

A User-Oriented Model for Expert Finding

Elena Smirnova¹ and Krisztian Balog^{2,*}

¹ INRIA Sophia Antipolis - Méditerranée, France
elena.smirnova@inria.fr

² ISLA, University of Amsterdam, The Netherlands
k.balog@uva.nl

Abstract. Expert finding addresses the problem of retrieving a ranked list of people who are knowledgeable on a given topic. Several models have been proposed to solve this task, but so far these have focused solely on returning the most knowledgeable people as experts on a particular topic. In this paper we argue that in a real-world organizational setting the notion of the “best expert” also depends on the individual user and her needs. We propose a user-oriented approach that balances two factors that influence the user’s choice: time to contact an expert, and the knowledge value gained after. We use the distance between the user and an expert in a social network to estimate contact time, and consider various social graphs, based on organizational hierarchy, geographical location, and collaboration, as well as the combination of these. Using a realistic test set, created from interactions of employees with a university-wide expert search engine, we demonstrate substantial improvements over a state-of-the-art baseline on all retrieval measures.

1 Introduction

Expert finding addresses the task of identifying the right person with the appropriate skills and knowledge [6]. Experts can be required for a variety of purposes: problem solving, question answering, providing more detailed information on a topic, to name a few. The expert finding task has attracted a great deal of interest within the Information Retrieval (IR) community over the past few years [3, 4, 8, 11, 12, 28]. The main focus has been on developing content-based algorithms, similar to document search. State-of-the-art expertise retrieval algorithms identify experts based on the content of documents that they are associated with [6, 19, 26]. While these approaches have been very effective in finding the most knowledgeable people on a given topic, they abstract away from the actual user performing the search. Such abstractions are common in IR, especially at world-wide evaluation campaigns, like TREC, CLEF, or INEX, as they help to simplify the process of building reusable test collections. On the other hand, the abstraction of the expert finding task, as defined at the TREC Enterprise track [8, 11], ignores important aspects that may play a role when people locate

* Current affiliation: Department of Computer and Information Science, Norwegian University of Science and Technology, Norway, krisztian.balog@idi.ntnu.no

and select experts. Behavioral studies of human expertise seeking have found that, besides topical knowledge, the factors most influencing the selection process are related to the time needed to contact a person—including accessibility within the organizational hierarchy and workload [1, 21, 29].

In this paper we focus on two factors that influence the user’s choice: (i) time to contact a person and (ii) the person’s knowledge about the topic, relative to that of the user, seeking for expertise. We propose a user-oriented model that is based on *rational user behaviour* [2]; it assumes that, when choosing an expert, individuals will balance between the time needed to contact an expert and the knowledge gain. To the best of our knowledge, this is the first study that utilizes a user-centered approach for expert search. This approach provides a general framework for expertise retrieval that encompasses existing expert finding methods as a special case, namely, where the cost of contacting a candidate expert is ignored. We define knowledge gain as the relative difference between the expertise levels of the user and a candidate expert on a given topic; this is estimated using an existing state-of-the-art content-based approach. Contact time between two people—the user and a candidate expert—is approximated using their distance in a social network. Specifically, we consider social graphs based on (i) organizational hierarchy, (ii) geographical location, and (iii) collaboration (co-authorship), as well as the combination of these three structures.

We perform evaluation using real user queries and graded relevance judgments, obtained from the interactions of employees with a university-wide expert search engine [18]. We demonstrate substantial improvements on all retrieval measures with respect to a knowledge-only baseline. Our model involves a single parameter that can be used to adjust the user’s preferences for knowledge gain versus contact time. We explore a range of values for this parameter and find that configurations with slightly more weight on contact time perform best.

2 Related Work

Expert finding has been addressed from different viewpoints, including *expertise retrieval*, which takes a mostly system-centered approach, and *expertise seeking*, which studies related human aspects [16]. Besides textual data, *social networks* have also been used for finding experts. Finally, improving the user’s search experience by providing customized results touches on the field of *personalization*.

