

Mining Expertise and Interests from Social Media

Ido Guy, Uri Avraham, David Carmel, Sigalit Ur, Michal Jacovi, Inbal Ronen

IBM Research Lab

Haifa, Israel

{ido, uria, carmel, sigalit, jacovi, inbal}@il.ibm.com

ABSTRACT

The rising popularity of social media in the enterprise presents new opportunities for one of the organization's most important needs—expertise location. Social media data can be very useful for expertise mining due to the variety of existing applications, the rich metadata, and the diversity of user associations with content. In this work, we provide an extensive study that explores the use of social media to infer expertise within a large global organization. We examine eight different social media applications by evaluating the data they produce through a large user survey, with 670 enterprise social media users. We distinguish between two semantics that relate a user to a topic: expertise in the topic and interest in it and compare these two semantics across the different social media applications.

Categories and Subject Descriptors

H.3.3 [Information Search and Retrieval]: search process; H.5.3 [Group and Organizational Interfaces]: Computer-supported cooperative work

General Terms

Experimentation; Human Factors; Measurement

Keywords

Enterprise; enterprise 2.0; expert finding; expert search; expert recommendation; expertise location; interest mining; people search; social business; social computing; social media; social software; web 2.0

1. INTRODUCTION

Expertise location aims at finding people knowledgeable in a given topic, typically within an organization. This challenge has been fertile ground for research over the years. Early approaches were based on explicit input from individuals about their own skills and expertise [28]. To avoid the extra burden of manually filling in and maintaining expertise profiles, most methods in the past two decades focused on implicit inference. Expertise is usually mined based on documents such as project descriptions, human resource databases, professional articles, program code, or the employee's own files (e.g., [7, 34, 36, 42]). Email has also become a popular source for expertise mining since it is used so often to communicate about work topics [5, 9].

The emergence of social media in the enterprise has introduced many new types of applications that allow users to share and interact, including blogging systems [44], wikis [8], communities [32], social bookmarking [31], microblogging [46], file sharing [41], and people tagging [13]. These applications encourage users to contribute and participate by commenting, tagging, joining, sharing, 'liking', and authoring documents. The diversity of both the content types and the user associations with content suggests that expertise information derived from social media data can be of great value. Social media also holds other advantages for expertise mining. The content and associations are typically public, as opposed to email or local files, which are personal and thus more sensitive to privacy issues [36]. The data is dynamic in nature, reflecting changes in user activity over time. It also reflects the network structure based on relations among people, in addition to the relation of people to content [14, 15]. Moreover, social media that resides behind a firewall is typically used by employees to discuss internal topics, and hence reflects the organization's unique vocabulary and areas of interest.

In this work, we explore the use of different enterprise social media applications as data sources for expertise mining. As social media becomes more prevalent, we believe that a better understanding of how expertise can be derived from social media data is vital and can contribute to the overall value-proposition of social media in the enterprise. To the best of our knowledge, this is the first study to provide a comprehensive comparison among social media applications as data sources for expertise location.

Our evaluation is based primarily on a large-scale user survey within our organization, which includes 670 participants, whose input serves as the basis for comparison. We explore expert mining over a variety of enterprise social media applications and compare the retrieved results based on participants' self-ratings of their expertise level for 342 topics in total. Our data consists of documents originating from blogs, wikis, forums, bookmarks, microblogs, communities, shared files, and people tags. We also examine different types of user associations with documents, such as author, member, commenter, 'liker', and sharer. We observe that a user can generally be related to a topic based on two semantics: 1) the user is an *expert* in the topic and 2) the user is *interested* in the topic. While expertise reflects the traditional semantics of being knowledgeable or skilled in the topic, interest reflects a weaker association with the topic, indicating curiosity and a desire to learn more. In our survey, we did not explicitly define expertise or interest, but let participants form their own interpretation. Each participant was asked two questions about a given topic: 'what is your level of expertise in the topic?' and 'what is your level of interest in the topic?' As far as we know, no previous research has explicitly compared between the two semantics.

Our findings indicate that the lists of individuals retrieved from the different social media applications are very diverse. As expected, self-ratings of topics based on interest are significantly higher in general than ratings based on expertise. Profile tags and communities return results that have the highest agreement with the self ratings, while files and bookmarks cover more experts, and microblogs yield the best combination of the two. Aggregation of data originating from multiple social media applications further improves the overall results. Surprisingly, authorship association with a document does not imply greater expertise in the document's topic than commenting on it or liking it, indicating that in social media, the individuals who provide feedback on a document may be just as expert as the authors.

The remainder of this paper is organized as follows. In the following section, we review related work. We then present our experimental setup, followed by a description of the results, comparing the different social media applications as data sources for expertise and interest mining. We conclude by extensively discussing our findings and suggesting future work.

2. RELATED WORK

Expertise location—sometimes referred to as expert finding, expert search, or expert recommendation—has been widely studied, especially by the human-computer interaction (HCI) and information retrieval (IR) communities. Yiman-Seid and Kobsa [45] identified two motives for seeking an expert: as a source of information and as someone who can perform a social or organizational role. Ehrlich and Shami [12] further enumerated four motives: getting answers to technical questions, finding people with specific skills, gaining awareness of “who is out there”, and providing information.

Many of the studies describe the use of expertise location systems within specific organizations, relying on data sources that were available before the social media era. For example, Expert Seeker [7] was used to locate experts within the NASA organization, relying on a human resource database, an employee performance evaluation system, a skills database, and a project resource management system. The Lotus Knowledge Discovery System [34] was based on a mix of corporate documents from organizational databases, files on the intranet, and external websites, to provide knowledge discovery and expertise location inside IBM. ExpertFinder matched novices and experts by analyzing Java files related to them [42] and Expertise Recommender [30] presented a general architecture for expertise location based on work products and byproducts, which was analyzed through a field study within a medium-size software company. Email has also been inspected in several studies (e.g., [5, 9]). Its massive use by the workforce makes it a rich source for inferring expertise, but its main drawback is the sensitivity to privacy issues. Reichling and Wulf [36] provide a good summary of the study of expertise location systems in practice and add a case study of their own within a large industrial organization.

To address the growing interest in expertise location within the IR community, TREC introduced an expert finding task at its Enterprise track in 2005 [11]. The task was based on the W3C corpus containing people's association to W3C workgroups and a diverse set of documents crawled from the W3C website. Many of the IR studies that followed were based on this collection. Balog et al. [4] defined and compared two models: the *candidate-based* model, which builds profiles for candidate experts and ranks them based on their similarity with the query; and the *document-based*

model, which first finds documents relevant to the query and then locates the associated experts based on these documents. In their evaluation, the document model consistently outperformed the candidate model. Macdonald and Ounis [25] referred to document-based expert search as a voting problem, where documents vote for candidates with expertise relevant to the query. They explored several ways to combine these votes to rank the candidate experts. Serdyukov et al. [39] suggested a graph-based approach for expert finding in large enterprises. The graph was built based on different associations among candidate experts and web documents. In this work, we apply a document-based approach to compare the various social media data types as sources for expertise location.

