

# The SIGIR 2008 Workshop on Future Challenges in Expertise Retrieval (fCHER)

Krisztian Balog  
ISLA, University of Amsterdam  
*K.Balog@uva.nl*

## 1 Introduction

At the TREC Enterprise Track in 2005 the need to study and understand expertise retrieval (ER) was recognized through the introduction of an Expert Finding task (as opposed to mere document retrieval). The task has generated a lot of interest in the IR community, and rapid progress has been made in terms of modeling, algorithms, and evaluation over the past 3 years.

The workshop on Future Challenges in Expertise Retrieval (fCHER) was held in conjunction with the 31<sup>st</sup> ACM SIGIR Conference on Research and Development on Information Retrieval (SIGIR 2008) in Singapore, on July 24, 2008. The main objective of the workshop was to bring people from different research communities together, to discuss recent advances in expertise retrieval, and to define a research agenda for the next years.

The workshop schedule accommodated regular papers (up to 8 pages long) along with position papers (up to 4 pages long). The program committee accepted 8 papers (4 full and 4 position papers). Each paper was reviewed by at least three members of the program committee, consisting of 21 researchers from academia and industry, representing a broad range of disciplines. In addition, the fCHER program also included an invited talk by Arjen P. de Vries. The workshop papers and presentations are available online at <http://ilps.science.uva.nl/fCHER>.

In the following section, we give an overview of the research problems and solutions presented at this workshop.

## 2 Presentations

The day was divided into an invited talk, two technical paper sessions, and a final discussion period. We summarize each below.

### 2.1 Invited Talk

The workshop opened with an invited talk entitled “Expert Finding = Finding People + Assessing Expertise,” delivered by Arjen P. de Vries. Arjen is a senior researcher at CWI and also holds a part-time full professorship at Delft University of Technology (The Netherlands). He has been a co-organizer of the TREC Enterprise track since its launch in 2005.

---

Arjen started his presentation by illustrating the complicating factors in expertise retrieval with a real life example; in many cases people (try to) search for expertise by referral via trusted people, but often do not know what exactly to ask for (what terms to use). Looking at the chain of e-mails in which a request for expertise is passed from one person to another, it is also clear that mere candidate mentions do not necessarily imply expertise. Further issues that complicate system support for expertise retrieval include: (1) by itself, the volume of communication or publication is not a reliable indication of expertise as certain topics engender more opinion than facts, (2) a lack of information about past performance of experts, (3) access to expertise is often controlled, informally or formally, by the experts or their management, (4) the referral network breaks when somebody new enters the organization, and (5) solutions to complex problems require diverse ranges of expertise.

The presentation continued with an overview of existing approaches and techniques used for expertise finding, both in research and in commercial systems. The focus of current systems is on *identifying* and *recommending* experts, while these systems do not really attempt to *validate* (the breadth and depth of) or to *classify* (the type and level of) expertise.

Arjen then went on with identifying two types of challenges, system and evaluation, to be overcome. System challenges include multi-lingual entity extraction, privacy management, and the interoperability with heterogeneous data sources. As to evaluation challenges, the main difficulty is to get realistic data for research purposes. Currently existing data sets (W3C, CERC, and the UvT Expert Collection) are all built from public-facing pages of organizational intranets. Therefore, these collections do not contain (private) communication, databases, click data, security information, and so forth. The conclusion so far is this: although expert finding methods could in principle use many more resources that indicate expertise, possibly more reliably, it is difficult to set up the research. An alternative benchmark model could be to distribute systems, instead of the standard data, queries, and relevance information. Such “in situ” evaluation would require an organization that is willing to evaluate its own intranet with its own users, and would require robust systems, capable of running off-site, unmonitored, from participants.

Arjen concluded his presentation with a discussion of alternative expertise finding related applications, such as locating communities of expertise or supporting formation of cross disciplinary teams. There are interesting further lines of work on assessing expertise and using social network structure and its role in expert finding.

## 2.2 Morning Paper Session

The first session focused on using additional sources of evidence for expertise retrieval. The first two papers look at evidence within an organization, and utilize internal corporate Web 2.0 data. The last two exploit data outside the organization, and acquire additional evidence from the Web.

