Evaluating Relation Retrieval for Entities and Experts

Jianhan Zhu
University College London & The Open University
Adastral Park Campus
Ipswich, IP5 3RE, UK
j.zhu@adastral.ucl.ac.uk

Arjen P. de Vries
CWI & Delft University of Technology
Kruislaan 413, 1098SJ
Amsterdam, The Netherlands
arjen@acm.org

Gianluca Demartini, Tereza Iofciu
Leibniz Universität Hannover
L3S Research Center
Germany
{demartini, iofciu}@L3S.de

ABSTRACT
Tremendous progress has been made in terms of retrieval models and user evaluation for expert finding. From 2007, INEX provides the XML Entity Ranking track (INEX-XER) as a forum for the study of entity retrieval, a research area closely related to expert finding. Here, instead of being restricted to finding people (a particular type of entity) with a specified expertise (the topic), any type of entity related to a given topic can be the target of the retrieval system. INEX-XER 2008 proposes a novel entity relation search task, which goes beyond entity retrieval by further establishing relations between entities. Based on the connections between expert and entity retrieval, we propose to explore a tentative expert relation search task in this position paper. Our proposal shows how we can bring expert and entity retrieval research together for developing approaches that could potentially be effective for both. We expect this proposal to inspire contributions to expert finding from other research areas than information retrieval, such as semantic web, information extraction, social network analysis, virtual communities, and question answering.

Categories and Subject Descriptors
H.3 [Information Storage and Retrieval]: H3.1 Content analysis and Indexing; H.3.3 Information Search and Retrieval

General Terms
Experimentation, Measurement, Performance

Keywords
Relation retrieval, expert finding, entity retrieval

1. INTRODUCTION AND MOTIVATION
Many user tasks would be simplified if search engines would support typed search, and return entities instead of 'just' web pages. As an example, expert finding, i.e., retrieving people (the entity) having a specified expertise (the topic), is a key task in enterprise search and has recently attracted lots of attention from both academic and industrial communities, as evidenced by the organization of the Expert Search Task in TREC [1, 6, 13]. Since 2005, tremendous progress has been made in terms of expertise modeling, algorithms, and evaluation strategies. The goal of expert finding is to identify a list of people who are knowledgeable about a given topic. Contrary to traditional IR systems, the target of expert finding is retrieving people (named entities) instead of documents. This task is usually addressed by uncovering associations between people and topics [6].

Future Challenges in Expertise Retrieval
SIGIR 2008 workshop, July 24, Singapore

Balog et al. [2] proposed the use of language modeling for expert finding, introducing a "Model 1" which directly represents the knowledge of an expert from associated documents and a "Model 2" which first locates documents on the topic and then finds the associated experts. Other expert finding approaches include the two-stage language model by Cao et al. [4], a generative probabilistic model by Fang and Zhai [8], a proximity-based document representation model by Petkova and Croft [10], data fusion models by Macdonald and Ounis [9], and an expert-centric language model by Serdyukov and Hiemstra [12] etc.

We observe a growing interest in extending the typed search introduced with expert finding to the retrieval of entities of other types. For example, [5] proposed the EntityRank algorithm that integrates local co-occurrence and global access information for entity search into a probabilistic estimation of entity and query association, which is quite similar to a two-stage expert finding approach. Also, INEX (INitiative for the Evaluation of XML retrieval) 2007 has started the XML Entity Ranking track (INEX-XER) to provide a forum where researchers may compare and evaluate techniques for engines that return lists of entities [7].

Expert finding and entity retrieval are closely related research areas. Since important progress has been made in expert finding since 2005, these expert finding models and techniques could also be applied to entity retrieval; as evidenced by the same random walk approach that was applied to expert finding [11] and entity retrieval [14].

In the upcoming INEX-XER 2008, we propose a new entity relation search (ERS) task investigating how well systems can not only find entities relevant to a topic but also establish correct relations between entities.