The rise of social media brought about studies that examined knowledge sharing and expert location through specific social media applications. Examples include forums [20], communities [47], blogs [6], collaborative tagging [33], and collaborative question answering systems [2]. All of these studies examined social media outside the firewall. A few others referred to particular types of social media applications inside the enterprise. Kolari et al. [21] presented an application for expertise location over corporate blogs that utilized the content of the blog posts, their tags, and comments. No evaluation was provided. Millen et al. [31] argued that one of the benefits of social bookmarking in the enterprise is increased awareness of the interests and expertise of other corporate employees. Amitay et al. [3] presented a unified approach that allowed searching for documents, people, and tags in the enterprise. Data was derived from applications for social bookmarking and blogging, but the two data sources were not compared and the system was evaluated as a whole. Our work focuses on comparing a wide variety of enterprise social media applications as data sources for expertise inference.

Expertise location in the enterprise is part of the broader domain of enterprise people search [16], defined as any search within the enterprise in which the returned entities are people. Enterprise people search, in its commonly used form, is a navigation task: the users know which person they are looking for and try to reach a piece of information about that person, such as the contact details, organizational environment, job title, or photo. This is done using a query that includes cues such as the first name and the first letter of the last name, or the last name and the organizational unit (see [16] for more details). In contrast, expertise location is a discovery task: the users seek to find the people related to a certain topic, with the topic being used as the query.

In this work, we make the distinction between a person's expertise and her interests. Mining user interests has been mostly studied by the user modeling and recommender system communities. Typically, a user model is built based on her interests, and is used for personalization and recommendation purposes. Traditional methods employed explicit user feedback, usually in the form of rating a set of items, to extract the user's interests. To avoid this extra burden on the user, leveraging implicit interest indicators [10], such as purchase history, views, clicks, or queries, has become the prevalent approach in recent studies, for example to measure user similarity [43] or personalize search results [35]. The emergence of social media introduced new forms of data and metadata that may indicate user interests, such as tags [23, 38], comments, votes ('likes') [22], or short microblog messages [1]. While the focus of mining user interests has been for

personalization purposes, in this work we examine interest from a global perspective, referring to it as a weaker form of expertise.

3. EXPERIMENTAL SETUP

3.1 Social Media Applications

We experimented with an enterprise social software application suite [18], which has been deployed in our organization for over three years and includes eight types of social media applications:

- A blogging system that allows users to write blog posts, comment on their own or other posts, and ‘like’ other posts or comments.
- A social bookmarking system that allows users to share bookmarks of both intranet and Internet pages, and annotate them with tags.
- A system for online communities, where users can create and join communities of interest. As part of these communities, other content can be shared, including blogs, bookmarks, files, forums, and wikis.
- A file sharing system that allows users to upload files they authored. Users can share a file with others even if they are not the authors, download the file, comment on it, or ‘like’ it.
- A forum system that allows users to open threads about different topics and comment on their own or other threads.
- A microblogging system that allows users to write short messages of up to 500 characters, either on their own profile page (“board”) or on others’.
- A people tagging system where users can annotate each other with descriptive tags.
- A wiki system that allows multiple users to co-author wiki pages and comment on these pages.

3.2 Data Modeling and Scoring

Our social media data is stored in a unified index [3], which is built on top of Lucene [29] and maps the relationships among people, documents, and terms. In the unified approach, all types of entities (people, documents, terms) are searchable and retrievable. In our case, the search query consists of terms, representing a specific topic, while the retrieved entities are people related to that topic. We apply a document-centric approach, where the retrieval of relevant experts is based on first retrieving documents that are similar to the query and then scoring related people by aggregating their association weights with those documents [4].

Table 1 details the applications we experimented with, the entity represented by a document in the index (e.g., a blog post, a wiki page), and the set of associations by which users can relate to the document. For the bookmark application, we modeled two types of documents to represent both the association of the user with the whole bookmarked web page, including its content and metadata, and the user’s association with the tags she used in the bookmarking application. For files, the sharer association refers to all files the user shared with others; sharedWith refers to all files that were shared with the user by others. For microblogs, boardOwner refers to all messages written on the profile page of the user.

Table 1 also provides an overview of the usage level of the different applications in our organization: the numbers in

Table 1. Social media applications, documents, and associations

Application	Document	Associations	#People
Blogs	Post (54,414)	Author (54,414) Commenter (22,265) Liker (13,816)	25,886
Bookmarks	Bookmark (980,693)	Bookmarker (980,693)	22,301
	Tag (3,029,791)	Tagger (3,029,791)	
Communities	Community (20,606)	Member (1,485,412)	290,237
Files	File (124,599)	Author (124,599) Commenter (8,620) Liker (3,125) Sharer (21,417) SharedWith (86,136) Downloader (298,149)	62,088
Forums	Thread (75,153)	Author (75,153) Commenter (397,765)	23,377
Microblogs	Message (353,554)	Author (353,554) BoardOwner (353,554)	36,430
Profile Tags	Profile Tag (23,847)	Tagger (45,265) TaggedWith (276,020)	50,671
Wikis	Page (191,493)	Author (1,479,426) Commenter (6,641)	21,053

parentheses indicate the amount of documents and associations of each type. The rightmost column indicates, per application, the number of unique people who are “covered” by it, i.e., the number of employees who are associated with at least one document (there are 400,000 employees in our organization overall). It can be seen that communities cover the most substantial amount of individuals. Since anyone can add members to their community, many individuals in our organization are members of at least one community. Other applications that allow users to associate other individuals with a document also cover more people: files allow sharing a file with others, profile tags allow tagging another person, and microblogs allow writing on the board of another person. Overall, each application covers at least 20,000 individuals, which provides a good basis for comparison.

We indexed all content and metadata available for each document type: for blog posts, forum threads, and wiki pages, we indexed both their content and metadata (title, description, tags); for a file, the metadata and actual content of the file; for a bookmark, the metadata and actual content of the web page; for a microblog, its content (no metadata exists); for a community, we indexed only its metadata (the content may consist of other blogs, bookmarks, wikis, forums, or files) and for a tag, just its own text.

Given a query, representing a specific topic, the relevant individuals were retrieved through the documents that are found relevant for the query. These individuals were scored using a standard voting mechanism [25], according to the following formula that scores a person p for query q :

$$S(q, p) = idf(p) \cdot \sum_{d \in D(q)} [e^{-\alpha T(d)} \cdot Pop(d) \cdot S(q, d) \cdot W(d, p)]$$

where $D(q)$ is the collection of all documents retrieved for query q ; $T(d)$ is the number of days passed since document d was last updated and α is a decay factor; $Pop(d)$ is a score describing the relative popularity, or authority, of d , which is based on the amount of feedback (comments, tags, likes) it contains [3]. These first two factors in the summed product can be considered as static (query-independent) scores of document d . $S(q, d)$ is the cosine similarity score between the query and the document as calculated by Lucene. $W(d, p)$ is a document-person association weight, determined based on the type of document and the association(s) between the document and the person. By default, it is set to the number of associations of p with d . When aggregating data from multiple social media applications, it is also used to assign different weights to documents based on their originating application. Finally, $idf(p) = \ln(N/Np)$ is the query-independent *inverse document frequency* of person p [3], computed as the logarithm of the ratio between the number of all documents in the system (N) and the number of documents associated with p (Np). Similar to the vector-space *idf* score for terms [27], the *idf* score for people penalizes individuals who are related to many documents in general, and are hence less specifically related to a given query.