Expertise Retrieval. To reflect the growing interest in entity ranking in general and expert finding in particular, TREC introduced an expert finding task at its Enterprise track in 2005 [11]. At this track it emerged that there are two principal approaches to expert finding [3, 8, 11, 28]. The two models have been first formalized and extensively compared by Balog et al. [5], and are called *candidate* and *document* models, or *Model 1* and *Model 2*, respectively. Candidate-based approaches (also referred to as profile-based [14] or query-independent [23] methods) build a textual (usually term-based) representation of candidate experts, and rank them based on a query/topic, using traditional ad-hoc retrieval models. The other type of approach, document-based models (also referred to as

query-dependent approaches [23]), first find documents which are relevant to the topic, and then locate the associated experts. Nearly all systems that took part in the 2005–2008 editions of the Expert Finding task at TREC implemented (variations on) one of these two approaches. Building on either candidate or document models, further refinements to estimating the association of a candidate with the topic of expertise have been explored. For example, instead of capturing the associations at the document level, they may be estimated at the paragraph or snippet level [7, 20, 24]. Other extensions incorporate additional forms of evidence through the use of priors [14], document structure [33], hierarchical, organizational, and topical context and structure [6, 23], and Web data [25].

Expertise Seeking. Several studies have identified factors that may play a role in decisions of which expert(s) to select or recommend; most importantly, factors related to the person’s expertise level, social and physical proximity [1, 9, 21, 27]. In [21] the importance of the accessibility of an expert is pointed out—people prefer to contact those who are physically or organizationally close. Shami et al. [27] find that users prefer to contact the persons they know, even when they could receive potentially more information from people located outside their social network. Hofmann et al. [16] include additional factors besides topical knowledge (including organizational structure, position, experience, reliability, up-to-dateness, and contacts) into a system that recommends similar experts. Our work is different from [16] both in the task and in that we take the social distance between the user and the expert into account.

Social Networks. The use of social information for expert search has mainly been investigated from two directions. One uses graph-based measures (such as HITS or PageRank) on social networks, extracted from email communications, to produce a ranking of experts [10, 13]. Alternatively, others assume homogeneity among neighbours in a social network (based on co-authorship or organizational hierarchy) and define a smoothing procedure to relevance-based expert scores [17, 32]. It is important to mention that in this paper we consider not only people, but also other types of entities (organizational units and geographical locations) as nodes.

Personalized Search. Providing personalized results for users, depending on their information needs, has attracted significant interest in recent years. User modeling is one approach to search personalization; it can affect search in different phases: during the retrieval process, as a re-ranking step or in the query pre-processing procedure [22]. Adding a personalization component to the retrieval process can considerably slow down the system and query expansion does not guarantee to affect the result list. Re-ranking, on the other hand, remains a relatively cheap and reliable way to improve search quality; Google Personalized Search provides a large-scale example of this approach [30]. In this work we take a re-ranking approach to personalization for an expert finding system.

3 Baseline Model

Probabilistic approaches to expertise retrieval have been both popular and successful at the TREC Enterprise track [3, 8, 11, 28]. We take one of the most widely used models as our baseline, the so-called “Model 2” approach, proposed by Balog et al. [5]. Model 2 is attractive because it is an easy-to-understand, yet very effective method, with solid theoretical foundations. The key idea behind the model is to simulate how a user may search for experts using a standard document search engine: first, finding documents which are relevant, and then, examining each of these documents for associated persons. By scanning through a number of documents the user can obtain an idea of which people are more likely to be experts on the query topic [6].

Formally, the task of expert finding is stated as estimating the probability of an expert e to be knowledgeable on a given query topic q : $p(e|q)$. After applying Bayes’ rule, this problem transforms into estimating the probability of a query given an expert, $p(q|e)$; assuming uniform prior on experts, $p(e)$:

$$p(e|q) \propto p(q|e) \cdot p(e). \quad (1)$$

The generative language modeling technique naturally follows from here as a way to estimate the probability of query terms to be generated by the expert. Under the Model 2 approach, the degree of topical association between an expert and a query, $p(q|e)$, is estimated as a weighted sum of relevance scores of documents $p(q|d)$ related to an expert:

$$p(q|e) = \sum_d p(q|d) \cdot p(d|e). \quad (2)$$

Weights $p(d|e)$ express the degree of association between a document and an expert. Probabilities $p(q|d)$ are provided by the document language model, using the standard language modeling approach [31]. Since uniform priors were assumed, Eq. 2 provides the final ranking of expert candidates.