The paper by Amitay et al. [1], *Finding People and Documents, Using Web 2.0 Data*, introduces a social search application that provides a list of related people to every search. Person-document relationships are mined from Web 2.0 data, such as blogs and social bookmarks, gathered from IBM’s intranet. A large scale user study with over 600 people showed high agreement between self-assessments and system recommendations.

In their position paper entitled *Expert Search using Internal Corporate Blogs*, Kolari et al. [6] discuss the potential usage of internal corporate blogs, as a source of evidence for expert search.

---

Based on over 23,000 blogs from IBM, covering nearly 3 years, they analyze the characteristics of blog growth and content as well as the characteristics of the social network, defined by blog interactions.

In the paper entitled *Expertise Retrieval Using Search Engine Results*, Jiang et al. [4] examine the effectiveness of a document-based expert search model on web search engine results. Among other research questions, the authors address how to build appropriate search queries in order to get relevant information about a person from the Web. Experimental results show that performance using search engine results is comparable to that of obtained from the intranet, however, the two result sets are different.

The final paper of the session, *Being Omnipresent To Be Almighty: The Importance of Global Web Evidence for Organizational Expert Finding* by Serdyukov and Hiemstra [7] also studies the use of web evidence. Various sources of expertise evidence available on the Web are identified and discussed. The authors experiment with obtaining different rankings of candidates by means of APIs of two major web search engines, and demonstrate that the combination of these sources with intranet results perform significantly better than rankings built on organizational data only.

### 2.3 Afternoon Paper Session

The second session featured a diverse set of topics, introducing several new aspects to ER.

First in the session, the position paper by Haruechaiyasak and Kongthon [2], entitled *Multidisciplinary Expertise Retrieval with an Application in R&D Management*, proposes a new and challenging ER task: finding a group of experts whose combined expertise is required to solve a multidisciplinary R&D problem. In addition, the scope of the search is not limited to a single organization. To support this task, a framework is proposed and, as an illustration, a case study is discussed.

Hofmann et al. [3] in their paper, entitled *Integrating Contextual Factors into Topic-centric Retrieval Models for Finding Similar Experts*, study the task of finding similar experts in the following real-world setting. Communication advisors, working at the public relations department of a university, get requests for topical experts from the media. There are cases when the top expert identified is not available and a similar expert has to be recommended. Given this task, a number of contextual factors are identified and their role is assessed through a small-scale user study. Experimental results show that integrating contextual factors with topic-centric expertise retrieval models can significantly improve retrieval performance.

The position paper by Ke and Mostafa [5], *Collaborative Expertise Retrieval: A Referral Approach to Finding Distributed Experts*, discusses an approach to the retrieval of distributed expertise in a networked environment. The task addressed is routing (referring) information needs to experts (information providers) that have the relevant expertise. A conceptual model, an agent simulation framework is proposed, where distributed agents, representatives of information consumers, providers (experts), and referrers learn to collaborate with each other for finding the experts.

Finally, the position paper by Zhu et al. [8], entitled *Evaluating Relation Retrieval for Entities and Experts*, explores entity retrieval, a research area that is closely related to expert finding. Based on the connections between the two tasks, the authors propose tentative guidelines for both entity and expert relation search tasks, in the context of the upcoming INEX Entity Ranking

---

track (INEX-XER 2008), and show how to bring expert and entity retrieval research together for developing approaches that could potentially be effective for both.

## 2.4 Discussion

The final discussion session was centered around the following issues: the notion of expertise, data, new tasks, evaluation, and links to other disciplines.

**The notion of expertise.** It became apparent that it is not possible to establish a clear definition for the concept of expertise, as it is often unclear what it takes for a person to be an expert. It depends on a number of factors, including the task, information need, context, and so forth. The working definition of an expert is a person who is knowledgeable or skilled in some area or field.

