

Modeling documents as mixtures of persons for expert finding

Pavel Serdyukov and Djoerd Hiemstra

Database Group, University of Twente,
PO Box 217, 7500 AE, Enschede, The Netherlands
{serdyukovpv, hiemstra}@cs.utwente.nl

Abstract. In this paper we address the problem of searching for knowledgeable persons within the enterprise, known as the expert finding (or expert search) task. We present a probabilistic algorithm using the assumption that terms in documents are produced by people who are mentioned in them. We represent documents retrieved to a query as mixtures of candidate experts language models. Two methods of personal language models extraction are proposed, as well as the way of combining them with other evidences of expertise. Experiments conducted with the TREC Enterprise collection demonstrate the superiority of our approach in comparison with the best one among existing solutions.

1 Introduction

In enterprises or in common web search settings users often experience the need not only for getting information, but for getting into the contact with those who could be the source of this information. The opportunity of interaction with a knowledgeable person is sometimes appreciated much higher than the access to a very relevant and clearly written document on the search topic [20]. An expert finding system helps to find individuals or even working groups possessing certain expertise and skills within an organization [6]. Quite like a typical document retrieval system, it uses a short user query and documents stored on personal desktops or within centralized databases as the input. The prediction of a personal expertise is made through the analysis of textual content of documents the person has relation to. The proof of relation can be authorship, simple occurrence in the text or just a fact that the document is stored locally at the PC (e.g. in the browser cache). For ensuring traceability, the system must return not only the ranking of people, but also the list of those documents that appeared to be the best indicators of expertness.

Apart from causing the boom on the enterprise search systems market [19], expert finding systems also compelled close attention of the IR research community. The expert search task is included into the Enterprise track of the Text REtrieval Conference (TREC) since 2005 [5]. The TREC community provided the experimental dataset and set up the standards for the evaluation.

The fundamental principle of state-of-art methods for expert finding is to infer personal expertise by studying the co-occurrence of personal identifiers

(names, email addresses etc.) and query terms in the scope of documents. The more often a person is detected in the documents containing many words describing the topic, the more likely we may rely on this person as an expert in this topic. However, all methods also consider that persons as well as terms occur in the document independently and do not influence the appearance of each other. Although, the assumption about independence among terms is a de facto standard in IR [7], the independence of terms from persons seems not quite adequate.

In this paper we consider that the occurrence of terms in the document can not be considered independent from the presence of a candidate expert. We propose a ranking model in which people are regarded as generators of the document's content. Our generative modeling method combines the features of both so-called profile- and document-centric approaches: it ranks candidates using their language models built from the retrieved documents and takes the frequency of appearance of a candidate in the top ranked documents as an additional evidence of his proficiency in a search topic.

The paper is organized as follows. In the next section we give a more detailed description of existing approaches to expert finding. In Section 3, we show how to utilize the assumption that persons mentioned in the document determine which terms it consists of. In Section 4, we explain how personal language models can be mined from retrieved documents and used further to build a good predictor of personal expertise. Experimental results supporting our assumptions are given in Section 5. Discussion of the paper outcome and a brief outline of potential future work can be found in Section 6.

2 Related work

Existing approaches to candidate experts modeling and ranking are basically variations of two kinds. The first approach is *profile-centric* [15, 21]. All documents related to a candidate expert are merged into a single personal profile prior to retrieval time. The personal profiles are ranked w.r.t a query as single documents using standard retrieval measures and corresponding best candidates are returned to the user. The second approach, *document-centric*, is based on the analysis of individual documents. It runs a query against all documents and ranks candidates by summarized scores of associated documents [1, 9, 17] or text windows surrounding the person's mentioning [16]. It is also suggested not to finish propagation of scores on the level of directly related persons and to propagate the scores further through reciprocal document-candidate links [23]. Document-centric approaches are claimed to be much more effective than profile-centric [1], probably due to the fact that they estimate the relevance of the text content related to a person on the much lower and hence less ambiguous level.

