

Knowledge-Based Query Expansion in Real-Time Microblog Search

Chao Lv, Runwei Qiang, Feifan Fan, and Jianwu Yang^(✉)

Institute of Computer Science and Technology,
Peking University, Beijing 100871, China
{lvchao,qiangrw,fanff,yangjw}@pku.edu.cn

Abstract. Since the length of microblog texts, such as tweets, is strictly limited to 140 characters, traditional Information Retrieval techniques usually suffer severely from the vocabulary mismatch problem such that they cannot yield good performance in the context of microblogosphere. To address this critical challenge, in this paper, we propose a new language modeling approach for microblog retrieval by inferring various types of context information. In particular, we expand the query using knowledge terms derived from Freebase so that the expanded one can better reflect the information need. Besides, in order to further answer users' real-time information need, we incorporate temporal evidences into the expansion methods so that the proposed approach can boost recent tweets in the retrieval results with respect to a given topic. Experimental results on two official TREC Twitter corpora demonstrate the significant superiority of our approach over baseline methods.

Keywords: Query expansion · Microblog search · Freebase

1 Introduction

Information Retrieval (IR) in the microblogosphere such as Twitter has attracted increasing research attention along with the fast development of social media. To explore the information seeking behavior in microblogosphere, TREC first introduced a Real-Time Search Task (RTST) in 2011 [15], which can be summarized as “At time T , give me the most relevant tweets about topic X ”.

However, it is inherently quite challenging to develop an effective real-time IR platform in the context of microblogosphere. First, in contrast to traditional web search techniques, real-time search task usually has to face to the problem of severe vocabulary mismatch. Since the tweets are very short, there is a large risk that query terms fail to match any word observed in relevant tweets. This problem is extremely severe especially when people search certain entities, which usually have several alternative aliases. Besides, real-time search usually indicates the information need concerning something that is happening right now. Thus, it is very crucial for the IR approach to boost the recent tweets relevant to the given topic. This real-time information need requires search engines to trade off between the recency and relevance score computed between the query and tweet.

Query Expansion (QE) methods based on pseudo-relevance feedback (PRF) [11, 13, 18] are widely used in microblog search to mitigate the problems mentioned above. However, these methods rely much on the assumption that the top ranked documents in the initial search are relevant and include good words for query expansion. Nevertheless, in real world, this assumption does not always hold in microblogosphere [4, 14], considering the example that the query contains proper nouns difficult to understand. What’s more, even if the top ranked documents are highly relevant to the topic, it is still very likely that they contain numerous topic-unrelated words due to the informality of the tweet content [14].

To overcome the limitations of existing methods, we utilize Freebase¹ as the knowledge source to infer more topic-related context information for each query. Freebase is a practical, scalable tuple database used to organize general human knowledge [3], covering a large amount of knowledge in different aspects (domains), into a hierarchical structure. In this paper, we propose a knowledge query generation method, in which we first match related concepts in Freebase with respect to the query, and then extract useful terms from different properties of the concepts to generate the knowledge query. By interpolating the original query with the knowledge query, we can better reflect the users’ information need.

To further incorporate the temporal evidence in microblogosphere, we follow the work of Li and Croft [9] and incorporate a prior distribution regarding to the recency of documents into the language modeling frameworks. More specifically, while selecting top knowledge terms from Freebase using an association based method, we assign each top ranked pseudo-relevance document with a time prior so that the words appearing in more recent documents are associated with higher probability.

The main contributions of this paper include: (1) we propose a novel approach to generate knowledge terms from Freebase to expand the original query, which can result in better understanding of information need; (2) the temporal evidence is incorporated in our QE method to trade off between relevance and recency; (3) we perform a set of experiments on two official twitter test collections published by TREC, to compare our proposed method with the state-of-the-art baseline methods. And, the experimental results demonstrate that our proposed approach can give rise to significant better retrieval performance.