4 Experimental Setup

The main research question we aim to answer is this: Does modeling the user, seeking for experts, lead to improvements over a strictly content-based method?

Data Collection. The data set we use for experimental evaluation is the UvT Expert Collection [6]. It was collected from a publicly accessible database of employees involved in research or teaching at Tilburg University (UvT), The Netherlands. The data set comprises of four types of documents: research descriptions, course descriptions, publications, and personal homepages—in total 38,422 documents. The collection contains information about 1,168 experts, including their geographical location (building and room no.) and position within the organizational hierarchy. The UvT collection is bilingual (English/Dutch), but we did not resort to any special (language-dependent) treatment of the documents, apart from removing a standard list of English/Dutch stopwords.

Topics and assessments. In recent work, Liebrechts and Bogers [18] performed the evaluation of an expert finding system developed for UvT. As part of their study, 30 randomly selected UvT researchers were asked to formulate queries on which they themselves were knowledgeable. These participants were then asked to rate their and their colleagues’ relative expertise levels on this area on a five-point scale (from 0 = “no expertise” to 4 = “high expertise”). This resulted in a “user-centered” set of 30 queries and 268 graded and realistic relevance judgements; we use this topic set for experimental evaluation¹.

Evaluation Measures. We use standard Information Retrieval measures, namely, Mean Average Precision (MAP), Mean Reciprocal Rank (MRR), and precision at rank 5 (P@5). We also report on a standard preference measure, Normalized Discounted Cumulative Gain at the number of relevant experts (NDCG@R), which emphasizes the quality at the top of the ranked list.

Associating people and documents. To compute the baseline model (Eq. 2) it is necessary to estimate the degree of association between an expert and a document, $p(d|e)$. Since each UvT document describes a particular expert’s activity (research, teaching, publication), these associations can be established unambiguously. Following [6] we set $p(d|e)$ to 1 if person e is an author of document d , and to 0 otherwise. The baseline model was computed using the EARS toolkit².

5 User-Oriented Model

In this section we introduce our user-oriented model for expert finding that is based on *rational user behaviour* [2]. Under this approach, a “rational” user u is optimizing between the time needed to contact an expert e : $time(e|u)$, and the expertise level or knowledge of that expert relative to the expertise of the user (that is, the knowledge gain) given a particular query q : $knowledge(e|u, q)$. We combine these two factors in a mixture model:

$$score(e|u, q) = \lambda \cdot knowledge(e|u, q) - (1 - \lambda) \cdot time(e|u), \quad (3)$$

where the parameter $\lambda \in [0; 1]$ controls the balance between the two components and reflects the user’s preference in terms of expertise level versus contact time. We consider a range of values for λ in Section 6. Next, we discuss the two main components of our model—knowledge gain and contact time—in detail.

5.1 Knowledge Gain

This component expresses the level of expertise that user u could gain from expert e on a specific topic q . We estimate this value by considering the difference between their knowledge on topic q :

$$knowledge(e|u, q) = p(q|e) - p(q|u), \quad (4)$$

¹ Note that Liebrechts and Bogers [18] use a newer version of the UvT Expert Collection (with more experts and documents) that is not yet publicly available. Therefore, our retrieval scores are not directly comparable with those reported in [18].

² <http://code.google.com/p/ears/>

where $p(q|e)$ and $p(q|u)$ denote the level of expertise of e and u , respectively, on a given topic q ; these values are estimated using the baseline model (Eq. 2). Note that using this component alone for ranking (i.e., by setting λ to 1 in Eq. 3) would produce a ranking identical to that of the baseline model.

5.2 Contact Time

Intuitively, the contact time between a user and an expert could be estimated by their “social distance.” Indeed, shortest-path can be meant as the “most efficient” or “fastest” connection between nodes in a social network [15]. The more intermediaries that separate a user and an expert, the longer it takes for them to contact each other due to the weakening of ties. Based on the data that is available in the UvT collection, we constructed multiple social networks of researchers, induced by particular types of relationships: organizational, geographical, and co-authorship.