A subfamily of document-centric methods exploits the social network built using links among persons extracted from top documents (e.g. by utilizing *from* and *to* fields of emails). The persons are ranked by popular centrality measures calculated on the acquired network. Campbell et al. [4] proposed the use of

HITS algorithm [13] which performed better than just ranking by candidate’s in-degree (related documents number). However, Chen et al. [10] found that a document-centric approach is still better than HITS based ranking. Query-independent experts discovery using links acquired from posts and replies at specialized forums was studied recently by Zhang et al. [27].

A number of advanced pseudo-relevance feedback techniques are consistently applied to the expert search task. Query expansion from the top retrieved documents performed quite well [2, 21]. Macdonald and Ounis also successfully experimented with different numbers of expansion terms received from the top ranked candidate profiles [18]. Serdyukov et al. [22] applied massive query expansion using the mixture of two pseudo-relevance language models: built on top ranked documents and top ranked profiles.

Expert finding is only a subcase of the entity ranking task. Generalization of search for other entity classes in the Web (*countries, cities, dates etc.*) is made by Zaragoza et al. [25] and Tsikrika et al. [24].

3 Person-centric expert finding

The key approaches to expert finding discussed in the previous section state that the level of personal expertise can be determined by the aggregation of document scores related to a person. As we show further, their intuition is based on measuring the co-occurrence degree of the query terms and personal id within the context of a document. In probabilistic terms, they suppose that our task comes to the estimation of the joint probability $P(e, q_1, \dots, q_k)$ of observing the candidate expert e together with query terms $q_1 \dots q_k$ in the documents ranked by the query. The methods which we describe here are graphically represented in Figure 1. We see that while in the typical document-centric method shown




Fig. 1. Dependence networks for two methods of estimating $P(e, q_1, \dots, q_k)$

on the left (see Section 3.1), the document is responsible for producing terms, in our method, shown on the right (see Section 3.2), the document requests a person to generate its terms. Below, we define these models formally.

3.1 Baseline approach

Let’s take a look at the well-known document-centric model by Balog and De Rijke (their Model 2) [1] using the principle shown in Figure 1a. According to

their approach, we have the following formulas for the total joint probability $P(e, q_1, \dots, q_k)$ over ranked documents set R :

$$P(e, q_1, \dots, q_k) = \sum_{D \in R} P(D)P(e, q_1, \dots, q_k|D) \quad (1)$$

$$P(e, q_1, \dots, q_k|D) = P(e|D) \prod_{i=1}^k P(q_i|D) \quad (2)$$

where $P(D)$ is a document prior, which is uniform. $P(e|D)$ is the probability of relation between person e and document D , calculated as:

$$P(e|D) = \frac{a(e, D)}{\sum_{i=1}^m a(e_i, D)}, \quad (3)$$

where m is a number of candidate experts in the system and $a(e, D)$ is a non-normalized association degree between the person and the document, which may depend on various factors: on the importance of the document part containing the person, on the number of occurrences of the personal identifier in the document, or on our confidence that a certain personal identifier found in the document matches person e .

The right part of the Equation (2) is a score of a document according to the language model based ranking principle [11], in which:

$$P(w|D) = (1 - \lambda_G) \frac{c(w, D)}{|D|} + \lambda_G P(w|G), \quad (4)$$

where $c(w, D)$ is the count of term w in document D , $|D|$ is its length, λ_G is the probability that term w will be generated from the global language model. $P(w|G)$ is the global language model estimated over the whole document collection.

As we may notice, this approach considers a candidate and the query terms to be conditionally independent given a ranked document (see Figure 1a). It is also similar to the popular query expansion method by Lavrenko [14] if only we consider the candidate expert as an expansion term. Since it's not only the most representative, but also the one of the most effective expert finding methods [1], it serves as a baseline in our experiments.

