profile to enhance the source query. Given the fact that the tags might not be the precise descriptions of resources, the resulting retrieval performance has been markedly low [4, 5]. Borrowing from the traditional Information Retrieval (IR) field, local analysis and co-occurrence based user profile representation have been adopted to expand the query according to a user's interaction with the system [6, 8]. However, in this case the selection of expansion terms is solely based on lexical matching between the query and the terms which exist in the user profile. If the terms are not found in the user's profile, the query cannot be expanded at all.

In this work, an **I**terative **P**ersonalized Query Expansion **A**lgorithm for **W**eb Search, called ***iPAW,*** is presented. This algorithm is based on individual user profiles. In the user profile, terms are modeled according to their relationships, which can be defined by co-occurrence statistics or defined by a tag-topic model, introduced below. Each term in the user profile will have an associated weighting score calculated based on its relationship with other terms in the profile and terms extracted from top-ranked documents. After calculation, the terms with the highest scores will be chosen to expand the original query. The intuition behind the model is the prior assumption of term consistency: *the most appropriate expansion terms for a query are likely to have the same weighting scores, be associated with, and influenced by terms extracted from the documents ranked highly for the initial query*. In other words, the selection of expansion terms for a given query is not solely based on lexical matching, but by iterative context enhancing and weight propagation. In addition, one important "adaptivity factor" was examined that could affect the expansion process.

[1] `http://www.delicious.com`

2 Related Work

Manual query expansion has been studied in early IR systems [11]. This approach demands user intervention and requires the user to be familiar with the search system, which is generally not true for the modern web. For these reasons, the overwhelming majority of search systems in existence today, function via automatic query expansion. One common technique employs a machine readable thesaurus to locate expansion terms in lists of synonyms [16]. Other approaches extract expansion terms from large collections of documents [13]. Local analysis involving relevance feedback is another popular category of approach. Explicit feedback is often difficult to obtain. An alternative method is implicit relevance feedback through pseudo relevance feedback (PRF) [3]. Web query logs are also used by researchers to bridge the gap between the user-centric query space and author-centric web page space [10]. However, in practice, acquiring web query logs is difficult for most researchers due to the various concerns of search companies.

Personalized web search has been extensively studied. There are approaches that utilize query log and click-through analysis [12]. There are systems that explore desktop data and external resources [8]. In personalized social media search, Personalization usually involves two general approaches. The first approach runs the unmodified original query for all users but re-ranks the returned results based on an individual user profile [7, 17]. Another group of work modifies or augments a user's original query, or query expansion. Researchers have frequently used co-occurring tags to enhance the source query [4, 5].

3 Personalized Query Expansion

3.1 Problem Definition

In social tagging systems such as del.icio.us, users can label interesting web pages with primarily short and unstructured *annotations* in natural language called *tags*. These web pages are denoted as a URL in the del.icio.us website. Textual content is crawled by following a URL that refers to a *document* or *web document*. Multimedia content is excluded in this research. In response to a *query*, an initial set of the most relevant documents is fetched. The top "*c*" ranked documents are assumed to be relevant, and therefore refer to the *top-ranked documents*. *Term* refers to a *word* in the vocabulary, these two terminologies are used interchangeably. Terms extracted from documents are specifically called *docTerm*, to be distinguished from general "terms" used in user profiles and from tags.

Formally, social tagging data can be represented by a tuple $\mathcal{P} \coloneqq (\mathcal{U}, \mathcal{D}, \mathcal{T}, \mathcal{A})$, where $\mathcal{U}, \mathcal{D}, \mathcal{T}$ are finite sets of users, web documents and tags, and $\mathcal{A} \subseteq \mathcal{U} \times \mathcal{D} \times \mathcal{T}$ is a ternary relation, whose elements are called tag assignments or annotations. The set of annotations of a user is defined as: $\mathcal{A}_u \coloneqq \{(t, d) | u, d, t \in \mathcal{A}\}$. The tag vocabulary of a user, is given as $\mathcal{T}_u \coloneqq \{t | (t, d) \in \mathcal{A}_u\}$. The user's set of documents is $\mathcal{D}_u \coloneqq \{d | (t, d) \in \mathcal{A}_u\}$. The docTerm vocabulary of a user is further defined to be

$docTerm_u \coloneqq \{w|w \in \mathcal{D}_u\}$ where w denotes the words in the document corpus. $\mathcal{T}_u$ is the full list of tags that the user has used, and $docTerm_u$ is the vocabulary extracted from the documents that the user has tagged. So that terms in a user profile could be chosen from $\mathcal{T}_u$, $\mathrm{doc}Term_u$, or $\mathcal{T}_u \cup \mathrm{doc}Term_u$.