The unified index allows flexible querying based on different types of user-document associations. We could therefore use it to retrieve relevant people for a given query, based on any type of social media application (e.g., all people related to ‘ajax’ based on forums), or based on a specific association (e.g., all people related to ‘user experience’ based on blog post authorship). We also experimented with searching over an aggregate of all eight applications and their respective associations as listed in Table 1.

3.3 User Survey

We chose topics for our experiments from the query log of a social search engine used in our organization to search for people, documents, and tags [37]. We randomly sampled 1,200 queries that were issued at least 5 times in the year that preceded our experiment. We filtered out duplicates and queries that do not fit expertise location, such as names of business units or brands, administrative processes, job titles, locations, or people’s names. The rest of the queries, 411 in total, included themes or domains (e.g., mobile, social media, user experience, customer segmentation), technologies (Dojo, faceted search, Perl), internal and external tools (Phonebook, Dropbox, Camtasia), companies (Verizon, Unica), products (Websphere Application Server, iPad), and internal projects.

We then extracted from the unified index (described in the previous section) 1,563 users that had at least 25 queries, out of the 411, to which they were related by at least one of the associations listed in Table 1, considering the top 100 individuals retrieved for each pair of query and association. For each of these users we selected at random exactly 25 such queries. In our survey, we asked participants two questions with regards to each of the 25 topics selected for them: 1) What is your level of expertise in the topic? and 2) what is your level of interest in the topic? Participants rated their replies on a 5-Likert scale, where 1 indicated ‘none’ and 5 indicated ‘very high’.

We sent invitations to participate in the survey to the 1,563 individuals via email. We received responses from 670 who provided a total of 16,750 topic ratings with respect to both their expertise and interest in each. Our participants originated from 35 countries, spanning the different organizational units: 30% sales, 25% software, 20% services, 9% headquarters, 5% systems, 5% operations, 3% research, and 3% others. This sample does not perfectly represent the entire population of our organization’s employees, but rather active users of enterprise social media, for whom rich data can be produced for expertise inference. However, we believe that the results of our study fairly reflect the potential value of expertise mining from social media, as its usage is likely to grow when more organizations deploy and encourage the use of social media applications behind their firewalls.

3.4 Result Quality Evaluation Measures

In the main part of our evaluation, we compared the results returned by different social media applications based on the survey ratings. The individuals retrieved by the unified index for a given application are only partially covered by the participants of our survey and may include many individuals who did not provide feedback for that topic. In our result quality analysis, we only considered queries for which at least 10 participants provided feedback, reducing the list of queries to 342 in total.

Due to the rating sparsity, common evaluation methods for ranking algorithms based on explicit relevance feedback were not applicable. Our experimentation with NDCG [19] indicated that due to the large number of missing ratings among the top k results, the ideal ordering, used to normalize the NDCG score, is hard to predict. NDCG is known to have difficulties when only partial relevance feedback exists [27]. We also experimented with *precision@k* [27], assuming a “good” result is rated either 4 or 5. However, since negative feedback (1 or 2) was not taken into account at all, applications that received more feedback in our survey were favored. We therefore opted to apply a measure that would consider all available ratings among the top k results and would distinguish among all 5 values on the Likert scale.

To this end, we define the *average rating at k* for a query q , denoted $AR@k(q)$, as the average rating of all results among the top k retrieved individuals for which we have feedback (i.e., the individuals who rated themselves with respect to the query). The *mean average rating at k*, $MAR@k$, is defined as the mean of $AR@k(q)$ over all queries, while only queries with at least k results, of which at least one has been rated, are considered. We only report $MAR@k$ results when relying on at least 100 queries with feedback, to ensure a broad enough sample for averaging. $MAR@k$ disregards missing feedback and thus practically assumes it is identical to the average feedback received for the query across the top k results. We also experimented with a variant of $MAR@k$ that sets any missing feedback to 3, thus assuming neutrality for individuals with no feedback.

The $MAR@K$ measure reflects the precision of the results; however, it does not distinguish between applications according to the number of experts they cover (their recall) for a given topic. To complement the picture, we also use the *recall@k* measure. Given a query q , the *recall@k(q)* is defined as the portion of individuals who rated themselves 4 or 5 among the top k results for q , out of the overall number of individuals who rated themselves 4 or 5 for q . *recall@k* is defined as the mean of *recall@k(q)* over all queries. The *recall@k* measure is affected by

Table 2. Mean *Match@100* values among the 8 applications

	microblogs	blogs	bookmarks	files	profile tags	wikis	communities	forums
microblogs		13.4	7.9	6.0	6.5	3.1	2.4	2.1
blogs	13.4		7.0	7.3	4.7	4.6	2.4	2.2
bookmarks	7.9	7.0		7.9	3.6	1.7	1.5	1.0
files	6.0	7.3	7.9		2.7	3.6	2.2	0.6
profile tags	6.5	4.7	3.6	2.7		2.9	3.4	1.6
wikis	3.1	4.6	1.7	3.6	2.9		2.4	1.3
communities	2.4	2.4	1.5	2.2	3.4	2.4		0.9
forums	2.1	2.2	1.0	0.6	1.6	1.3	0.9	
average	5.9	5.9	4.4	4.3	3.6	2.8	2.2	1.4

the number of overall results returned by the application, since if an application returns only few results, it cannot reach a high value of recall even if it has perfect precision.

4. EXPERIMENTAL RESULTS

4.1 Overlap Analysis

We first examine the overlap among the results returned by our different social media applications (regardless of the survey input). This analysis strives to understand the diversity of our data and the potential value of collecting and aggregating expertise data across a wide variety of social media applications. To this end, we use the *match@100* measure [15], which calculates the percentage of people co-occurring at the top 100 results returned by any pair of the eight social media applications for a specific query. We consider the mean *match@100* over all queries that return at least 100 results for both applications; each pair of applications had at least 100 such queries. Calculating the match at values other than 100 showed very similar results. We therefore report only the *match@100* results for simplicity of presentation.

Table 2 shows the mean *match@100* over the 411 queries for each pair of applications. Higher overlap values are indicated with stronger shades of blue. The bottom row shows the average overlap of each of the applications over the other seven applications. Blogs and microblogs have the highest average overlap (5.9% each) and they are also the pair with highest overlap out of all pairs (13.4%). Communities and forums are the most distinct, showing low overlap with the other applications. Generally, the low values of overlap show that our social media data is highly diverse. One might expect that each social media application would return similar people for most topics; however, this is clearly not the case. Rather, the low intersection shows that each social media application brings different people to the table and further motivates the exploration of the differences among the applications and of the potential value in their aggregation.