3.2 Putting persons in the middle

The person-centric method, which is the contribution of our paper, can be viewed as a hybrid method combining the features of both document- and profile-centric methods. It builds its prediction by analysing the top retrieved documents and summarizing the expertise evidence over them. However, the estimation of a personal language model (see next section) becomes a crucial step in this prediction.

Our approach is based on the assumption of dependency between the query terms and a candidate. We suppose that candidates are actually responsible for

the generation of terms within retrieved documents. According to the model presented in Figure 1b, we calculate the required joint probability as follows:

$$P(e, q_1, \dots, q_k) = \sum_{D \in R} P(q_1, \dots, q_k | e) P(e | D) P(D) = P(q_1, \dots, q_k | e) \sum_{D \in R} P(e | D) P(D) \quad (5)$$

where $P(q_1, \dots, q_k | e)$ is the probability of generating the query from the personal language model. It reflects the amount of relevant knowledge the candidate has. The sum in the right part of this formula can be considered as a person's prior $P(e)$:

$$P(e) = \sum_{D \in R} P(e | D) P(D), \quad (6)$$

which measures the influence/activity of the candidate in the topic area. It is proportional to the frequency of appearance of the candidate in the topical documents. We take a ranked document prior to be inversely dependent on the document rank: $P(D) = 1/\text{rank}(D)$ in order to distinguish the importance of a document in covering the aspects of the query topic. In our experiments we also show the performance with uniformly distributed $P(e) = 1/m$, where m is a number of candidate experts in the system.

We also consider that query terms occur independently given a candidate experts, what results in:

$$P(q_1, \dots, q_k | e) = \prod_{i=1}^k P(q_i | e) \quad (7)$$

Now we present our algorithm of mining for personal language models from the top retrieved documents.

4 Mining for personal language models

As we see, the personal query term generation probabilities $P(q_i | e)$ is the only part we miss so far. Of course, we can get them in the way similar to the one which profile-centric methods use: merge those retrieved documents that relate to the person e into one and calculate corresponding term frequencies. However, it would be justifiable if only there was only one person per document. Since we have already postulated that all candidates may be responsible for generating query terms in the documents they are mentioned in, such approach would give us only very rough approximation of a personal language model in most cases. Guided by these considerations, we represent a document as a mixture of personal models and the global language model. In formal terms, we define the likelihood of the top retrieved documents set R as:

$$P(R) = \prod_{D \in R} \prod_{w \in D} ((1 - \lambda_G) (\sum_{i=1}^m P(e_i | D) P(w | e_i)) + \lambda_G P(w | G))^{c(w, D)} \quad (8)$$

Here e_1, \dots, e_m are the persons occurring in the documents from R , $c(w, D)$ is the count of term w in document D , $(1 - \lambda_G)$ is a probability that a term will be generated from one of the personal models and not from the global language model. λ_G controls the ability of the algorithm to build personal models which are discriminative only for the terms which are topic-specific. Those terms which have high probability in the collection in total will get low generation probabilities over all persons.

Our approach to candidate experts modeling is based on the similar hypothesis with one used in pseudo-relevance feedback method for document retrieval by Zhai and Lafferty [26]. It also considers that the topical model of a user query can be mined from the top retrieved documents. The significant difference is that we define this model as a mixture of models of candidate experts. They are those who actually hold and share the knowledge which can meet the user information need. To say the truth, the personal language model which we get from top documents is only one of many the person uses. If we analyze the whole collection in the same way, we could get much more detailed personal term distribution. However, it would be much more difficult to distinguish candidate experts because the ambiguity of their expertise would increase dramatically in this case. Since we are interested only in the language model the person uses while generating documents that cover the query topic to some extent, it is reasonable to get it dynamically: at query execution time from retrieved documents. Our approach also shows some resemblance with Probabilistic Latent Semantic Indexing [12] with a distinction that our semantic topics are not 'latent', but personified and hence 'visible' in documents.