Given a source query q, a set of terms/words in the user profile $\{w_1, w_2 \dots w_n\}$, and a set of initial top-ranked documents $\mathcal{D}^{top} = \{d_1, d_2 \dots d_c\}$ the goal is to return a ranked list of profile terms to be added to the query, regularized by terms extracted from the top-ranked documents.

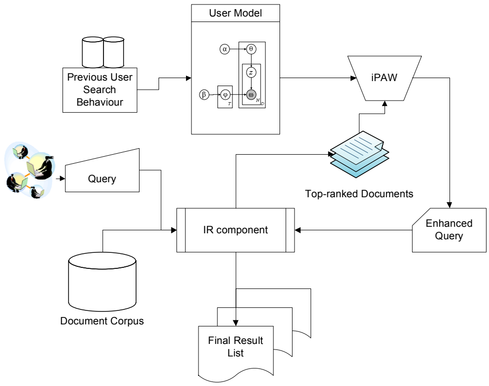

Fig. 1. Solution Architecture

3.2 Solution Overview

Figure 1 shows a sketch of the proposed solution architecture where personalized query expansion consists of two phases: (1) The construction of user models by using a co-occurrence matrix built for all the tags and documents in the training set or a tag-topic model where the tags and web documents are modeled simultaneously. This phase will be detailed in the next section. (2) The expansion of the given query by using the ***iPAW*** algorithm where the global term consistency is assumed over the word graph in addition to leveraging the top-ranked documents retrieved by the initial query. This process is described in the next section.

3.3 *iPAW* Algorithm

Let $G = (V, E)$ be a connected graph, where nodes V correspond to the n words in the user profile, and edges E correspond to the association strengths between words. An $n \times n$ symmetric weight matrix A on the edges of the graph is given, where a_{ij} denotes the weight between words w_i and w_j and M is a diagonal matrix with entries $M_{ii} = \sum_j a_{ij}$. A $n \times c$ matrix F is also defined with $F_{ij} = f(w, d)$ if a word w is presented in a document d and $F_{ij} = 0$ otherwise, where $f(w, d)$ denotes the weight of w in d.

Inspired by semi-supervised learning methods [19], here an iterative personalized query expansion algorithm for web search is developed. The algorithm is shown in Figure 2.

1. Form the affinity matrix A, which is defined by co-occurrence statistics or relationships that are calculated by the tag-topic model proposed in the user model construction section.
2. Construct the matrix $S = M^{-1/2}AM^{1/2}$.
3. Iterate $F(t+1) = \mu SF(t) + (1-\mu)F^0$ until convergence, where F^0 is the initial weighting matrix obtained by terms extracted from the top-ranked documents (*tf-idf* weighting used in the current paper (Jones 1988)) and μ is a parameter in (0,1).
4. Let F^* denote the limit of the sequence $\{F(t)\}$. Compute the final weighting scores for each word w as $w = \sum_{i=1}^{c} f(w, d_i)$, from which the top γ words can be acquired from the final ranked list of profile words to be added to the query.

Fig. 2. iPAW algorithm

The algorithm can be understood intuitively in terms of spreading activation networks [14]. A pairwise relationship is first defined, which can be viewed as the user model construction, and then it is symmetrically normalized, which is necessary for the convergence of the iteration step 3. During each iteration of step 3, terms in the user model will receive the information (weights) spreading from its neighbors, but also influenced by terms extracted from top-ranked documents (initial information). The parameter μ specifies the relative amount of the information from its neighbors and its initial weighting information.

Term consistency is assumed in the iteration phase. The first term of the iteration function in step 3 is the global consistency constraint, which means that a good weighting function should not change too much between nearby points. In this paper, nearby points are refined weighting scores with respect to initial relationships between words and context information (top-ranked documents) obtained by the initial query. They are likely to have the same effect over the graph. The second term is the fitting constraint, which means the weighting of words should fit the weighting scores of words extracted from the top-ranked documents retrieved by the given query.