Formally speaking, we assume an undirected connected graph $G(V, U, E)$, where nodes V represent people (UvT employees), nodes U represent auxiliary nodes, and edges E reflect relations. The weight of an edge between nodes $v_i, v_j \in V \cup U$ is denoted by $w(v_i, v_j)$ and by default is set to 1. In a connected graph there exists a path $p = \{v_0, \dots, v_k\}$ between any two nodes. The weight of the path is computed as a sum of the weights of its constituent edges:

$$w(p) = \sum_{i=1}^k w(v_{i-1}, v_i). \tag{5}$$

We define the contact time between a user and an expert $time(e|u)$ to be the length of the shortest path between the corresponding network nodes, normalized by the diameter of the network:

$$time(e|u) = \min_{u \rightarrow e} w(p) / \max_{e, u} \min_{u \rightarrow e} w(p). \tag{6}$$

The proposed approach is very general, and can be applied to any graph that defines relationships between people. In the remainder of this section we present

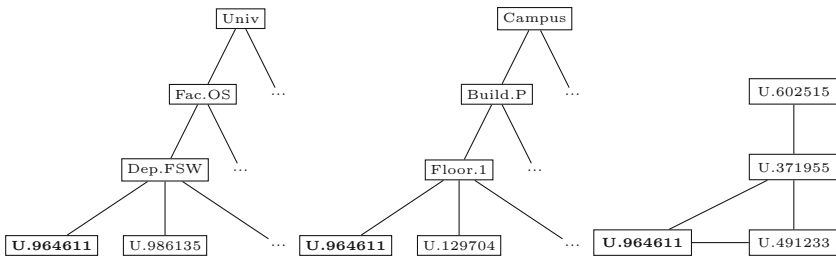


Fig. 1. Illustration of social network structures built from various relationships: (Left) organizational, (Middle) geographical, and (Right) co-authorship

the actual construction of three different types of networks; while these are discussed in the context of a specific organization—Tilburg University—, these types of structures are typically available in organizational settings.

Organizational network. Organizational network of the Tilburg University takes the form of hierarchy—see Figure 1 (Left). Therefore, in the graph $G(V, U, E)$, nodes U constitute to organizational units (departments, faculties, university), and edges E reflect organizational belonging. At the bottom level each individual belongs to one or more departments of the university. Then, departments are hierarchically embedded into corresponding faculties. Finally, faculties are linked to form the university. Employees with missing information about organizational position are considered to be inside the university and each belong to an artificial department and faculty as the only member (thereby, we set the maximum distance value to researchers without organizational information). This network defines a connected graph.

Geographical network. The geographical network also forms a hierarchy—see Figure 1 (Middle). In the graph $G(V, U, E)$, nodes U constitute to geographical units (floors, buildings, campus) and edges reflect geographical belonging. Researchers on the same floor are grouped at the lowest level of the hierarchy. Next, floors become a part of the corresponding building. Finally, the campus combines all buildings. Employee profiles with missing information about the individual’s location are considered to be inside the campus and belong each to an artificial building and floor as the only member. Since it is a hierarchy, the geographical network is connected.

Collaboration network. In this network there are no auxiliary nodes (the set U is therefore empty), links exist only between people (i.e., nodes in V); two researchers are connected with an edge if they co-authored a paper—see Figure 1 (Right). To ensure that the graph is connected, we add an edge between any pairs of disconnected nodes in V , and set its weight equal to the diameter of the largest connected component.

The standard properties of these networks are summarized in Table 1. We note that in a realistic setting a user makes use of different types of relationships at the same time in order to reach an expert. Therefore, we also considered combining different types of relations into one network. As a result, nodes in a combined network can be connected by mixed-type paths. Table 1 shows that combined networks indeed exhibit shorter average shortest paths³.

³ We also experimented with alternative ways of calculating the “social distance” between the user and an expert besides the length of the shortest path. Namely, the inverse value of the number of all paths connecting two nodes and the average length of these paths. The former distance favors larger number of paths in the network connecting the user and an expert and captures the premise that an increasing number of alternative paths increase the chances to reach the expert. The latter distance accounts for the fact that the user may use a non-optimal way to contact an expert. Our results showed that the shortest-path distance outperformed the above mentioned alternatives on all network configurations and for all retrieval measures.