2 Related Work

QE methods based on PRF assume that most frequent terms in the pseudo-relevance documents are useful, which may not hold in practice. Cao *et al.* [4] then integrated a term classification process to predict the effectiveness of expansion terms. Miyanishi *et al.* [14] proposed a manual tweet selection feedback to improve the retrieval performance. However, this method sometimes fails due to the content redundancy of tweets, which contain meaningless words that may degrade search results. Several approaches have been proposed to use the external resource such as Wikipedia, WordNet and ConceptNet to improve

¹ <http://www.freebase.com>

query expansion [7, 17]. Medelyan *et al.* [10] explored the possibilities of using Wikipedia’s articles as an external corpus to expand ad-hoc queries and demonstrated Wikipedia especially useful to improve weak queries which PRF is unable to improve. Pan *et al.* [16] proposed using Dempster-Shafer’s Evidence Theory to measure the certainty of expansion terms from the Freebase structure. To the best of our knowledge, query expansion based on Freebase knowledge in microblog search is novel and effective.

Previous works showed that temporal evidence can be incorporated into IR [5, 6]. Li and Croft [9] adopted a prior distribution regarding to the recency of documents into language modeling frameworks for retrieval. Liang *et al.* [11] proposed a temporal re-ranking component to evaluate the temporal aspects of documents. Miyanishi *et al.* [14] assumed that similar temporal models share similar temporal property and proposed a query-document dependent temporal relevance model. Albakour *et al.* [1] introduced a decay factor to balance the short-term and long-term interests for a given topic.

3 Proposed Methods

Given the RTST, we assume that a query Q is obtained as a sample from a generative model $\hat{\theta}_Q$, while the document D is generated by model $\hat{\theta}_D$. If $\hat{\theta}_Q$ and $\hat{\theta}_D$ are the estimated query and document language model respectively, according to [8], the relevance score of D with respect to Q can be computed by the following negative KL-divergence function:

$$S(Q, D) = -D(\hat{\theta}_Q || \hat{\theta}_D) \propto \sum_{w \in V} P(w | \hat{\theta}_Q) \cdot \log P(w | \hat{\theta}_D) \quad (1)$$

Within this ranking formula, the retrieval problem is essentially equivalent to the problem of estimating $\hat{\theta}_Q$ and $\hat{\theta}_D$. In principle, we can use any language model for the query and document, which is very flexible.

The start point of our study is to infer more topic-related context for the query with the help of Freebase. In this section, we first elaborate on why we choose Freebase as our knowledge base. Based on the characteristics of Freebase, we describe our proposed method of generating knowledge query in detail. Further improvements can be obtained by combining the knowledge-based query expansion with model-based pseudo-relevance feedback method.

3.1 Why We Choose Freebase

Freebase is a large collaborative knowledge base consisted of data harvested from sources such as the Semantic Web and Wikipedia, as well as individually contributed data from community members. In Freebase, human knowledge is described by structured categories, which are also known as *types* and each type has a number of defined *properties*. In this way, Freebase merges the scalability of structured databases with the diversity of collaborative wikis into a structured

general human knowledge [3]. This structured knowledge shows two superiorities compared with the semi-structured or plain contents: (1) when searching in the Freebase (with API), different types can be integrated for a more accurate concept, and some types such as name and alias are more important; (2) when generating knowledge terms, we can treat different types and the corresponding properties as different evidence sources.

3.2 Generation of Knowledge Query

We generate the knowledge query based on types, aiming extracting terms from different properties for a given query. The basic procedures of our proposed method include:

- **Concept Match.** We select the topic-related concepts with the help of Freebase API. Taking the query “Mila Kunis in Oz Movie” as an example, we match two concepts “Mila Kunis” and “The Wizard of Oz” in Freebase.
- **Term Selection.** Freebase describes the human knowledge of a given concept using types and properties. For some important meta types such as alias, name and notable_for, we directly add terms from these type properties to the knowledge query. For other types (i.e. description and domain-specific types), we adopt an association based term selection method to extract the topic-related top K terms.
- **Domain Filtering.** The concept in Freebase may be involved in several domains, such as TV, book and celebrities. When searching microblog about an entity (or concept) during a specific time period, users usually focus on one specific domain of that entity. Thus, aside from the common domain, we only use Freebase knowledge from one specific domain.