After simplifying, a closed form solution can be derived as (see also [19, 20]):

$$F^* = (1-\mu)(I-\mu S)^{-1}F^0$$

An important feature of such computation is that the weightings calculated here share similarities with the entries obtained in the F^* calculation in Zhou et al's paper [19], where they try to find the largest entry in order to get the corresponding label, while this approach attempts to find select terms with large added weights. The actual values of weights are not so critical as far as they have discriminative power to separate high potential words from low potential words.

4 User Model Construction

In order to capture accurate information for the construction of the user model, in this paper the tags and web documents are modeled simultaneously. Initially, a simple user model is defined using a co-occurrence matrix. A new model is then introduced where a latent graph is built among terms.

1. For each tag $t \in \mathcal{T}$, choose $\theta_t \sim Dirichlet(\alpha)$
   For each topic $o \in \mathcal{O}$, choose $\varphi_o \sim Dirichlet(\beta)$
2. For each document $d \in \mathcal{D}$
   Given the vector of tags $\boldsymbol{t}_d$
   For each word w_i indexed by $i = 1, \dots, N_d$
   Conditional on $\boldsymbol{t}_d$ choose an tag $x_i \sim Uniform(\boldsymbol{t}_d)$
   Conditional on x_i choose a topic $z_i \sim Discrete(\theta_{x_i})$
   Conditional on z_i choose a word $w_i \sim Discrete(\varphi_{z_i})$

Fig. 3. Generative process of the Tag-Topic model

4.1 Co-occurrence Matrix

Firstly, a co-occurrence matrix was built according to [6, 8]. For all the tags and documents in the training set, important terms (or keywords) with high *tf-idf* scores (20% was used in the experiment described below) were selected. The cosine similarity between two words w_i and w_j was then calculated using:

$$cos(w_i, w_j) = DF_{w_i,w_j} / \sqrt{DF_{w_i} \cdot DF_{w_j}}$$

Where $DF_{w_{i,}}$ is the document frequency of word w_i. Note that here tags and doc-Terms are modeled into the matrix together. Graph G is built in which the nodes denote the terms and the edges E are weighted by their co-occurrence similariy.

4.2 A Tag-Topic Model

In the current paper, an Author-Topic model, introduced by Steyvers et al. [15], was adopted and a Tag-Topic model was proposed in order to learn topic-word and tag-topic distributions from the annotation data in an unsupervised manner. Then a latent graph is built based upon the features derived. This can be achieved through important docTerms, tags or a mixture of both.

To run the original Author-Topic model on the social tagging data at an individual user level, the tags can be viewed as authors in the new proposed model. When generating a document, a tag is chosen at random for each individual word in the document. This tag picks a topic from its multinomial distribution over topics, and then samples a word from the multinomial distribution over words associated with that topic. This process is repeated for all words in the document. This process is summarized in Figure 3.

In the figure, θ and φ denote topic-word distributions and tag-topic distributions respectively, while α and β denote Dirichlet priors. From the count matrices obtained during the modeling process [15], θ and φ can be easily estimated. The algorithm assigns words to random topics and tags (from the set of tags annotated to the document), and then repeats the Gibbs sampling process to update topic assignments for several iterations.

After the topic-word distributions and tag-topic distributions have been obtained, the adjacency graph of word associations and tag associations can be constructed. To illustrate how the model could be used in this respect, taking tags as an example, the distance between tags t_i and t_j was defined as the symmetric *KL divergence* between the topic distributions conditioned on each of the tags:

$$symKL(t_i, t_j) = \sum_{o=1}^{O} \left[\theta_{io} \log \frac{\theta_{io}}{\theta_{jo}} + \theta_{jo} \log \frac{\theta_{jo}}{\theta_{io}} \right]$$

Similarly we can compute associations between words.

So a latent graph G is defined using the latent feature obtained from the tag-topic model, where the nodes denote the terms and the edges E are weighted by $symKL$. After normalization, matrix S can be calculated. This process is executed offline, and then matrix S is saved for the query expansion model. Since terms extracted from the documents are not equally important, only the top δ terms from each topic are retained to form the graph. Tags are usually regarded as high quality descriptors of the web pages' topics and a good indicator of web users' interests, so they are all retained in the user profile. To compare the use of docTerms extracted from documents and tags assigned to the documents for query expansion, three sets of user profiles have been defined: selected docTerms, tags and a mixture of both.

5 Empirical Evaluation

5.1 Experimental Setup

In order to evaluate the methods on real-world data a crawl was conducted on the popular social tagging site delicious during 2010. A total of 5,943 users, 1,190,936 web pages and 283,339 tags were obtained. The average number of tags per user and pages per user in the corpus are 47.68 and 200.39, respectively. Users with less than 10 personal tags were removed from the sample.