It is especially interesting to inspect the line that shows the overlap of the profile tags with the other applications, since profile tags placed on an individual directly reflect the perception of that individual by others. If we assume profile tags represent the ground truth for a person’s expertise, the overlap of the rest of the applications with them, in spite of being generally low, gives

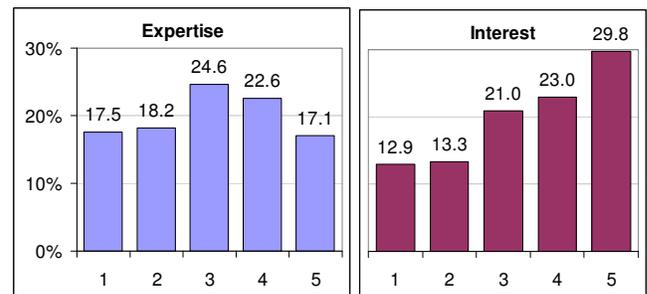


Figure 1. Expertise and interest general rating distributions.

some indication to their relative quality for expertise mining. As can be seen, microblogs have the highest overlap with profile tags, followed by blogs, while forums have the lowest overlap.

4.2 Expertise and Interest Rating

We now move to analyze our survey’s results. First, we inspect the general differences between the ratings of expertise versus interest. Figure 1 shows the rating distribution for each of the two semantics across all 16,750 rated topics. Expertise ratings are distributed quite symmetrically around the center of 3, with a very slight inclination towards positive rating (4, in particular). The average expertise rating was 3.04 (stdev: 1.34). Interest ratings, on the other hand, are significantly higher (two-tailed unpaired t-test, $p < 0.001$) with an average of 3.43 (stdev: 1.37). More than 50% of the interest ratings are positive and nearly 30% are 5. These results indicate that the participants of our survey are generally interested in the topics they were asked about, while their expertise in them is approximately distributed normally.

There was a positive correlation between the expertise rating and the interest rating by a given participant to a given topic (Pearson coefficient of 0.7), indicating that people are usually interested in topics in which they have expertise and vice versa. Yet, there was also a considerable difference between the two ratings: the average absolute value of this difference for a given topic by a given person was 0.72 (stdev: 0.86). 39.3% of the 16,750 rated topics had higher interest rating (for 34.3% of these the difference was greater than 1), while 11.7% had higher expertise rating. Topics for which the expertise rating was higher than the interest rating mainly included projects that were no longer in use or have been sunset and tools employees were forced to use or administrate as part of their everyday work. Overall, these results indicate that users indeed distinguish between expertise and interest and rate them quite differently. Interest is generally rated higher than expertise. We assume that participants felt more comfortable rating themselves as interested rather than as experts in a topic. But it is also likely that users are related to a wider set of topics in which they are interested than topics in which they consider themselves experts.

4.3 Application Comparison

We used the *MAR@k* measure to compare the precision of expertise and interest mining using different social media applications as data sources. Figure 2(a) shows the *MAR@k* comparison of the eight social media applications based on expertise rating. For each application, we retrieved only documents of its type, as detailed in Table 1. Naturally, the values of *MAR@k* decrease as k increases, since the ratings of more individuals that are ranked lower are taken into account. The

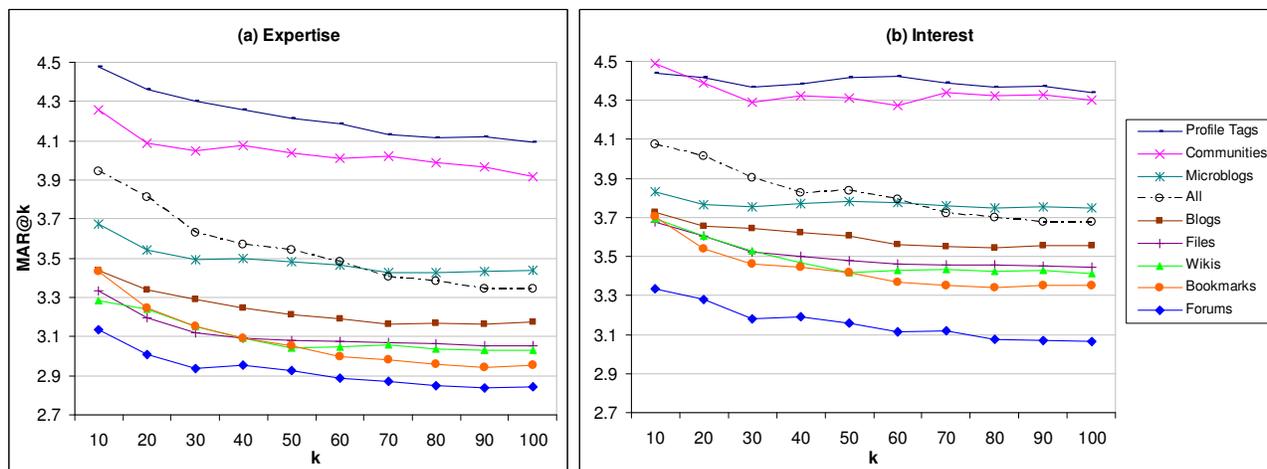


Figure 2. $MAR@k$ as a function of k for expertise and for interest over all 8 applications and their aggregate.

picture is quite clear for the top four applications. Profile tags yield the most precise expertise results with an average rating of almost 4.5 for the top 10 results and 4.1 for the top 100 results. These high values can be associated with the fact that the people tagging application is directly used to annotate other people and map their characteristics. As mentioned, this data also reflects the “wisdom of the crowd”, since tags applied to a person by other employees indicate their collective perception of that person.

Communities are the second most precise source for expertise inference. The high values can be attributed to the high number of user associations with communities (see Table 1), which indicates that each user is a member of many communities. Apparently, a user’s areas of expertise are well propagated through the large number of communities to which they belong.

Microblogging is also a precise data source. This relatively new form of social media is characterized by short but frequent messages. Apparently, the conciseness of content and lack of metadata do not affect the quality of the source. On the other hand, it could be that the high occurrence of messages renders a more comprehensive dataset. This finding also shows that enterprise microblogging is not as noisy as might be expected. While outside the firewall people often update about their everyday activities, microblogging in the workplace is mainly used to discuss work-related topics, promote ideas, or converse about areas of interest [46]. Blogs are the next most accurate source. Compared to microblogs, blogs occur less frequently but are richer in content and metadata such as title and tags. Nevertheless, microblogs appear to outperform blogs when it comes to expertise mining, despite their lack of metadata and shorter text.

For bookmarks, the $MAR@k$ for expertise decreases sharply between $k=10$ and $k=100$. Wikis and files behave similarly to each other, while forums appear as the least precise data source. It could be that forums are used to address ad-hoc problems (e.g., with a programming language or a tool being used) that are not within the user’s key areas of expertise. Generally, the applications that seem to best indicate expertise, such as blogs, microblogs, and to some extent profile tags and communities, can be considered as ‘socializing’ data sources, compared to wikis, files, and forums, used for ‘collaborating’ or getting things done.

This distinction between socializing and collaborating data sources has been previously made in the context of mining social network information [15]. While collaborating sources were found more accurate for extracting network information, it appears that socializing sources are more precise for implying expertise.

Figure 2(b) presents the $MAR@k$ results for interest ratings. The overall values are higher due to the higher average rating of interest compared to expertise, as discussed in Section 4.2. Generally, the decrease with k is milder than for expertise, probably since there are more individuals who are interested in a topic than individuals who are experts, leading to a broader set of relevant results. The order of values between the applications is quite similar to the expertise case, with a few differences. First, communities are very close to profile tags as indicators for interest and even outperform them for $k=10$. This indicates that the user’s set of communities is an even stronger indication for interest than for expertise, while profile tags are relatively stronger for expertise. Microblogs are still third, but with a smaller difference from blogs. Forums are even lower relative to the rest of the applications, suggesting that they are an even weaker indication for interest than for expertise. The overall similarity of the applications’ relative ranking between expertise and interest indicates that sources that accurately relate users to topics do so effectively for both the expertise and interest semantics. That said, some sources, such as profile tags, forums, and microblogs, are slightly better at reflecting expertise, while communities, wikis, and bookmarks tend to better reflect interest.