Table 1. Statistical properties of social networks. Columns from left to right: total number of nodes in the network (Nodes), number of edges (Edges), maximum (Diam.) and average (AvSPath) shortest paths between two nodes in the network.

Network type	Nodes	Edges	Diam.	AvSPath
Org	1263	3210	6	4.76
Geog	1266	3033	6	4.92
Collab	1168	635	16	15.49
Org+Geog	1361	6243	6	4.09
Org+Collab	1263	3838	6	4.62
Collab+Geog	1266	3667	6	4.76
Org+Collab+Geog	1361	6870	6	3.921

6 Experimental Evaluation

In this section we present an experimental evaluation of our user-oriented model. This model involves a parameter λ that controls the interplay between the knowledge of an expert and contact time (see Eq. 3). In Table 2 we report retrieval performance for the baseline and for the user-oriented models, considering various types of networks (organizational, geographical, and collaboration), as well as their combinations. The scores we present here are achieved using optimized λ settings; we come back to the setting of this parameter in Section 7.1.

The results obtained show that the user-oriented model markedly outperforms the baseline; all measures are improved, significantly and substantially, independent of the type of the network used. The observed improvements confirm that considering the social distance between people leads to a much more accurate model of selecting experts in an organizational context, and that our approach can effectively incorporate this information into the retrieval model.

We find that the user-oriented model performs better on combined networks. The largest improvement against the baseline is witnessed for the Collab+Geog network; the relative performance increase is over 50% for all metrics.

Table 2. Comparison of retrieval performance. The optimal λ values are displayed in brackets. Significance is tested against the baseline using a two-tailed paired t-test; \dagger and \ddagger reflect significant changes for $\alpha = 0.01$ and $\alpha = 0.001$, respectively. Best results for each metric are in boldface.

	MAP	MRR	P@5	NDCG@R
Baseline	0.2419	0.5845	0.3067	0.3466
<i>User-oriented model</i>				
Org	0.3727 \ddagger ($\lambda=0.2$)	0.8944 \ddagger ($\lambda=0.4$)	0.4333 \ddagger ($\lambda=0.6$)	0.5321 \ddagger ($\lambda=0.4$)
Geog	0.3731 \ddagger ($\lambda=0.4$)	0.8944 \ddagger ($\lambda=0.4$)	0.4533 \ddagger ($\lambda=0.6$)	0.5549 \ddagger ($\lambda=0.4$)
Collab	0.3387 \ddagger ($\lambda=0.5$)	0.8944 \ddagger ($\lambda=0.3$)	0.4000 \ddagger ($\lambda=0.4$)	0.4944 \dagger ($\lambda=0.2$)
Org+Collab	0.3782 \ddagger ($\lambda=0.4$)	0.8944 \ddagger ($\lambda=0.4$)	0.4400 \ddagger ($\lambda=0.6$)	0.5331 \ddagger ($\lambda=0.4$)
Org+Geog	0.3809 \ddagger ($\lambda=0.2$)	0.8944 \ddagger ($\lambda=0.4$)	0.4400 \ddagger ($\lambda=0.6$)	0.5414 \ddagger ($\lambda=0.4$)
Collab+Geog	0.4035 \ddagger ($\lambda=0.4$)	0.8944 \ddagger ($\lambda=0.4$)	0.4800 \ddagger ($\lambda=0.4$)	0.5694 \ddagger ($\lambda=0.4$)
Org+Collab+Geog	0.3932 \ddagger ($\lambda=0.3$)	0.8944 \ddagger ($\lambda=0.4$)	0.4667 \ddagger ($\lambda=0.4$)	0.5446 \ddagger ($\lambda=0.4$)

Augmenting this network with information about the organizational hierarchy does not bring in further improvements (Collab+Geog vs. Org+Collab+Geog). We note that the best performing run is significantly different from all other runs in terms of MAP and also (with the exception of Org+Collab+Geog) in terms of NDCG@R. Interestingly, the best performance was almost always (with the exception of P@5) achieved by putting slightly more weight on the contact time aspect over knowledge (i.e., $\lambda < 0.5$).

7 Analysis

So far, we presented retrieval results using optimized parameter settings, averaged over all queries. Below, we analyze the sensitivity of our model w.r.t. its parameter, and take a closer look at differences on the level of individual topics.