Four groups of users were created according to the number of bookmarks (less than 50, 50-100, 100-500 and more than 500) associated with the users (see [17]). 50 randomly selected users from each group together with their tagging records were extracted to form a total collection of 200 test users. The English terms were processed in the usual way, i.e. down-casing the alphabetic characters, removing the stop words and stemming words using the Porter stemmer. All the pre-processed web pages were used in the experiments as the document corpus. For each user, 75% of his/her tags with annotated web pages were used to create the user profile and the other 25% were used as a test collection.

The evaluation method used by previous researchers in personalized social search [7, 17] is employed. The main assumption is as follows: Any documents tagged by u with t are considered relevant for the personalized query (u, t) (u submits the query t).

The following evaluation metrics were chosen to measure the effectiveness of the various approaches: the precision of the top 5 documents (P@5), mean reciprocal rank (MRR), mean average precision (MAP) and the recall of the top 5 documents (R@5). The first three measurements are commonly used to evaluate search algorithms while the last one is useful for evaluating query expansion systems as this method has been shown to improve both recall and precision in the past. Statistically-significant differences in performance were determined using a paired t-test at a confidence level of 95%.

Table 1. Overall Results

	MAP	MRR	P@5	R@5
BM25	0.0354	0.0411	0.0128	0.0483
BM25Prf	0.0278	0.034	0.0113	0.0398
LexMatch	0.0341	0.0432	0.0129	0.0471
iCo	0.042	0.0498	0.015	0.0533
iTerms	0.0413†*	0.0495†*	0.0155†*	0.0543†*
iTags	0.038*	0.0456†*	0.0138†*	0.0506†*
iMix	0.0448†*^{pl}	0.053†*^{pl}	0.0158†*^{pl}	0.0574†*^{pl}

5.2 Experimental Runs and Parameter Tuning

In order to usefully evaluate the performance of the personalized query expansion framework 2 different non-personalized baselines were selected: *BM25* – a popular and quite robust probabilistic retrieval method, and *BM25Prf* – a pseudo-relevance feedback oriented query expansion method based on the Divergence from Randomness theory [2]. In addition to non-personalized baselines, we also have several personalized baselines. Firstly, a method was adopted which processes the user model using lexical matching between query terms and terms that exist in the user model (constructed by using co-occurrence matrix) [8]. This method is denoted as *LexMatch*.

For our proposed approach, the method using ***iPAW*** and co-occurrence matrix as its user model representation is denoted as *iCo* (this can be viewed as an upper baseline). Finally, there are three variant approaches which use ***iPAW*** and the proposed tag-topic model. For those user models which only contain terms extracted from the documents, the algorithm is denoted as *iTerms*, for those user models which only contain tags, it is denoted as *iTags*. *iMix* is used to represent the method, which utilizes user models that contain a mixture of terms and tags.

For the tag-topic modeling, $\mathcal{O}$ and δ were set to 5 and 20 empirically. In the expansion framework, the optimal values were obtained when $\mu = 0.9$ and $\gamma = 10$ for terms (used in *LexMatch* as well) and $\gamma = 1$ for tags. The number of top documents $\mathcal{D}^{top}$ used in the query expansion framework was set to 10. The parameters for *BM25Prf* were also set to 10 documents and 10 terms.

5.3 Personalization vs Non-personalization

Firstly, we examine the experimental results that describe the performance of the three personalized query expansion runs (***iPAW*** with the tag-topic model) proposed in this paper together with two non-personalized baselines on the overall test users, which are shown in Table 1. The statistically significant differences are marked as † w.r.t to the *BM25* baseline and * w.r.t to the *BM25Prf* baseline.

As illustrated by the results, the *BM25Prf* model was the lowest performer for all evaluation metrics. This result is not surprising because the evaluation described in this paper is based upon a personalized-approach rather than the non-personalized evaluation model normally employed in large evaluation campaigns. This further demonstrates that merely borrowing common techniques from traditional IR will not solve the personalized search problem. Pleasingly, the three personalized query expansion-based search models all outperform the simpler text retrieval model with the highest improvement of 28.95% (In terms of the *iMix* method with the MRR metric when compared to *BM25*), which is statistically significant. It should be noted that all three personalized query expansion methods provide an average improvement of 40.74% compared to the traditional query expansion method with the highest improvement at 61.15%.