We also experimented with a version of $MAR@K$ that assigns a neutral value of 3 to each missing feedback. While the range of the results was much closer to 3 in this version, the order of the applications was kept similar to the one shown on Figure 2, with the only difference being that for high values of k , microblogs had higher MAR than communities and bookmarks had higher MAR than wikis, both for expertise and interest. We do not report the full results as they do not provide further insights, however, the similarity in application ranking between the two versions gives another indication for the robustness of the $MAR@K$ measure.

Figure 3 shows the $recall@k$ comparison of the eight social media applications. The values generally increase as k increases, since

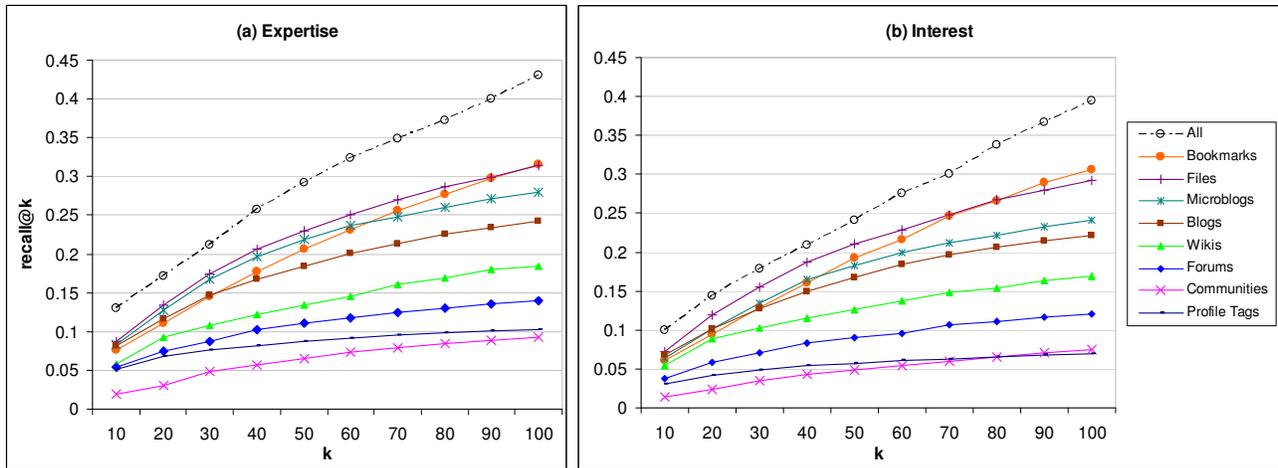


Figure 3. $Recall@k$ as a function of k for expertise and for interest over all 8 applications and their aggregate.

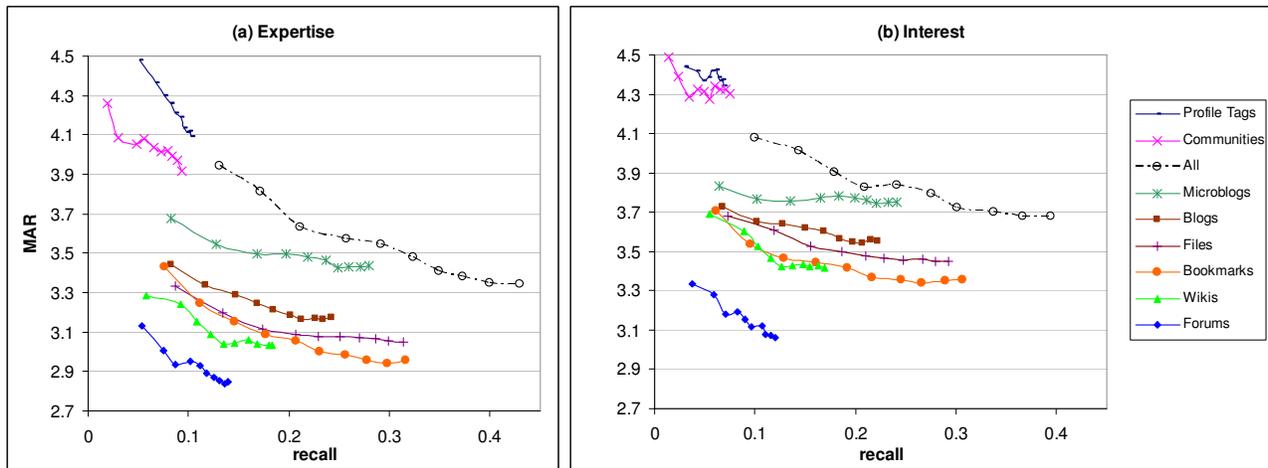


Figure 4. Recall-precision graph for expertise and for interest over all 8 applications and their aggregate.

more experts are covered. Both profile tags and communities, which had the highest precision values as reflected in the $MAR@k$ results, have the lowest values of recall, with communities lower than profile tags for expertise, but higher for interest when k is high. The applications with highest recall are files, bookmarks, microblogs, and blogs.

We note that for both $MAR@k$ and $recall@k$, the order of applications is fairly stable across the different k values (i.e., is not sensitive to k), which strengthens the reliability of both as measures for the quality of a source for expert searching.

Figure 4 presents a view of the overall results through a recall-precision graph that provides a visualization of both the $MAR@k$ and $recall@k$ measures (for $k=10,20,\dots,100$)¹. It can be seen that

¹ Our attempt to produce a combined measure of $MAR@k$ and $recall@k$ using some kind of an F-measure did not succeed, mostly because the two are spanned across different scales and reflect different meanings that cannot be smoothly integrated. We therefore opted to combine the two on a (variant of) two-dimension recall-precision graph.

profile tags and communities are especially strong with regards to precision, with a very low recall, however. Microblogs provide the best combination of recall and precision out of all single applications, followed by blogs. On the other hand, wikis and especially forums yield the weakest combination of precision and recall.

As mentioned before, we also examined an aggregate of all eight applications. To identify an optimal aggregate, we applied coordinate descent [24], a commonly used optimization technique, which iteratively optimizes the objective function by repeatedly cycling through the values of one parameter, while holding the other parameters fixed. In our case, the parameters were the respective weights for each application (eight in total), applied as part of the person-document factor, $W(d,p)$, as explained in Section 3.2, while the objective was set to the product of $MAR@100$ and $recall@100$.