7.1 Parameter Sensitivity Analysis

The parameter λ plays an important role in the user-oriented model: it regulates the balance between knowledge gain and time to contact an expert. Figure 2 shows retrieval scores for different values of this parameter. We display results for the three single-typed social networks (Org, Collab, Geog) and for the best combination (Collab+Geog); other combinations show very similar behaviour to that of Collab+Geog, but are not reported in the interest of readability.

It is clear from the plots that the combination always performs better than considering either knowledge ($\lambda = 1$) or contact time ($\lambda = 0$) alone. While there are differences in terms of absolute values, all network types exhibit very similar behavior given a metric. We also find that—with the exception of P@5—the optimal combination is shifted towards contact time relatively to the knowledge of an expert, and retrieval performance tops around $\lambda = 0.4$. In fact, we observe relatively flat curves when λ is in the range of $[0.1, 0.4]$; this indicates the stability and robustness of the model with respect to the choice of this parameter.

7.2 Topic-Level Analysis

We looked at results aggregated over all queries in Section 6, and now we continue our comparison by contrasting the performance of the baseline and user-oriented




Fig. 2. The effect of varying λ on retrieval scores using different types of social networks




Fig. 3. Topic-level differences between the baseline and the best performing user-oriented model in terms of (Left) AP and (Right) NDCG@R

Table 3. Topics displaying the largest difference compared to the baseline. Columns: user provided ranking (exp. level) with the corresponding position in the baseline (BL) and user-oriented (UO) rankings; geographical information: building (Build) and room (Room) number; whether an expert collaborated with the user (Collab.).

Expert ID	Build.	Room	Collab.	exp. level	ranked BL	ranked UO
521705	Y	232	No	3	232	1
270229	Y	231	No	3	12	2
171409	Y	252	No	2	25	3
198560	Y	238	No	2	67	16
289604	D	335	No	1	153	20

(a) Positive example #1: user #521705, query “philosophy of science.”

Expert ID	Build.	Room	Collab.	exp. level	ranked BL	ranked UO
523860	P	811	No	4	4	150
961051	K	1008	No	4	2	28
975900	K	1007	No	3	38	203
985430	K	1014	No	3	11	177
316326	K	1006	No	3	23	189
749796	D	405	No	3	4	43
329827	E	116	No	3	14	180

(b) Negative example #1: user #749796, query “marketing communications.”

Expert ID	Build.	Room	Collab.	exp. level	ranked BL	ranked UO
655248	P	1164	No	4	202	20
964611	P	1165	No	4	141	1
800353	P	1162	No	3	999	53
720437	P	3107	No	3	367	99
371955	P	1159	Yes	3	137	2
120146	P	3101	No	1	161	89
491233	P	3104	Yes	1	165	3
578111	P	1161	No	1	826	74

(c) Positive example #2: user #964611, query “customer satisfaction.”

Expert ID	Build.	Room	Collab.	exp. level	ranked BL	ranked UO
890847	M	315	No	4	3	99
265543	K	302	No	3	7	264
700940	K	415	No	3	31	470
968270	M	309	No	3	17	100
938920	P	1180	No	2	1	1
551309	P	1133	No	2	6	54
303267	M	312b	No	2	9	267
917125	K	316	No	2	26	467
859044	M	613	No	2	77	488

(d) Negative example #2: user #938920, query “sociale zekerheid.”

approaches on the level of individual queries. Figure 3 presents differences in average precision (AP) and NDCG@R scores against the baseline. It shows that the user-oriented approach considerably improved performance for more than half of the topics (17 for both AP and NDCG@R), for about a third it has remained the same (8 for AP and 9 for NDCG@R) and declined for the rest (5 for AP and 4 for NDCG@R). With the exception of a handful of topics (3 for AP and 2 for NDCG@R), the rate of decrease, however, is barely noticeable.