There were noticeable improvements in retrieval effectiveness when using user profiles that consisted of terms and a mixture of terms and tags in query expansion, but a more modest increase for the user profiles consisting tags alone. This reinforces the earlier finding that using tags alone for expanding queries is not sufficient. Another exciting observation is that in many cases, the personalized query expansion methods can outperform the baselines for all the evaluation metrics, with statistically significant improvements in almost entire runs.

5.4 Personalization via Lexical Matching and Co-occurrence statistics

The goal of the second set of experiments was to evaluate the performance of personalized query expansion using lexical matching and co-occurrence statistics (*LexMatch*), in comparison with the *iCo*. The statistically significant differences in the table are marked as l w.r.t to the *LexMatch* baseline and p w.r.t to the *iCo* for the *iMix* method only.

As we can see from Table 1, query expansion solely based on co-occurrence statistics and lexical matching is unsatisfactory. Although the performance is better in terms of MRR and P@5 metrics when compared to the non-personalized baseline, however, in the MAP metric the performance is even lower. After examining the expanded terms in the *LexMatch* model, it was found that because of the nature of social tagging systems, many tags are freely chosen and different from the terms stored in the user models. This results in a large number of queries being left un-expanded. Furthermore, the expanded terms sometimes show noise, resulting in lower performance than the *BM25* baseline.

However, using the same co-occurrence matrix as in *LexMatch*, the *iCo* method works much better, with performance just slightly lower than the *iMix* method. In some metrics it appears to work better than *iTerms*. This shows the power of using pseudo-relevance feedback documents to enhance the word graphs. Also the effectiveness of using the Tag-Topic model is also empirically confirmed (in terms of *iMix* which works better than *iCo*). It should be noted that the improvements achieved are statistically significant.

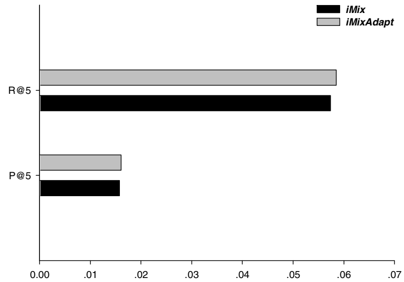

Fig. 4. Adaptivity Factor

5.5 Groups and Recall Results

A comparison of the results was also conducted across all the user groups. It is worth mentioning that there are improvements for users with low levels of activity. This demonstrates that the method is very effective even for users who do not have much available data in online social tagging systems.

In addition to the precision-based measurements, the personalized query expansion methods also showed significant improvements w.r.t the recall-based metric R@5. These improvements are on a similar scale when compared to the non-personalized and various personalized baselines. This reveals the benefits of adopting this method even in situations where recall is important.

6 Adaptivity Factor

In the previous section personalized query expansion was applied to expand all queries for each user. However, an optimal personalized query expansion algorithm should automatically adapt itself to various aspects of each query by leaving non-ambiguous queries un-expanded. This section examines one important factor that could affect the expansion process. Initial experimental results when using the factor, namely query clarity, are also presented.

It has been long known that the success of IR systems clearly varies across different topics. Therefore, a widely accepted method was adopted, called query clarity [9], to automatically tweak the amount of personalization fed into the algorithm. It measures the divergence between the language model associated with the user query and the language model associated with the collection. A simplified version was employed here [8] and it is defined as:

$$QC = \sum_{w \in q} P_{ml}(w|q) \cdot log \frac{P_{ml}(w|q)}{P_{coll}(w)}$$

where $P_{ml}(w|q)$ is the probability of the word w within the submitted query, and $P_{coll}(w)$ is the probability of w in the entire collection of documents.

In general, when there is no relevance information available for queries, Cronen-Townsend et al. [9] proposed using the scale of possible clarity scores for the collection at hand to heuristically set the clarity score threshold. In this research, the threshold is set heuristically to 90%. Simply put, a query is deemed “clear enough” if an estimated 90% or more of all queries would have a lower clarity score. If a query is defined to be clear, it is left unaltered in the retrieval process.

The same experimental setup was used as previously, however, only the *iMix* method is evaluated (where adaptation method is denoted as *iMixAdapt*), with P@5 and R@5. It can be seen from the results (Figure 4), there is more significant improvement in terms of recall than precision. In fact, the difference between the two methods in terms of R@5 is statistically significant. The major reason for this difference appears to be the arbitrary selection of queries to be expanded.