Then, we view the top ranked knowledge terms from selected properties equally to form a new knowledge query Q_{fb} . After that, the knowledge query model $\hat{\theta}_{Q_{fb}}$ is interpolated with the original query model $\hat{\theta}_Q$:

$$P(w|\hat{\theta}_{Q_1}) = (1 - \alpha) \cdot P(w|\hat{\theta}_Q) + \alpha \cdot P(w|\hat{\theta}_{Q_{fb}}) \quad (2)$$

where $\alpha \in [0, 1]$ is the weighting parameter to control the influence of the knowledge query. Both $\hat{\theta}_Q$ and $\hat{\theta}_{Q_{fb}}$ are estimated according to the maximum likelihood estimator.

Concept Match. We then describe our concept match algorithm in detail, which can be concluded as two steps:

1. **Noun Phrase Detection.** For a given query Q , we first split Q by space and receive a sequence of words q_1, q_2, \dots, q_n . Part-of-speech Tagging is then performed on each word, and all the noun phrases are extracted with rule-based method [2] from the original query.

2. **Maximum Match.** For each noun phrase, we regard it as a new query and get the related concepts as described in **Algorithm 1**. The `FreebaseSearch` function searches the given query in the Freebase and returns the top ranked concept if found. The match process ends if a related concept is found or none of the separate words can find a match.

Algorithm 1. `GetConcept(NQ)`

Require:Noun Phrase Query $NQ = q_1 q_2 \cdots q_n$.**Ensure:**Candidate Concept Set $CSet$.

- 1: $CSet = \text{FreebaseSearch}(NQ)$
 - 2: **if** $CSet$ is empty **then**
 - 3: **if** $n == 1$ **then**
 - 4: **return** \emptyset
 - 5: **end if**
 - 6: $NQ_1 \leftarrow q_1 q_2 \dots q_{n-1}$
 - 7: $NQ_2 \leftarrow q_2 q_3 \dots q_n$
 - 8: $CSet_1 \leftarrow \text{GetConcept}(NQ_1)$
 - 9: $CSet_2 \leftarrow \text{GetConcept}(NQ_2)$
 - 10: **return** $CSet_1 \cup CSet_2$
 - 11: **else**
 - 12: **return** $CSet$
 - 13: **end if**
-

Term Selection. For each returned concept from Freebase API, different types and corresponding properties which reflect the different aspects of the concept are provided by the search result. Some types (i.e. meta types) are very general and precise, such as `alias`, `name` and `notable_for` in the common domain, we directly add the property terms to the knowledge query for these types. When it comes to other types such as `description` and domain specific ones whose properties contain a lengthy text, an association based term selection method is utilized to extract the topic-related knowledge terms. Effective term selection is an important issue for an automatic query expansion technique. In microblog retrieval, a good expansion term should satisfy the following criteria:

1. The term should be semantically associated with the concept from the original query;
2. The term extracted from Freebase should be suitable for the local corpus and it should provide useful information for the query as well;
3. As the user's intent may change and events related to the given topic will develop over the time, the ranking function should favor the short-term words that are mostly used in recent tweets.

The candidate terms extracted from Freebase have met the first criterion to some extent. In order to satisfy the second criterion, we score the candidate terms with an association based method on the basis of the top ranked N pseudo-relevance documents (PRD):

$$Score(w) = \sum_{D \in PRD} P(D) \cdot P(w|D) \cdot \prod_{i=1}^n P(q_i|D) \quad (3)$$

where $P(D)$ is the document prior which is usually assumed to be uniform, and $\prod_{i=1}^n P(q_i|D)$ is the query likelihood given the document model, which is traditionally computed using Dirichlet smoothing. To further meet the third criterion, we follow the work of [9] and incorporate the temporal evidence into the document prior in Eq. 3 by using an exponential distribution:

$$P(D|T_D) = r \cdot e^{-r(T_Q - T_D)} \quad (4)$$

where r is the exponential parameter that controls the temporal influence, T_Q is the query issue time and T_D is the tweet post time. Both T_Q and T_D are measured in fractions of days. Note that T_D is constantly less than T_Q as we cannot use the future evidence. Finally, we select the top scored K words from the common description and domain specific properties, to form the knowledge query Q_{fb} along with the terms extracted from meta properties.