4.1 Using fixed personal contribution probabilities

Considering that all parameters, including $P(e_i|D)$ are given, we are able to calculate the maximum likelihood estimates of term generation probabilities from personal language models $P(w|e_i)$. In order to do that, we apply the EM algorithm [8], traditionally used to estimate unknown parameters. We propose the following formulas updating likelihood of the document set R (see Equation (8)) to be used recursively for its maximization:

E-step:

$$P(e|w, D) = \frac{(1 - \lambda_G)P(e|D)P(w|e)}{(1 - \lambda_G)(\sum_{i=1}^m P(e_i|D)P(w|e_i)) + \lambda_G P(w|G)} \quad (9)$$

M-step:

$$P(w|e) = \frac{\sum_{D \in R} c(w, D)P(e|w, D)}{\sum_w \sum_{D \in R} c(w, D)P(e|w, D)} \quad (10)$$

4.2 Measuring personal contribution

So far we relied on the assumption that probabilities $P(e|D)$ are fully determined by a person-document type of association. This practically means that if we have

some document with probability distribution $P(e|D)$, then for some another document with the same set of persons having the same kind of associations with it, the probability $P(e|D)$ will be distributed likewise. However, in our method we extract not only personal language models, but also probability distributions $P(e|w, D)$, which show who is the most probable generator of the term w in the document D . It allows us to estimate the probability of contribution for each person solely based on the document’s content.

For that purpose, we no more fix probabilities $P(e|D)$ and calculate them at every M-step of EM algorithm presented in Section 4.1 as follows:

$$P(e|D) = \frac{1 + \sum_{w \in D} c(w, D)P(e|w, D)}{m + \sum_{i=1}^m \sum_{w \in D} c(w, D)P(e_i|w, D)}, \quad (11)$$

where m is the number of candidate experts extracted from the retrieved documents in total, used here for the purposes of Laplace smoothing.

5 Experiments

5.1 Experimental setup

For the evaluation we utilize the W3C corpus - the data from the expert search task in the Enterprise track of the TREC used in 2005 and 2006 - and its largest (1.8 GB) ‘lists’ part containing discussions within the W3C consortium. We focus our experiments on this part of the collection for several reasons. At first, this data has a standardized format (emails of average length 450 words) what means that its properties should not change significantly across different enterprises. Moreover, it allows to accomplish persons recognition using unique email addresses and hence to avoid uncertainty in determining person-document relations. Since these email addresses always occur in a specific email field, we are able to differentiate the types of person-document relations as well. The data is parsed and indexed using Java and the Lucene open-search engine.

TREC also provided a list of 1092 candidate experts with supplemented full names, email addresses and unique ids. Experiments were conducted by considering only these candidates as person entities. We also tested inclusion of other person entities by taking any unique email found in the collection as a new person id. This caused only small degradation of performance, probably due to the rapid increase of noisy features with each new document retrieved, so we do not report these results here.

We provide results separately for two sets of TREC queries with relevance judgments: used in 2005 (50 queries) and in 2006 (49 queries). These query sets are somewhat different in nature. In 2005 queries were made up using names of working groups in W3C as titles and members of these groups as experts on the query topic. In 2006 the the TREC community manually judged each candidate for each query using the provided list of documents where a person id occurred. While queries from 2006 allow to reproduce a classic expert search scenario, queries from 2005 partly simulate the search for sub-groups within organization (a search for any person in the group working on the query topic problem).

5.2 Results discussion

First of all, we do candidates recognition by finding their email addresses in *from*, *to*, *cc* and *body* email fields. We additionally search for candidates in *body* fields using their full names. Association scores are $a(e, D^{from}) = 1.5$, $a(e, D^{to}) = 1.0$, $a(e, D^{cc}) = 2.5$ and $a(e, D^{body}) = 1.0$, what is the best combination according to recent studies of W3C 'lists' subcollection [3]. If a person appears in several fields, the highest association score is taken. The standard language model based IR approach, as defined in Equations (2) and (4), was used for the retrieval of documents.