**Data.** There was common agreement on the fact that obtaining real organizational data is problematic. Organizations are not likely to share their private, sensitive data. One solution could be to give working systems to an organization, thus perform “in situ” evaluation. Another direction could be to look outside, i.e., on the Web; at the time of writing a proposal for a “web people search” track at TREC is being discussed. Acquiring the data from the Web has its own difficulties, because of limitations on search engine APIs, but it is hoped that the document collection being released as part of the revived web track can be used for a people finding task.

**New tasks.** Research could be informed by talking to people with expertise search (related) problems. Obvious candidates would be consulting, hiring and head-hunting companies, as well as intelligence agencies (both business and defense). There were concrete proposals for new tasks, these concerned (1) finding (cross-disciplinary) teams, i.e., a group of people as opposed to an individual, (2) routing proposals to reviewers, (3) finding the person who knows the expert (not the expert itself), and (4) finding similar experts. Expert finding can be casted as a special case of entity retrieval, therefore, the more general task of finding entities could also be of interest to those working on expertise retrieval.

**Evaluation.** Current test collections all use different types of assessments. These have been provided by an external person (for the W3C collection), by colleagues (for the CERC collection), or by experts themselves (for the UvT collection). The question that cannot be answered yet is the reliability of these judgements, since no collection exists for which more than one type of assessments is available and can be compared.

Participants seemed enthusiastic about the idea to see whether the UvT collection could be expanded, e.g., by adding (more full-text) publications of employees’ or by adding another institute. Most of the discussion converged on three user tasks at universities: who to contact for a PR opportunity (science communicators), team building for proposals (researchers), and distributing funding opportunities (staff and even funding agencies). Another natural task on this type of data is “advisor finding.”

With respect to measures, it was briefly discussed that one could try to measure the performance of expert finding by making explicit by how much the six degrees of separation would be

---

reduced with an expert finding system (somewhat like the motivation behind Cooper's expected search length).

**Links to other disciplines.** While most of the workshop attendees were coming from the IR community, the paper by Ke and Mostafa [5] provided a nice example for links with other fields by using methods from Information Organization, Machine Learning, and Multi-agent Systems. Another example was given by Kolari et al. [6], by analyzing network properties of corporate blogs and blog interactions using methods from Social Network Analysis. During the workshop, links to many other disciplines besides these were discussed.

ER has obvious connections to Knowledge Management, where expertise is often represented in a "skills matrix," with people on one dimension and skills (or knowledge areas) on the other dimension. Given a list of knowledge areas (topics), current expert search methods can be used to fill out this matrix, by using topics as queries. However, it is not clear how to obtain the list of knowledge areas by automatic means from a collection.

Another obvious link is the one to Social Network Analysis, where a task could be, for example, to improve navigation within a social network. Getting real data is a recurring theme, and is a possible obstacle here again. There are also efforts to (web) people search underway within the Semantic Web community.

Much of the research in ER has been carried out by abstracting the user away. Future research offers opportunities for developing user models, which is a sub-area of Human-computer Interaction.

Finally, finding not only people but also properties of people, or more generally, recognizing entities and properties and relations of entities are problems studied in Natural Language Processing that are, increasingly, of interest to people finding tasks in general and to expert finding in particular. An important relevant evaluation effort in this area that is currently being set up is the WePS: Web People Search effort, whose attribute extraction subtask aims to identify CV-type information, including educational, work-related and contact attributes of individuals from web data.<sup>1</sup>

### 3 Conclusions and the Future

The fCHER workshop was a successful gathering to discuss recent advances and to bring new insights on the future research topics in the field of expertise retrieval. The workshop program featured an invited talk and eight high quality papers, and offered a highly interactive environment with lively discussions throughout the whole workshop.

From the workshop and the discussion by participants, a number of themes emerged to both challenge researchers in this area and suggest new avenues of research and development:

- Having three test collections, representing different organizations, already available, it would not be very exciting to run existing expert search algorithms on a public web crawl of a fourth organization. Getting real organizational data, however, is highly problematic, and is not likely to happen.

---

<sup>1</sup>See <http://nlp.uned.es/weps/>.