Domain Filtering. One advantage of Freebase is that it provides the domain information of a given concept. A domain is a collection of types which share a namespace (e.g. TV, book and celebrities). As mentioned above, only properties from common domain and top ranked domain are used for term selection. Hence, a domain ranking function should be defined to rank all the domains of a given concept. We first define the domain language model $\hat{\theta}_{dm}$ as follows:

$$P(w|\hat{\theta}_{dm}) = \frac{c(w, dm)}{\sum_{w' \in dm} c(w', dm)} \quad (5)$$

where $c(w, dm)$ is the count of word w occurred in the domain related properties. Then, we compute the KL-divergence score of the domain language model with the document language model estimated on the top ranked N PRD using empirical word distribution.

$$Score(dm) = -D(\hat{\theta}_{dm} || \hat{\theta}_{PRD}) \propto \sum_{w \in dm} P(w|\hat{\theta}_{dm}) \cdot \log P(w|\hat{\theta}_{PRD}) \quad (6)$$

Note that we use Dirichlet smoothing method while estimating $\hat{\theta}_{PRD}$. With this formula, we can rank all the domains and choose terms from the top ranked domain for term selection.

3.3 Mixture Feedback Model

With the knowledge query environment, we believe the information need is more understandable, which could lead to a high precision in the top retrieved documents. Based on this hypothesis, we further utilize a model-based feedback to update the query representation. More specifically, we update the $\hat{\theta}_{Q_1}$ with the simple mixture model $\hat{\theta}_F$ which is widely used in microblog retrieval [11, 18].

$$P(w|\hat{\theta}_{Q_2}) = (1 - \beta) \cdot P(w|\hat{\theta}_{Q_1}) + \beta \cdot P(w|\hat{\theta}_{Q_F}) \quad (7)$$

where $\beta \in [0, 1]$ is a weighting parameter to control the amount of model-based feedback.

The model-based feedback model generates a feedback document by mixing the query topic model $\hat{\theta}_F$ with the collection language model $\hat{\theta}_C$. Under this simple mixture model, the log-likelihood of feedback documents F is:

$$\log P(F|\hat{\theta}_F) = \sum_w c(w, F) \cdot \log((1 - \lambda) \cdot P(w|\hat{\theta}_F) + \lambda \cdot P(w|\hat{\theta}_C)) \quad (8)$$

where $c(w, F)$ is the count of word w occurred in the set of feedback documents F . Then we follow the work of [18] and implement the EM algorithm with the fixed smoothing parameter $\lambda = 0.5$. No matter whether or not the query finds its knowledge terms in Freebase, the query environment will be updated by the model-based feedback.

4 Evaluation

4.1 Experimental Setup

Data Set. Two corpora (i.e. Tweets11 and Tweets13 collection) are used in our experiments, which are crawled via the official API released by TREC organizers [12]. Tweets11 collection has a sample of about 16 million tweets, ranging from January 24, 2011 to February 8, 2011 while Tweets13 collection contains about 259 million tweets, ranging from February 1, 2013 to March 31, 2013. Tweets11 is used for evaluating the effectiveness of the proposed real-time Twitter search systems over 50 official topics in the TREC’11 Microblog track as well as 60 official topics in the TREC’12 Microblog track, respectively. And, Tweets13 is used in evaluating the proposed real-time Twitter search systems over 60 official topics in the TREC’13 Microblog track. In our experiments, topics for Tweets13 are used for tuning the parameters and then we use the best parameter settings to evaluate our methods with topics for Tweets11.

The tweets and their corresponding topic information were preprocessed in several ways. We first discarded the non-English tweets using a language detector, named *ldig*². Second, in conformance with the track’s guidelines, all simple retweets were removed by deleting documents beginning with the string ‘RT’. Moreover, we crawled all the shortened URLs and add the web title to the origin tweet. In the end, each tweet was stemmed using the Porter algorithm and stopwords were removed using the InQuery stopwords list.

² <http://github.com/shuyo/ldig>.

Evaluation Metric. In TREC Microblog Track, tweets were judged on the basis of the defined information using a three-point scale [15]: irrelevant (labeled as 0), minimally relevant (labeled as 1), and highly relevant (labeled as 2). The main evaluation metric is Mean Average Precision (MAP) for top 1000 documents and Precision at N (P@N), which are widely used in IR. MAP and P@30 with respect to *allrel* (i.e. tweet set judged as highly or minimally relevant) are used in this paper.