Table 3 shows the “optimal” aggregation weights, reached via the coordinate decent process. The weights are generally consistent with the ranking of the individual applications, as shown in Figure 4, i.e., applications with higher recall-precision values are

Table 3. Optimal aggregation weights for expertise and interest

	blogs	bmarks	cmmnts	files	forums	mblogs	tags	wikis
Exp	0.22	0.11	0.03	0.13	0.11	0.17	0.16	0.07
Int	0.25	0.1	0.12	0.13	0.06	0.2	0.1	0.04

assigned with higher weights. Interestingly, blogs are assigned with the highest weight, followed by microblogs, in spite of the fact that as an individual source microblogs were found superior. Forums, in spite of being the least effective source, are assigned with a slightly higher weight than wikis. Some differences emerge in expertise versus interest optimal weights, most noticeably for communities (higher weight for interest) and for profile tags and forums (higher weight for expertise). The fact that no application was assigned with a weight of zero indicates that each contributes to the overall result improvement.

The $MAR@k$ and $recall@k$ results achieved by the “optimal” aggregate are depicted in Figures 2, 3, and 4 (marked ‘All’). From Figure 2 it can be seen that in terms of precision the aggregate does not exceed profile tags or communities. For higher values of k , the aggregate’s precision is also lower than for microblogs. For recall, on the other hand, the aggregate clearly reaches a substantially higher value than any other single application, as shown in Figure 3. The gap between the aggregate and the single applications is higher for expertise than for interest. Overall, as clearly reflected in Figure 4, the aggregate outperforms all single applications in terms of the combined recall-precision measurement.

4.4 Association Comparison

In the next analysis, we restricted the individuals retrieved to those associated with content by a specific association. Table 4 shows the $MAR@100$ and $recall@100$ results for all association types, by their corresponding applications. The relative results between the associations were similar for other values of k .

The $MAR@100$ results reveal a few subtle differences between expertise and interest ratings. For the blog application, both commenter and liker yield slightly better results than the author association for expertise. This is somewhat unexpected, as one would assume authorship of a blog indicates stronger expertise than feedback in the form of a ‘like’ or a comment. For interest, liking is slightly more indicative than commenting and authorship. For bookmarks, there is a substantial difference in favor of tagger over bookmarker for both expertise and interest. While the tagger association only considers tags, the bookmarker association considers all content and metadata. This finding suggests that the bookmark’s content itself is noisy and that focusing solely on the tags actually yields more accurate results.

The Files application has the largest number of association types—six in total. Liker most precisely indicates expertise, while authorship only comes fourth, preceded also by commenter and sharer. For interest, however, authorship yields the best results, while liker is the least effective. For forums, authorship is more precise than commenting for both expertise and interest. Since many of the authored threads in forums are questions, these results are more intuitive for interest than for expertise. For profile tags, the tags applied to an individual by others more accurately represent the person’s own expertise and interests than the tags used by that person to annotate others. For wikis, commenting is stronger than authorship for both expertise and interest, showing again that feedback on the content can be a more accurate

indication than authorship. For wikis this can be explained by the fact that there may be multiple authors per page.

Overall, we observe that for many of the applications, authorship is not the most precise evidence for expertise. Other forms of feedback, such as liking and commenting, often indicate more accurately the user’s areas of expertise. This is an interesting finding with regards to the use of social media for expertise location, which should be further validated and explored. It could be that social media provides a medium for conversations and discussions in which the authors of files, blogs, and wikis, do not necessarily have more expertise than the users who provide them feedback. In fact, authorship is sometimes a better indicator for interest, as can be observed for files.

Some of the associations are found to be relatively more precise for expertise than for interest, such as file liking, bookmark tagging, and the associations related to profile tags, forums, and microblogs. Others reflect interest more strongly, e.g., file authorship, downloading, and sharedWith, as well as bookmarking. The use of tags in general is a good indicator of expertise, as indicated by the results of both bookmark tagger and profile tagger. These results also imply that tags used to annotate other individuals are a more accurate source for both expertise and interest than tags used to annotate web pages.

$Recall@100$ results are listed in the two rightmost columns of Table 4. These were highly affected by the amount of data extracted for each association type, as detailed in Table 1. Authorship generally has higher recall than other types of associations, since it is more common. Other common associations, such as file downloading or forum commenting also

Table 4. $MAR@100$ and $recall@100$ for expertise and interest over all 20 associations

Association		$MAR@100$		$Recall@100$	
		Exp	Int	Exp	Int
Blog posts	author	3.3	3.68	0.19	0.16
	commenter	3.41	3.71	0.16	0.14
	liker	3.4	3.77	0.14	0.14
Bookmarks	bookmarker	2.92	3.32	0.31	0.3
	tagger	3.46	3.7	0.22	0.2
Communities	member	3.91	4.3	0.09	0.07
Files	author	3.13	3.78	0.24	0.22
	commenter	3.26	3.53	0.13	0.12
	liker	3.44	3.47	0.1	0.09
	sharer	3.2	3.57	0.2	0.19
	sharedWith	3.08	3.61	0.13	0.12
	downloader	3.03	3.48	0.24	0.23
Forum threads	author	3.16	3.32	0.1	0.08
	commenter	2.9	3.09	0.13	0.11
Microblogs	author	3.51	3.83	0.25	0.22
	boardOwner	3.56	3.81	0.24	0.21
Profile tags	tagger	3.94	4.2	0.05	0.06
	taggedWith	4.25	4.45	0.09	0.07
Wiki pages	author	3.07	3.42	0.16	0.15
	commenter	3.37	3.58	0.06	0.05

have relatively higher recall. For bookmarks, bookmarking has higher recall than tagging despite having fewer associations, since each bookmarking association contains considerably richer data.

5. DISCUSSION

Our evaluation provides a comprehensive overview of the potential use of different social media applications for expertise location. We found that the results across the different applications are very diverse and the overlap among the individuals retrieved based on each application is very low. In terms of precision, reflected through the $MAR@k$ measure, profile tags and communities emerged as the two most effective sources for expertise and interest mining. While for the expertise scenario, profile tags were superior, for interest, the two produced similar results. Both profile tags and communities, however, produced the lowest results in terms of recall. Files, bookmarks, and microblogs had the highest recall results. Specifically, microblogs were shown to be effective in terms of both precision and recall, having a good combination of the two. At the other extreme, forums were the least precise source and also had a low recall. Wikis were also among the sources of lower precision and recall.

We observe that the three applications that turn out to be most precise are relatively sparse in content. Profile tags are based solely on tags, communities are based on metadata only, and microblogs are based on short content limited to 500 characters per message. Applications with heavier content, such as blogs, files, wikis, and forums were shown to yield less accurate results. For bookmarks, considering just people associations with tags produced higher precision than taking into account the association with the whole content. On the other hand, using content typically contributes to recall. The pros and cons of using the content should be further investigated. Rich content may produce a wider set of results and allow deeper analysis, but it also seems to increase noise, while social media metadata may often provide precise results on its own.

The aggregate of all eight sources yielded a substantially higher recall than any single application for both expertise and interest. In terms of precision, however, it did not exceed all other single applications: profile tags, communities, and for higher values of k also microblogs, produced higher precision results. Nevertheless, these results indicate that aggregation of multiple sources can help improve the recall, while also obtaining good precision. Aggregation may also improve other desired qualities, such as diversity and serendipity. We used a coordinate descent technique to find the “optimal” weight configuration of the aggregate’s applications. The fact that all eight applications were assigned with a nonzero weight implies that a wide variety of commonly-used applications can help improve an organization’s expertise and interest mining capabilities.