- 
- The use of web evidence has not been researched extensively yet. Preliminary work shows highly encouraging results, both when using web data for smoothing or re-ranking purposes (using, say, data crawled from an intranet as the primary source of evidence of expertise) and when used as a source of evidence of expertise itself [4, 7]. There is obvious potential in bringing in semantic web data and technology for ER purposes, for instance in the form of friend-of-a-friend networks as made accessible through Google’s Social Graph API.
  - Identification of new tasks and real world use cases for ER technology. The WePS evaluation effort is an interesting example as its task definition is informed by Spock, a people search engine.
  - New evaluation methodologies which evaluate the whole system including aspects relating to system utility (not just accuracy of ranked lists).
  - We need interface and presentation methods geared towards expertise retrieval. For instance, for general web search, query biased summaries (“Google snippets”) have proved to be a valuable aid to help people decide whether or not to click on and read a search result. What is a natural counterpart in the case of expert finding? How do we give sufficient evidence of expertise for a user to at least examine a search result in expert finding?
  - Finally, there is a clear relation between expertise retrieval and entity retrieval as it is being evaluated at INEX or in other forms such as product search. But what *is* the relation? And what sets experts and expertise apart?

The TREC Enterprise Search track will be put on hold in 2009 due to the evaluation challenges discussed before. The track organizers are investigating whether a new TREC track focusing on people search on the web would be viable, assuming that it would attract researchers that have performed expert finding evaluation in the enterprise track setting in recent years.

The final conclusion of the workshop was that ER research has so far provided valuable insights into IR in general, and the community should be optimistic about the future despite the current obstacles.

## 4 Acknowledgments

We extend our sincere thanks to ACM and SIGIR, to the authors and presenters, to our invited speaker, and to the members of the program committee for their contributions to the material and productive discussions that formed an outstanding workshop. The author would like to express special thanks to Arjen P. de Vries and Maarten de Rijke for their valuable comments and input.

## References

- [1] E. Amitay, D. Carmel, N. Golbandi, N. Har’El, S. Ofek-Koifman, and S. Yogev. Finding people and documents, using web 2.0 data. In *Proceedings of the SIGIR 2008 Workshop on Future Challenges in Expertise Retrieval (fCHER)*, pages 1–6, 2008.

- 
- [2] C. Haruechaiyasak and A. Kongthon. Multidisciplinary expertise retrieval with an application in r&d management. In *Proceedings of the SIGIR 2008 Workshop on Future Challenges in Expertise Retrieval (fCHER)*, pages 25–28, 2008.
- [3] K. Hofmann, K. Balog, T. Bogers, and M. de Rijke. Integrating contextual factors into topic-centric retrieval models for finding similar experts. In *Proceedings of the SIGIR 2008 Workshop on Future Challenges in Expertise Retrieval (fCHER)*, pages 29–36, 2008.
- [4] J. Jiang, S. Han, and W. Lu. Expertise retrieval using search engine results. In *Proceedings of the SIGIR 2008 Workshop on Future Challenges in Expertise Retrieval (fCHER)*, pages 11–16, 2008.
- [5] W. Ke and J. Mostafa. Collaborative expertise retrieval: A referral approach to finding distributed experts. In *Proceedings of the SIGIR 2008 Workshop on Future Challenges in Expertise Retrieval (fCHER)*, pages 37–40, 2008.
- [6] P. Kolari, T. Finin, K. Lyons, and Y. Yesha. Expert search using internal corporate blogs. In *Proceedings of the SIGIR 2008 Workshop on Future Challenges in Expertise Retrieval (fCHER)*, pages 7–10, 2008.
- [7] P. Serdyukov and D. Hiemstra. Being omnipresent to be almighty: The importance of global web evidence for organizational expert finding. In *Proceedings of the SIGIR 2008 Workshop on Future Challenges in Expertise Retrieval (fCHER)*, pages 17–24, 2008.
- [8] J. Zhu, A. P. de Vries, G. Demartini, and T. Iofciu. Evaluating relation retrieval for entities and experts. In *Proceedings of the SIGIR 2008 Workshop on Future Challenges in Expertise Retrieval (fCHER)*, pages 41–44, 2008.