Next, we zoom in on four topic examples—two displaying the largest amount of performance increase against baseline, while the other two are taken from the opposite end of the spectrum. From the positive examples (Tables 3a and 3c) we see that almost all ranked experts belong to the same building and floor as the

user (the first digit of the room number indicates the floor). For these topics the user-oriented system correctly places colleagues sharing the same floor higher, as well as the two collaborators in the case of user #964611. In case of the negative examples, the user-oriented system again tried to promote people closest to the user. In one case (Table 3b) the user chose several experts, as most knowledgeable, mostly from a neighbouring building and located on the same floor; here, some personal connections between the user and those experts might play a role, but we do not have access to that type of social information from inside the organization. As to the other negative example (Table 3d), it appears that this particular user’s preference was highly shifted towards knowledge and away from proximity; addressing this issue is a question of setting the knowledge/time balance parameter λ in accordance with the user’s preference.

8 Conclusions and Future Work

In this work, we addressed the task of expert finding—ranking people with respect to their expertise on a given topic. We argued that, in a real-world organizational setting, the notion of the “best expert” depends on the individual user performing the search, and we proposed a user-oriented model that incorporates user-dependent factors. This model is based on the assumption that a user’s preference for an expert is balanced between the time needed to contact the expert and the knowledge value gained after. We defined the contact time between the user and an expert as their distance in a social graph, and examined different types of social relations that the user can engage in. Our approach provides a general framework for expert finding that encompasses existing approaches (which focus only on returning the most knowledgeable persons). We performed evaluation against a state-of-the-art baseline on the UvT Expert Collection using graded relevance judgements collected from real users, and demonstrated that the user-oriented approach significantly and substantially outperforms the baseline for all retrieval measures. We completed our investigation with parameter sensitivity examination and a topic-level success and failure analysis.

In future work, we plan to complement our current approach with automatic parameter learning. We also wish to extend this work with user studies that explore additional user-related factors that may play a role in expert finding.

Acknowledgments. Balog was supported by the Center for Creation, Content and Technology (CCCT). Smirnova was supported by INRIA, CORDI fellowship and AxIS team via the EC funded Elliot project. Smirnova thanks K. Avrachenkov and B. Trousse for their valuable comments and feedback.

References

- [1] Ackerman, M., Wulf, V., Pipek, V.: *Sharing Expertise: Beyond Knowledge Management*. MIT Press, Cambridge (2002)
- [2] Anderson, J.: *The adaptive character of thought*. Lawrence Erlbaum Assoc., Mahwah (1990)