We analyze the performance using the classic IR evaluation measures: Mean Average Precision (MAP), Mean Reciprocal Rank (MRR) and precision at top 5 ranked candidates (P@5). In our opinion, P@5 is more relevant to our task than precisions at greater ranks. The cost of a false recommendation in expert search is much higher than in document search: a conversation with an ignorant person or even reading all documents supporting the incorrect system's judgment takes much longer time than reading one irrelevant document. If we consider that the user can be satisfied with any single expert on the topic, than MRR becomes a decisive measure: it shows the ability of the system to present an expert as soon as possible if to go down by person's ranking one by one.

In order to demonstrate the quality of the mined personal language models (see Section 4), we start from presenting the performance of our methods considering that person's priors $P(e)$ are uniformly distributed and then using non-uniform priors, as in Equation (5), with the best of them. So, the following methods are evaluated:

- **Baseline**: the baseline document-centric method (see Section 3.1),
- **PCFix**: the person-centric method using fixed person-document association scores and uniform personal priors (see Sections 3.2 and 4.1),
- **PCUnf**: the person-centric method using unfixed dynamically calculated association scores and uniform personal priors (see Sections 3.2 and 4.2),
- **PCUnfNonUniPriors**: the person-centric method using unfixed dynamically calculated association scores and non-uniform personal priors (see Sections 3.2 and 4.2).

We have only two parameters in all models including the baseline model: λ_G , used in Equations (4) and (8), and the number of retrieved documents. Different values for λ_G between 0.1 and 0.9 showed negligible differences in performance, but 0.8 was slightly better than others. The second parameter was much more influential. It is always rather unclear how many top documents describe each query topic to the sufficient extent. So, a good algorithm should be robust to the size of a query result set. We vary its size from 1000 to 6000 of top ranked documents. We show MAP, P@5 and MRR values for both sets of queries in Figures 2, 3 and 4 respectively.

We see that the **PCFix** method performs similarly to the **Baseline** in average, except that it is notably better on P@5 for queries from 2005 (see Figure

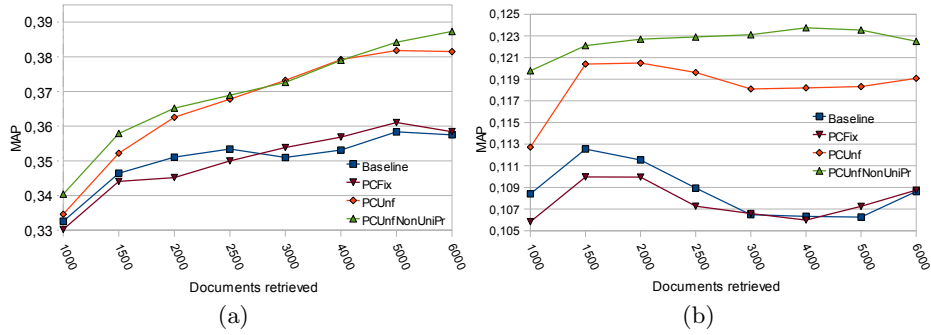


Fig. 2. MAP over different numbers of documents retrieved, for the queries from 2006 (a) and for the queries from 2005 (b)

3b). For other measures/queries, although it's better in half cases, it is worse in another half too.

However, the **PCUnf** method shows notably better performance than both the **Baseline** and the **PCFix** methods on all measures/queries, especially for MRR measure. It demonstrates that query-specific and purely content-based estimation of personal contribution to the document is crucial in personal language modeling.




Fig. 3. Precision at 5 over different numbers of documents retrieved, for the queries from 2006 (a) and for the queries from 2005 (b)

Moreover, using non-uniform priors $P(e)$, as in Equation 6, with the **PCUnf** method (the method **PCUnfNonUniPriors**) improves performance further for all MAP and P@5 measures at almost all numbers of retrieved documents. The frequency of participation in discussions on the topic is of course a significant evidence of personal expertise. However, from a statistical point of view, this prior penalizes the score of those candidates whose models are built using insufficient amount of training data, i.e. related documents. Both effects in total




Fig. 4. MRR over different numbers of documents retrieved, for the queries from 2006 (a) and for the queries from 2005 (b)

prevent incidental persons from getting high scores. However, using non-uniform priors spoils the performance of the **PCUnf** in case of MRR measure. So, if the user information need can be effectively satisfied with only one expert (and she is always available for requests), then the **PCUnf** is more preferable.