Baselines. To demonstrate the performance of our proposed method, we compare our knowledge-based query expansion methods with several baseline methods. The simple KL-divergence retrieval model (denoted as **SimpleKL**) [19] is used as our first baseline. That is, we estimate $\hat{\theta}_Q$ and $\hat{\theta}_D$ with empirical word distribution, and we choose Dirichlet smoothing method for document model estimation. Throughout this paper, we set the Dirichlet smoothing parameter $\mu = 100$. We use the Simple Mixture Model [18] (denoted as **QESMM**) as our second baseline, and optimize the number of feedback documents to 5 and the number of terms in the feedback model to 100. The smoothing parameter β is set as 0.9. We also compare our method with the state-of-the-art real-time ranking model (denoted as **RTRM**) under language modeling framework, proposed by [11]. **RTRM** approach utilized a two-stage pseudo-relevance feedback query expansion to estimate the query language model and expand documents with shortened URLs in microblog. We tune all the parameters of these models with TREC’13 topics on Tweets13 corpus.

4.2 Experimental Results

We conduct several experiments to measure the effects of our query expansion methods. For our knowledge-based query expansion method, we label the method with query model $\hat{\theta}_{Q_1}$ as **QEFB**, and the one with $\hat{\theta}_{Q_2}$ as **QEFB+SMM**. When selecting knowledge terms from Freebase description and domain specific property, we set the top ranked PRD number N to 100 and the expanded term number K to 10. Following Li and Croft [9], unless otherwise specified, we set the parameter r as 0.01. α in Eq. 2 is set as 0.5, which means we regard the original query and the knowledge query equally important. The query expansion parameters in the mixture feedback model are set like **QESMM**. All the parameters are tuned with TREC’13 topics. Then we test the optimized models on Tweets11 corpus with TREC’11 and TREC’12 topics.

Table 1 shows the performance comparison of different query expansion methods. To test for statistical significance, we used a paired t-test. † and ‡ indicate that the corresponding improvements over **SimpleKL** and **QESMM** are statistically significant ($p < 0.05$), respectively. Note that all the methods listed in the table estimate the document model as **SimpleKL**. As we can see, all of the query expansion methods have significant MAP and P@30 improvements compared with the **SimpleKL** method, which indicates the importance of query expansion in microblog retrieval. When the query is expanded with the Freebase knowledge query, our approach can retrieve more relevant documents in

Table 1. The performance comparison of different query expansion methods. The best performances are marked in bold.

Topics	TREC'13		TREC'12		TREC'11	
Method	MAP	P@30	MAP	P@30	MAP	P@30
SimpleKL	0.2926	0.4939	0.2681	0.3872	0.3572	0.3773
QESMM	0.3260	0.5117	0.2955	0.3989	0.3878	0.4133
RTRM	0.3486	0.5567	0.3282	0.4439	0.3975	0.4300
QEFB	0.3149 †	0.5117	0.3144†	0.4262	0.4125†	0.4207
QEFB+SMM	0.3652 ††	0.5750 ††	0.3320 ††	0.4450 ††	0.4241 †	0.4300 †

the top results. Thus, we can further improve the retrieval performance by combining the knowledge-based expansion method with mixture feedback model. Our knowledge-based query expansion method **QEFB+SMM** achieves the best retrieval performance in the three topic sets. More specifically, for TREC'12 topics, our method **QEFB+SMM** improves the MAP over **SimpleKL** and **QESMM** by 23.83 % and 12.35 %, respectively; while the corresponding increments in terms of P@30 are 14.93 % and 11.56 %, respectively. For TREC'11 topics, the **QEFB+SMM** raises the MAP over **SimpleKL** and **QESMM** by 18.73 % and 9.36 %, respectively; while the corresponding P@30 improvements are 13.97 % and 4.04 %, respectively. Our method also beats the state-of-the-art baseline **RTRM**, which has a good performance in Microblog Track.

4.3 Discussion

Many parameters in our proposed approach can affect the system performance. In this section, we analyze the robustness of the parameter settings in knowledge-based query expansion method. All these experiments in this section are run on TREC'13 topics, which are used for parameter selection.