- [3] Bailey, P., Craswell, N., de Vries, A.P., Soboroff, I.: Overview of the TREC 2007 enterprise track. In: *The Sixteenth Text REtrieval Conference Proc.* (2008)
- [4] Balog, K.: The SIGIR 2008 workshop on future challenges in expertise retrieval (fCHER). *SIGIR Forum* 42(2), 46–52 (2008)
- [5] Balog, K., Azzopardi, L., de Rijke, M.: Formal models for expert finding in enterprise corpora. In: *SIGIR 2006*, pp. 43–50 (2006)
- [6] Balog, K., Bogers, T., Azzopardi, L., de Rijke, M., van den Bosch, A.: Broad expertise retrieval in sparse data environments. In: *SIGIR 2007*, pp. 551–558 (2007)
- [7] Balog, K., Azzopardi, L., de Rijke, M.: A language modeling framework for expert finding. *Inf. Processing and Management* 45(1), 1–19 (2009)
- [8] Balog, K., Soboroff, I., Thomas, P., Craswell, N., de Vries, A.P., Bailey, P.: Overview of the TREC 2008 enterprise track. In: *TREC 2008* (2009)
- [9] Borgatti, S., Cross, R.: A relational view of information seeking and learning in social networks. *Management Science* 49(4), 432–445 (2003)
- [10] Campbell, C.S., Maglio, P.P., Cozzi, A., Dom, B.: Expertise identification using email communications. In: *CIKM 2003*, pp. 528–531 (2003)
- [11] Craswell, N., de Vries, A.P., Soboroff, I.: Overview of the TREC-2005 Enterprise Track. In: *The Fourteenth Text REtrieval Conference Proceedings* (2006)
- [12] CriES. Cross-lingual Expert Search workshop at CLEF (2010), <http://www.multipa-project.org/cries>
- [13] Dom, B., Eiron, I., Cozzi, A., Zhang, Y.: Graph-based ranking algorithms for e-mail expertise analysis. In: *Proceedings of the 8th ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery*, pp. 42–48 (2003)
- [14] Fang, H., Zhai, C.: Probabilistic models for expert finding. In: Amati, G., Carpineto, C., Romano, G. (eds.) *ECIR 2007*. LNCS, vol. 4425, pp. 418–430. Springer, Heidelberg (2007)
- [15] Hanneman, R., Riddle, M.: *Introduction to social network methods*. Cambridge University Press, Cambridge (2005)
- [16] Hofmann, K., Balog, K., Bogers, T., de Rijke, M.: Contextual factors for finding similar experts. *Journal of the American Society for Information Science and Technology* 61(5), 994–1014 (2010)
- [17] Karimzadehgan, M., White, R.W., Richardson, M.: Enhancing expert finding using organizational hierarchies. In: Boughanem, M., Berrut, C., Mothe, J., Soule-Dupuy, C. (eds.) *ECIR 2009*. LNCS, vol. 5478, pp. 177–188. Springer, Heidelberg (2009)
- [18] Liebrechts, R., Bogers, T.: Design and evaluation of a university-wide expert search engine. In: Boughanem, M., Berrut, C., Mothe, J., Soule-Dupuy, C. (eds.) *ECIR 2009*. LNCS, vol. 5478, pp. 587–594. Springer, Heidelberg (2009)
- [19] Macdonald, C., Ounis, I.: Voting for candidates: adapting data fusion techniques for an expert search task. In: *CIKM 2006*, pp. 387–396 (2006)
- [20] Macdonald, C., Hannah, D., Ounis, I.: High quality expertise evidence for expert search. In: Macdonald, C., Ounis, I., Plachouras, V., Ruthven, I., White, R.W. (eds.) *ECIR 2008*. LNCS, vol. 4956, pp. 283–295. Springer, Heidelberg (2008)
- [21] McDonald, D.D., Ackerman, M.: Just talk to me: a field study of expertise location. In: *CSCW 1998*, pp. 315–324 (1998)
- [22] Micarelli, A., Gasparetti, F., Sciarone, F., Gauch, S.: Personalized search on the world wide web. In: *The Adaptive Web: Methods and Strategies of Web Personalization*, pp. 195–230. Springer, Heidelberg (2007)
- [23] Petkova, D., Croft, W.B.: Hierarchical language models for expert finding in enterprise corpora. In: *ICTAI 2006*, pp. 599–608 (2006)

- [24] Petkova, D., Croft, W.B.: Proximity-based document representation for named entity retrieval. In: CIKM 2007, pp. 731–740 (2007)
- [25] Serdyukov, P., Hiemstra, D.: Being omnipresent to be almighty. In: SIGIR 2008 Workshop on Future Challenges in Expertise Retrieval, pp. 17–24 (2008)
- [26] Serdyukov, P., Hiemstra, D., Fokkinga, M.M., Apers, P.M.G.: Generative modeling of persons and documents for expert search. In: SIGIR 2007, pp. 827–828 (2007)
- [27] Shami, S., Ehrlich, K., Millen, D.: Pick me!: link selection in expertise search results. In: CHI 2008, pp. 1089–1092 (2008)
- [28] Soboroff, I., de Vries, A., Crawell, N.: Overview of the TREC 2006 Enterprise Track. In: The Fifteenth Text REtrieval Conference Proceedings (2007)
- [29] Woudstra, L.S.E., Van den Hooff, B.J.: Inside the source selection process: Selection criteria for human information sources. *Information Processing and Management* 44, 1267–1278 (2008)
- [30] Zamir, O., Korn, J., Ficks, A., Lawrence, S.: Personalization of placed content ordering in search results. Google Inc., Assignee (2005)
- [31] Zhai, C., Lafferty, J.: A study of smoothing methods for language models applied to information retrieval. *ACM Trans. Inf. Syst.* 22(2), 179–214 (2004)
- [32] Zhang, J., Tang, J., Li, J.: Expert finding in a social network. In: Kotagiri, R., Radha Krishna, P., Mohania, M., Nantajeewarawat, E. (eds.) DASFAA 2007. LNCS, vol. 4443, pp. 1066–1069. Springer, Heidelberg (2007)
- [33] Zhu, J., Huang, X., Song, D., Ruger, S.: Integrating multiple document features in language models for expert finding. *Knowl. and Inf. Systems* 23(1), 29–54 (2010)