Our survey results show a significant difference between ratings of expertise and interest. While there was a positive correlation between the two, interest ratings were significantly higher, indicating that a relation between a user and a topic may stem from interest in the topic rather than expertise in it. This distinction is especially relevant for social media, where participation is made easy through creating content and providing feedback in many forms. Despite the general rating difference between expertise and interest, the order of applications ranked by their ability to reflect expertise and interest was very similar, indicating that sources that reflect the user’s expertise in a topic

are also likely to reflect the user’s interest in it and vice versa. Some applications were still found to be relatively more precise for expertise inference than for interest (profile tags, microblogs, forums), while others were found relatively more precise for interest inference (communities, wikis, and bookmarks). Additionally, differences emerged between expertise and interest with regards to the most precise association for the files, blogs, and microblogs applications.

The suggested distinction between expertise and interest can have practical implications. Search applications that display people related to a topic (e.g., [3, 37]) may display these two types separately—experts on the topic and people who are interested in it—since each may present a different value for the searcher. The usefulness of locating experts has been widely studied [12, 45], for example to get answers to technical questions or find people with certain skills. In many of these cases, individuals who are interested in the topic but are not experts might not be good enough. On the other hand, experts are a scarce resource and as indicated by our survey results, the number of people who are interested in a topic is likely to be much higher than the number of experts. Thus the list of people interested in a topic may be useful, for example, when establishing a community around a topic, creating a distribution list, or arranging a brainstorm meeting around the topic. The set of people interested in a topic can also help to distribute an idea, a question, a project, or another piece of information that relate to that topic. This set is likely to be larger than the set of experts and these people may be more motivated to disseminate the topic than the actual experts.

When inspecting the different associations, our assumption was that authorship would represent expertise, while feedback such as commenting, liking, and sharing, would indicate interest. However, our results draw a different picture: authorship may not be the strongest indication for expertise when it comes to social media. For many of the applications, feedback-based associations, such as commenting and liking, yielded slightly more precise results for expertise. This result should be further examined and validated. It suggests that within the social media ecosystem, which includes many forms of participation, users who give feedback on a piece of content, may have more expertise than the author(s) who wrote it. Yardi et al. [44] found that employees expect to receive attention when they author a post in a corporate blog. It could be that when it comes to social media, people feel comfortable authoring about topics of interest without being experts, while the more established experts are engaged with different forms of feedback that require less effort. Another possible explanation is that employees in certain roles, such as those in the marketing and communications departments, evangelists, and administrative people, may often publish and share content to raise awareness and generate discussions. Such employees might not be experts on the topic, but may have a high interest in it.

5.1 Directions for Future Work

Previous work has shown that when seeking an expert, users consider both the relevance of the person to the topic and the network topology, for example their social distance to the expert (e.g., [30]). One of the benefits of social media data is that it also reflects social relationships among users [14, 15]. In this work, we focused on the person’s relevance to the topic and did not consider social relationship factors. The social media applications found to be more accurate for expert finding are not necessarily

the ones that contain the most accurate social network information. For example, it has been shown that collaborating sources, such as files and wikis, more accurately reflect a person's strong ties than socializing sources, such as people tags and blogs [15]; yet the latter are found in this work to more accurately reflect expertise and interests. Social media-based expertise location systems can take advantage of different applications to effectively extract both topic relevance and social relationship data.

Several previous studies have tried to identify the documents that better serve as evidence for expertise [26, 36, 40]. This study can be viewed as an attempt to determine the quality of documents as expertise evidence based on the social media application they belong to and the user association to the content. In fact, the public nature of social media data allows to provide the supportive documents as evidence for the expertise of retrieved experts, similar to the way explanations are used in the recommender systems domain [17]. For example, it can be shown that an individual is returned as an expert because she is a member of three relevant communities, authored two related blog posts, and commented on three wikis on the subject. Our future plans include incorporating detailed expertise evidence in our system and further studying of the topic.

The results presented in this work were influenced by the use of social media in our specific organization. We opted to take advantage of the variety of social media applications and the popularity of their use in our organization to provide a first study that broadly examines the use of social media for expertise location. We chose applications that represent the most popular types of social media that exist today. We call for more studies on the topic in other organizations that can further validate and extend the findings of this work. Our future plans include the exploration of social media data sources outside the firewall from which employees' expertise can also be derived. As employees are often active on external blogs, forums, business-related social network sites, and other applications, the combination of enterprise and external data is likely to further enhance expertise location [40].

Our scoring formula for the relation between topics and people is general enough to serve as a common ground for all types of social media applications. Yet, scoring can be further enhanced based on the characteristics of the specific application. For example, for forums, further analyses can distinguish between questions and answers and try to estimate the best answer. For communities, a distinction between members can be made based on their level of activity or their role in the community. It would be interesting to examine to what degree application-specific analysis can enhance expertise inference from social media.

Our survey is based on participant's self-rating, which may give different interpretations to the Likert scale. Moreover, self-esteem ratings may reflect the level participants would want to be associated with a topic rather than the level of expertise as perceived by others. A more objective picture might have been revealed if participants were asked to judge the expertise of others. However, previous attempts to ask participants about others' expertise in a large organization have not been productive, since most respondents simply did not know enough about each other [3]. We therefore opted to count on the breadth of our survey to eliminate noise and bias from individual participants. The high agreement of the profile tags with the survey ratings and

the fact that the other applications' overlap values with profile tags were similarly ordered to their quality as reflected in our survey reinforce the credibility of the self ratings.

Our study presents the potential of expertise mining as social media becomes more popular in organizations and as younger generations, more accustomed to everyday use of social media, join the workforce. Using social media for expertise location can serve as a catalyst for the use of social media in organizations, encouraging more individuals to become more active. On the other hand, dealing with a new form of spam—aimed to improve one's expertise ranking—may become one of the future challenges for social media-based expertise location systems.

6. CONCLUSION

Social media applications are becoming more popular in many organizations and provide new opportunities to mine expertise data. Our evaluation provided an overview comparison of the usefulness of eight different social media applications for mining expertise and interests. Our work lays the foundation for further research on the topic to validate the results in other organizations and extrapolate to other aspects, such as integration with external social media, examination of application-specific expertise mining methods, and combination with network analysis techniques.

7. REFERENCES

- [1] Abel, F., Gao, Q., Houben, G.J., & Tao K. 2011. Analyzing user modeling on twitter for personalized news recommendations. *Proc. UMAP '11*, 1-12.
- [2] Adamic, L.A., Zhang J., Bakshy, E., & Ackerman, M.S. 2008. Knowledge sharing and yahoo answers: everyone knows something. *Proc. WWW '08*, 665-674.
- [3] Amitay, E., Carmel, D., Har'el, N., Soffer, A., Golbandi, N., Ofek-Koifman, S., & Yogev, S. 2009. Social search and discovery using a unified approach. *Proc. HT '09*, 199-208.
- [4] Balog, K., Azzopardi, L., & de Rijke, M. 2006. Formal models for expert finding in enterprise corpora. *Proc. SIGIR '06*, 43-50.
- [5] Balog, K. & de Rijke M. 2006. Finding experts and their details in e-mail corpora. *Proc. WWW '06*, 1035-1036.
- [6] Balog, K., de Rijke, M., & Weerkamp, W. 2008. Bloggers as experts: feed distillation using expert retrieval models. *Proc. SIGIR '08*, 753-754.
- [7] Becerra-Fernandez, I. 2000. Facilitating the online search of experts at NASA using expert seeker people-finder. *Proc. PAKM '00*.
- [8] Buffa, M. 2006. Intranet wikis. *IntraWebs Workshop, WWW '06*.
- [9] Campbell, C.S., Maglio, P.P., Cozzi A., & Dom, B. 2003. Expertise identification using email communications. *Proc. CIKM '03*, 528-531.
- [10] Claypool, M., Le, P., Wased, M., & Brown, D. 2001. Implicit Interest Indicators. *Proc. IUI '01*, 33-40.
- [11] Craswell, N., de Vries, A., & Soboroff, I. Overview of the trec-2005 enterprise track. *Proc. TREC '05*, 199-205.
- [12] Ehrlich K. & Shami, N.S. 2008. Searching for expertise. *Proc. CHI '08*, 1093-1096.