Effects of Knowledge Query. For the query modeling, we propose using knowledge query to make the information need more comprehensible. Many factors affect the quality of the knowledge terms: (1) whether the maximum match algorithm can get topic-related concept from the Freebase; (2) whether the Freebase structured information (i.e. type information and domain information used in term selection and domain filtering) is helpful; (3) how the number of knowledge terms K affects the retrieval performance and (4) how the number of pseudo-relevance documents N used for term selection and domain selection affects the retrieval performance.

To answer the first question, we create the run **QEManualFB**, which means we manually select the concept from Freebase for each query. For the second question, we dismiss the structured information of Freebase and select terms from all the properties and regard all the properties equally. We label this run as **QEWiki**.

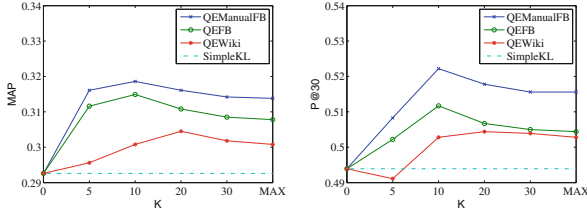


Fig. 1. Sensitivity to the selected knowledge term number K .




Fig. 2. Sensitivity to the PRD number N for knowledge term selection.

Figure 1 shows the MAP and P@30 scores of all the models for different K and fixed $N = 100$. *MAX* means all the candidate terms that satisfy $Score(w) > 0$ in Eq. 3 are selected. We can see that though **QEFB** is not better than **QEManualFB**, the performance gap between them is not large, which verifies the effectiveness of our concept match algorithm. **QEFB** is consistently better than **QEWiki**, which proves the importance of the structured information in Freebase. Moreover, when K is set around 10, **QEFB** can get its optimal retrieval performance and is significantly better than that of **SimpleKL**, which indicates the effectiveness of the association based term selection method.

We then fix the term number K to 10 and change the PRD number N for term selection and domain filtering. Figure 2 shows the MAP and P@30 scores of our **QEFB** model against different values of N . We can observe that when N is larger than 100, the performance changes slightly. It means that top 100 pseudo-relevance documents can provide enough information for selecting good knowledge terms from description property.

Effects of the Interpolation Coefficients. Recall that we first expand the query with knowledge query, and further expand the updated query $\hat{\theta}_{Q_1}$ with model-based feedback. The first-stage query expansion is controlled by a coefficient α , while the second-stage expansion is controlled by β . Figure 3 (left) shows the performance changing of **QEFB** ($N = 100, K = 10$) against different values of α . When $\alpha = 0$, the **QEFB** reduces to the baseline method **SimpleKL**. When $\alpha = 1$, we completely ignore the original query and only use the knowledge query. We can observe that the performance of **QEFB** is better than **SimpleKL** when α is no greater than 0.7. The optimal performance can be obtained when α is set




Fig. 3. Sensitivity to the first-stage knowledge query expansion coefficient α (left) and second-stage mixture feedback interpolation coefficient β (right).

around 0.5. Figure 3 (right) shows the performance changing of **QEFB+SMM** against different values of β . When $\beta = 0$, the **QEFB+SMM** reduces to **QEFB**. We can see the second-stage expansion seems to be more robust and constantly better than **QEFB**. After knowledge-based query expansion, the query can be more comprehensible and get more top related tweets, which lead to a further improvement with traditional model-based feedback.

5 Conclusion

In this study, we proposed using knowledge-based query expansion to solve the problems in microblog search. With the knowledge terms derived from the Freebase, the queries in microblogosphere can be more comprehensible and thus more relevant documents can be retrieved. The knowledge terms from Freebase should also co-occur with query terms in PRD, which has the potential to alleviate topic drift often induced by knowledge-based QE. Freebase’s structured information is also well utilized in knowledge query generation procedure. Moreover, we incorporated the temporal evidence into query representation. Thus the proposed method favors recent tweets which satisfy the real-time information need in microblog retrieval. Our thorough evaluation, using two standard TREC collections, demonstrates the effectiveness of the proposed method.