To sum things up, the presented results imply that our person-centric model is built on more realistic assumptions than the baseline document-centric model.

6 Conclusions and Further Work

We have presented the method for expert finding based on modeling of retrieved documents as mixtures of personal language models. Our approach assumed that terms in documents are generated by those persons who are mentioned in them. For the final ranking it combined two evidences of personal expertise: the probability of generation of the query by the personal language model and a prior probability of candidate experts expressing her level of activity in the important discussions on the query topic. We proposed two ways of personal models extraction from top ranked documents. In one case, we considered that person-document relation probabilities are fixed and fully depend on the field of a document where the person appeared. In another case, we obtained these probabilities dynamically by predicting the real contribution of persons to a document considering their intermediately calculated language models. When our method used this second way of modeling, it outperformed one of the best state-of-art approaches which we used as a baseline.

Several directions of improvement can be followed in the future. Certainly, the core person’s modeling part can be extended up to higher complexity. We may imagine that a person is a mixture of sub-persons representing different fields of her expertise. These inside experts can be used differently across documents and their probability of use may even depend on the set of other persons appearing in the document. A document can be also represented not only as a mixture of persons, but also as a mixture of global latent topics, which in turn appear to

be mixtures of persons, accumulating knowledge in the corresponding fields. Or we can even suppose that terms and persons are independent given such latent topic, which generates both these kinds of entities.

It is also reasonable to find more use of specific data properties. Particularly, we can consider that persons in the email document appear non-independently: the occurrence of persons in the *to* and the *cc* fields depends on the email sender in the *from* field, who is selecting them for communication. It is promising to take document links into account: for instance, by regarding emails relating to one thread as a single document.

References

1. K. Balog, L. Azzopardi, and M. de Rijke. Formal models for expert finding in enterprise corpora. In *SIGIR '06: Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 43–50, 2006.
2. K. Balog, T. Bogers, L. Azzopardi, M. de Rijke, and A. van den Bosch. Broad expertise retrieval in sparse data environments. In *SIGIR '07: Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 551–558, New York, NY, USA, 2007. ACM Press.
3. K. Balog and M. de Rijke. Finding experts and their details in e-mail corpora. In *15th International World Wide Web Conference (WWW2006)*, 2006.
4. C. S. Campbell, P. P. Maglio, A. Cozzi, and B. Dom. Expertise identification using email communications. In *CIKM '03: Proceedings of the twelfth international conference on Information and knowledge management*, pages 528–531, New York, NY, USA, 2003. ACM Press.
5. N. Craswell, A. de Vries, and I. Soboroff. Overview of the trec-2005 enterprise track. In *Proceedings of TREC-2005*, Gaithersburg, MD, 2005.
6. N. Craswell, D. Hawking, A.-M. Vercoustre, and P. Wilkins. Panoptic expert: Searching for experts not just for documents. In *Ausweb Poster Proceedings*, Queensland, Australia, 2001.
7. F. Crestani, M. Lalmas, C. J. V. Rijsbergen, and I. Campbell. "Is this document relevant? : Probably": a survey of probabilistic models in information retrieval. *ACM Comput. Surv.*, 30(4):528–552, 1998.
8. A. Dempster, N.M.Laird, and D.B.Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1):1–38, 1977.
9. H. Fang and C. Zhai. Probabilistic models for expert finding. In *Advances in Information Retrieval, 29th European Conference on IR Research, ECIR 2007, Rome, Italy, April 2-5, 2007, Proceedings*, pages 418–430, 2007.
10. J. X. S. T. Haiqiang Chen, Huawei Shen and X. Cheng. Social Network Structure behind the Mailing Lists: ICT-IIIS at TREC 2006 Expert Finding Track. In *Proceedings of the 15th Text REtrieval Conference (TREC 2006)*, 2006.
11. D. Hiemstra and F. M. G. de Jong. Statistical language models and information retrieval: Natural language processing really meets retrieval. *Glott international*, 5(8):288–293, 2001.
12. T. Hofmann. Probabilistic latent semantic indexing. In *SIGIR '99: Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*, pages 50–57, New York, NY, USA, 1999. ACM Press.