- [13] Farrell, S., & Lau T. 2006. Fringe Contacts: people tagging for the enterprise. *Workshop on Collaborative Web Tagging, WWW '06*.
- [14] Gilbert, E. & Karahalios, K. 2009. Predicting tie strength with social media. *Proc. CHI '09*, 211-220.
- [15] Guy, I., Jacovi, M., Meshulam, N., Ronen, I., & Shahar, E. 2008. Public vs. private: comparing public social network information with email. *Proc. CSCW '08*, 393-402.
- [16] Guy, I., Ur, S., Ronen I., Weber, S., & Oral, T. 2012. Best faces forward: a large-scale study of people search in the enterprise. *Proc. CHI '12*, 1775-1784.
- [17] Herlocker, J. L., Konstan, J.A., & Riedl, J. 2000. Explaining collaborative filtering recommendations. *Proc. CSCW '00*, 241-250.
- [18] IBM Connections – Social Software for Business: <http://www.ibm.com/software/lotus/products/connections>
- [19] Jarvelin, K. & Kakalainen, J. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Trans. Inf. Syst.* 20, 4 (October 2002), 422-446.
- [20] Kardan, A., Garakani, M., & Bahrani, B. 2010. A method to automatically construct a user knowledge model in a forum environment. *Proc. SIGIR '10*, 717-718.
- [21] Kolari, P., Finin, T., Lyons, K., & Yesha, Y. 2008. Expert search using internal corporate blogs. *Workshop on Future Challenges in Expertise Retrieval, SIGIR '08*, 2-5.
- [22] Lerman, K. 2007. Social Networks and Social Information Filtering on Digg. *Proc. ICWSM '07*.
- [23] Li X., Guo, L., & Zhao, E.Y. 2008. Tag-based social interest discovery. *Proc. WWW '08*, 675-684.
- [24] Luo, Z.Q. & Tseng, P. 1992. On the convergence of the coordinate descent method for convex differentiable minimization. *J. Optim. Theory Appl.* 72, 1 (Jan. 1992), 7-35.
- [25] Macdonald, C. & Ounis I. 2006. Voting for candidates: adapting data fusion techniques for an expert search task. *Proc. CIKM '06*, 387-396.
- [26] Macdonald, C., Hannah, D., & Ounis I. 2008. High quality expertise evidence for expert search. *Proc. ECIR '08*, 283-295.
- [27] Manning, C., Raghavan, P., & Shtze, H. 2008. Introduction to information retrieval. *Cambridge University Press*.
- [28] Maron, M.E., Curry, S. & Thompson, P.: An inductive search system: theory, design, and implementation. *IEEE Trans. on Systems, Man and Cybernetics* 16, 1 (1986) 21-28.
- [29] McCandless, M., Hatcher, E., & Gospodneti, O. 2010. Lucene in action, 2nd edition. *Manning Publications Co.*
- [30] McDonald, D.W. & Ackerman, M.S. 2000. Expertise recommender: a flexible recommendation system and architecture. *Proc. CSCW '00*, 231-240.
- [31] Millen, D.R., Feinberg, J., & Kerr, B. 2006. Dogear: social bookmarking in the enterprise. *Proc. CHI '06*, 111-120.
- [32] Muller, M., Ehrlich, K., Matthews, T., Perer, A., Ronen, I., & Guy, I. Diversity among enterprise online communities: collaborating, teaming, and innovating through social media. *Proc. CHI '12*, 2815-2824.
- [33] Noll, M.G., Yeung, A.C., Gibbins, N., Meinel, C., & Shadbolt, N. 2009. Telling experts from spammers: expertise ranking in folksonomies. *Proc. SIGIR '09*, 612-619.
- [34] Pohn W., Pinder, G., Dougherty, C., & White, M. 2001. The lotus knowledge discovery system: tools and experiences. *IBM Systems Journal* 40, 4, (2001), 956-966.
- [35] Qiu, F. & Cho, J. 2006. Automatic identification of user interest for personalized search. *Proc. WWW '06*, 727-736.
- [36] Reichling, T. & Wulf, V. 2009. Expert recommender systems in practice: evaluating semi-automatic profile generation. *Proc. CHI '09*, 59-68.
- [37] Ronen, I., Shahar, E., Ur, S., Uziel, E., Yogev, S., Zwerdling, N., Carmel, D., Guy, I., Har'el, N., & Ofek-Koifman, S. 2009. Social networks and discovery in the enterprise (SaND). *Proc. SIGIR '09*, 836.
- [38] Sen, S., Vig, J., & Riedl, J. 2009. Tagommenders: Connecting users to items through tags. *Proc. WWW '09*, 671-680.
- [39] Serdyukov, P., Rode, H., & Hiemstra, D. 2008. Modeling multi-step relevance propagation for expert finding. *Proc. CIKM '08*, 1133-1142.
- [40] Serdyukov, P. & Hiemstra, D. 2008. Being omnipresent to be almighty: the importance of global web evidence for organizational expert finding. *Workshop on Future Challenges in Expertise Retrieval, SIGIR '08*, 17-24.
- [41] Shami, N.S., Muller M.J., & Millen, D.R. 2011. Browse and discover: social file sharing in the enterprise. *Proc. CSCW '11*, 295-304.
- [42] Vivacqua, A. & Lieberman, H. 2000. Agents to assist in finding help. *Proc. CHI '00*, 65-72.
- [43] Xiao, J., Zhang, Y., Jia, X., & Li, T. 2001. Measuring similarity of interests for clustering web-users. *Proc. ADC '01*, 107-114.
- [44] Yardi, S., Golder, S. A., and Brzozowski, M. J. 2009. Blogging at work and the corporate attention economy. *Proc. CHI '09*, 2071-2080.
- [45] Yiman-Seid, D. & Kobsa, A. 2003. Expert-finding systems for organizations: Problem and domain analysis and the DEMOIR approach. *JOCEC* 13, 1 (2003), 1-24.
- [46] Zhang, J., Qu, Y., Cody, J., & Wu, Y. 2010. A case study of micro-blogging in the enterprise: use, value, and related issues. *Proc. CHI '10*, 123-132.
- [47] Zhang, J., Ackerman, M.S., & Adamic, L. 2007. Expertise networks in online communities: structure and algorithms. *Proc. WWW '07*, 221-230.