7 Conclusion

In this paper an **I**terative **P**ersonalized Query Expansion **A**lgorithm for **W**eb Search, called ***iPAW,*** was described which is based on individual user profiles mined from the annotations and resources the users bookmarked. The intuition behind the model is the prior assumption of term consistency. A tag-topic model for the construction of user models was also introduced which simultaneously integrates the annotations and web documents through a statistical model in a latent space graph. The proposed personalized technique performed well on the social data crawled from the web, delivering statistically significant improvements over non-personalized and personalized representative baseline systems. An adaptivity factor was applied to the proposed algorithm and also demonstrated improvements in a separate set of experiments.

Acknowledgements. This research is supported by the Science Foundation Ireland (Grant 07/CE/I1142) as part of the Centre for Next Generation Localisation (www.cngl.ie) at Trinity College Dublin.

References

1. Agichtein, E., Brill, E., Dumais, S.: Improving web search ranking by incorporating user behavior information. In: Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 19–26. ACM, Seattle (2006)
2. Amati, G., Rijsbergen, C.J.V.: Probabilistic models of information retrieval based on measuring the divergence from randomness. ACM Trans. Inf. Syst. 20(4), 357–389 (2002)
3. Baeza-Yates, R.A., Ribeiro-Neto, B.: Modern Information Retrieval. Addison-Wesley Longman Publishing Co., Inc. (1999)
4. Bender, M., Crecelius, T., Kacimi, M., Michel, S., Neumann, T., Parreira, J.X., Schenkel, R., Weikum, G.: Exploiting social relations for query expansion and result ranking. In: IEEE 24th International Conference on Data Engineering Workshop, ICDEW 2008, April 7-12, pp. 501–506 (2008)
5. 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, pp. 7–12. ACM, Nuremberg (2009)
6. Biancalana, C., Micarelli, A.: Social Tagging in Query Expansion: A New Way for Personalized Web Search. In: Proceedings of the 2009 International Conference on Computational Science and Engineering, vol. 04, pp. 1060–1065. IEEE Computer Society (2009)
7. 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: Proceeding of the 18th ACM Conference on Information and Knowledge Management, pp. 1227–1236. ACM, Hong Kong (2009)
8. Chirita, P.-A., Firan, C.S., Nejdl, W.: Personalized query expansion for the web. In: Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 7–14. ACM, Amsterdam (2007)
9. Cronen-Townsend, S., Zhou, Y., Croft, W.B.: Predicting query performance. In: Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 299–306. ACM, Tampere (2002)
10. Cui, H., Wen, J.-R., Nie, J.-Y., Ma, W.-Y.: Query Expansion by Mining User Logs. IEEE Trans. on Knowl. and Data Eng. 15(4), 829–839 (2003)
11. Harter, S.P.: Online information retrieval: concepts, principles, and techniques. Academic Press Professional, Inc. (1986)
12. Joachims, T., Granka, L., Pan, B., Hembrooke, H., Gay, G.: Accurately interpreting clickthrough data as implicit feedback. In: Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 154–161. ACM, Salvador (2005)
13. Qiu, Y., Frei, H.-P.: Concept based query expansion. In: Proceedings of the 16th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 160–169. ACM, Pittsburgh (1993)
14. Shrager, J., Hogg, T., Huberman, B.A.: Observation of phase transitions in spreading activation network. Science 236, 1092–1093 (1987)
15. Steyvers, M., Smyth, P., Rosen-Zvi, M., Griffiths, T.: Probabilistic author-topic models for information discovery. In: Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 306–315. ACM, Seattle (2004)
16. Voorhees, E.M.: Query expansion using lexical-semantic relations. In: Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 61–69. Springer-Verlag New York, Inc., Dublin (1994)

17. 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, pp. 155–162. ACM, Singapore (2008)
18. Zhang, Y., Koren, J.: Efficient bayesian hierarchical user modeling for recommendation system. In: Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 47–54. ACM, Amsterdam (2007)
19. Zhou, D., Bousquet, O., Lal, T.N., Weston, J., Scholkopf, B.: Learning with local and global consistency. In: Advances in Neural Information Processing Systems, vol. 16, pp. 321–328 (2004)
20. Zhu, X., Ghahramani, Z., Lafferty, J.: Semi-supervised learning using Gaussian fields and harmonic functions. In: The 20th International Conference on Machine Learning (ICML 2003), pp. 912–919 (2003)