Acknowledgments. The work reported in this paper was supported by the National Natural Science Foundation of China Grant 61370116.

References

1. Albakour, M., Macdonald, C., Ounis, I., et al.: On sparsity and drift for effective real-time filtering in microblogs. In: Proceedings of the 22nd ACM International Conference on Information & Knowledge Management, pp. 419–428. ACM (2013)
2. Bird, S., Klein, E., Loper, E.: Natural Language Processing with Python. O’Reilly Media, Inc., Newton (2009)

3. Bollacker, K., Evans, C., Paritosh, P., Sturge, T., Taylor, J.: Freebase: a collaboratively created graph database for structuring human knowledge. In: Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data, pp. 1247–1250. ACM (2008)
4. Cao, G., Nie, J.Y., Gao, J., Robertson, S.: Selecting good expansion terms for pseudo-relevance feedback. In: Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 243–250. ACM (2008)
5. Dakka, W., Gravano, L., Ipeirotis, P.G.: Answering general time-sensitive queries. *IEEE Transactions on Knowledge and Data Engineering* **24**(2), 220–235 (2012)
6. Dong, A., Zhang, R., Kolari, P., Bai, J., Diaz, F., Chang, Y., Zheng, Z., Zha, H.: Time is of the essence: improving recency ranking using twitter data. In: Proceedings of the 19th International Conference on World Wide Web, pp. 331–340. ACM (2010)
7. Kotov, A., Zhai, C.: Tapping into knowledge base for concept feedback: leveraging conceptnet to improve search results for difficult queries. In: Proceedings of the fifth ACM International Conference on Web search and Data Mining, pp. 403–412. ACM (2012)
8. Lafferty, J., Zhai, C.: Document language models, query models, and risk minimization for information retrieval. In: Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 111–119. ACM (2001)
9. Li, X., Croft, W.B.: Time-based language models. In: Proceedings of the Twelfth International Conference on Information and Knowledge Management, pp. 469–475. ACM (2003)
10. Li, Y., Luk, W.P.R., Ho, K.S.E., Chung, F.L.K.: Improving weak ad-hoc queries using wikipedia asexternal corpus. In: Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 797–798. ACM (2007)
11. Liang, F., Qiang, R., Yang, J.: Exploiting real-time information retrieval in the microblogosphere. In: Proceedings of the 12th ACM/IEEE-CS Joint Conference on Digital Libraries, pp. 267–276. ACM (2012)
12. Lin, J., Efron, M.: Overview of the TREC-2013 microblog track. In: Proceedings of TREC. vol. 2013 (2013)
13. Lv, Y., Zhai, C.: A comparative study of methods for estimating query language models with pseudo feedback. In: Proceedings of the 18th ACM Conference on Information and Knowledge Management, pp. 1895–1898. ACM (2009)
14. Miyanishi, T., Seki, K., Uehara, K.: Improving pseudo-relevance feedback via tweet selection. In: Proceedings of the 22nd ACM International Conference on Information and Knowledge Management, pp. 439–448. ACM (2013)
15. Ounis, I., Macdonald, C., Lin, J., Soboroff, I.: Overview of TREC-2011 microblog track. In: Proceedings of the 20th Text REtrieval Conference (TREC 2011) (2011)
16. Pan, D., Zhang, P., Li, J., Song, D., Wen, J.-R., Hou, Y., Hu, B., Jia, Y., De Roeck, A.: Using Dempster-Shafer’s evidence theory for query expansion based on freebase knowledge. In: Banchs, R.E., Silvestri, F., Liu, T.-Y., Zhang, M., Gao, S., Lang, J. (eds.) AIRS 2013. LNCS, vol. 8281, pp. 121–132. Springer, Heidelberg (2013)

17. Xu, Y., Jones, G.J., Wang, B.: Query dependent pseudo-relevance feedback based on wikipedia. In: Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 59–66. ACM (2009)
18. Zhai, C., Lafferty, J.: Model-based feedback in the language modeling approach to information retrieval. In: Proceedings of the Tenth International Conference on Information and Knowledge Management, pp. 403–410. ACM (2001)
19. Zhai, C., Lafferty, J.: A study of smoothing methods for language models applied to ad hoc information retrieval. In: Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 334–342. ACM (2001)