13. J. M. Kleinberg. Authoritative sources in a hyperlinked environment. *J. ACM*, 46(5):604–632, 1999.
14. V. Lavrenko and W. B. Croft. Relevance based language models. In *SIGIR '01: Proceedings of the 24th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 120–127, New York, NY, USA, 2001. ACM Press.
15. X. Liu, W. B. Croft, and M. Koll. Finding experts in community-based question-answering services. In *CIKM '05: Proceedings of the 14th ACM international conference on Information and knowledge management*, pages 315–316, New York, NY, USA, 2005. ACM Press.
16. W. Lu, S. Robertson, A. Macfarlane, and H. Zhao. Window-based Enterprise Expert Search. In *Proceedings of the 15th Text REtrieval Conference (TREC 2006)*, 2006.
17. C. Macdonald and I. Ounis. Voting for candidates: adapting data fusion techniques for an expert search task. In *CIKM '06: Proceedings of the 15th ACM international conference on Information and knowledge management*, pages 387–396, New York, NY, USA, 2006. ACM Press.
18. C. Macdonald and I. Ounis. Using relevance feedback in expert search. In *Advances in Information Retrieval, 29th European Conference on IR Research, ECIR 2007, Rome, Italy, April 2-5, 2007, Proceedings*, pages 431–443, 2007.
19. M. T. Maybury. Expert finding systems. Technical Report MTR06B000040, MITRE Corporation, 2006.
20. D. W. McDonald and M. S. Ackerman. Just talk to me: a field study of expertise location. In *CSCW '98: Proceedings of the 1998 ACM conference on Computer supported cooperative work*, pages 315–324, New York, NY, USA, 1998. ACM Press.
21. D. Petkova and W. B. Croft. Hierarchical language models for expert finding in enterprise corpora. In *ICTAI '06: Proceedings of the 18th IEEE International Conference on Tools with Artificial Intelligence*, pages 599–608, Washington, DC, USA, 2006. IEEE Computer Society.
22. P. Serdyukov, S. Chernov, and W. Nejdl. Enhancing expert search through query modeling. In *Advances in Information Retrieval, 29th European Conference on IR Research, ECIR 2007, Rome, Italy, April 2-5, 2007, Proceedings*, pages 737–740, 2007.
23. P. Serdyukov, H. Rode, and D. Hiemstra. University of Twente at the TREC 2007 Enterprise Track: Modeling relevance propagation for the expert search task. In *Proceedings of the 16th Text REtrieval Conference (TREC 2007)*, 2007.
24. T. Tsikrika, P. Serdyukov, H. Rode, T. Westerveld, R. Aly, D. Hiemstra, and A. de Vries. Structured Document Retrieval, Multimedia Retrieval, and Entity Ranking Using PF/Tijah. In *Proceedings of INEX 2007*, 2007.
25. H. Zaragoza, H. Rode, P. Mika, J. Atserias, M. Ciaramita, and G. Attardi. Ranking very many typed entities on wikipedia. In *CIKM '07: Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*, pages 1015–1018, New York, NY, USA, 2007. ACM.
26. C. Zhai and J. Lafferty. Model-based feedback in the language modeling approach to information retrieval. In *CIKM '01: Proceedings of the tenth international conference on Information and knowledge management*, pages 403–410, 2001.
27. J. Zhang, M. S. Ackerman, and L. Adamic. Expertise networks in online communities: structure and algorithms. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, pages 221–230, New York, NY, USA, 2007. ACM Press.