The motivation of the ERS task is that user information needs are often not satisfied with 'just' a list of entities relevant to a query, because the user would like to know more details about these entities, such as their relations with other entities, and their attributes. In a similar direction, Balog et al. [3] proposed expert profiling to complement expert finding in enterprise environment. They define the profile of an expert as her "topical profile" consisting of her skills and areas of expertise, and "social profile" in the form of her collaboration network.

We think that entity relation search could be applied to expert finding in terms of finding relations between experts and entities. We think that the INEX XML Entity Ranking track and TREC Expert Search task can complement each other in terms of task designs, retrieval models, and result evaluation etc. In this paper, we introduce the INEX XML Entity Ranking Track, explore its connections to expert finding, and propose an expert relation search task that can be carried out practically. In Section 2, we introduce the Wikipedia dataset for entity ranking. We give an overview of INEX-XER in Section 3. In Section 4, we propose the new Entity Relation Search (ERS) task. In Section 5, we propose expert relation search on the basis of entity relation search.

2. WIKIPEDIA DATASET FOR ENTITY RANKING

The Entity Ranking track uses the INEX Wikipedia XML collection, exploiting the category metadata about the pages to loosely define the entity sets. Given preferred categories, relevant entities are assumed to loosely correspond to those Wikipedia pages that are labeled with these preferred categories (or perhaps sub-categories of these preferred categories). Retrieval methods need to handle the situation where the category assignments to Wikipedia pages are not always consistent, and also far from complete. For example, given a preferred category 'art museums and galleries' (10855), an article about a particular museum such as the 'Van Gogh Museum' (155508) may not be labeled by 'art museums and galleries' (10855) but labeled by a sub-category of the preferred category instead, such as category 'art museums and galleries in the Netherlands' (36697). Therefore, when searching for "art museums in Amsterdam", correct answers may belong to other categories close to this category in the Wikipedia category graph, or may not have been categorized at all by the Wikipedia contributors.

3. ENTITY RANKING

Entity ranking concerns tuples of type <query, category>. The category (the entity type) specifies the type of 'objects' to be retrieved. The query (consisting of title, description, and narrative fields) attempts to capture the information need. Examples of entity ranking topics include "find European countries where I can pay with Euro", "find cities in the world where a summer Olympic game has been organized" etc. Here we can see that the set of entities to be ranked is assumed to be loosely defined by a generic category, which is often implied in the query itself, e.g., entities of type "European countries" and "cities" are desired in the above two examples, respectively. Another example of an INEX-XER topic is given in XML format.

```
<title>Impressionist art in the Netherlands</title>
<description>I want a list of art galleries and museums in the Netherlands that have impressionist art.</description>
<narrative>Each answer should be the article about a specific art gallery or museum that contain impressionist or post-impressionist art works.</narrative>
<categories>
    <category id="10855">art museums and galleries</category>
</categories>
```

We can treat expert finding as a special case of entity retrieval where we use the semantic notion of 'people' as its core category, and the query would specify 'expertise on T' for expert finding topic T. Of course, not all entity ranking queries with target category 'people' are expert finding topics; the 2007 test collection included also topics searching for presidents, tennis players and composers.

One important difference between the TREC Expert Search task and the INEX-XER bound to experts is the context and the dataset. The former focuses on the enterprise settings where the goal is to extract evidence from a dataset of e-mails and web pages, while the latter uses an encyclopedia as description of people's expertise and the queries can spread much more over all the possible topics.

4. ENTITY RELATION SEARCH

In some cases a search engine user might want to find relations between entities. In this Section we propose a new search task built on top of entity ranking. In entity relation search, we try to model a more exploratory search scenario, where people are interested in exploring the different aspects of entity ranking results. This corresponds to a view on entity relation search where the tasks are divided into an entity ranking stage, followed by the relation search stage. Given the entity ranking results, the motivation of entity relation search is here to retrieve further details about relevant entities found in entity ranking.

We call the entities found in entity ranking the main entities. Further details about the main entities are retrieved in the form of relations between each of these main entities and its related entities, which we call the target entities. The relations between main entities and target entities can be either 1 to 1, i.e., one main entity is related to one target entity, or 1 to n (n>1), i.e., one main entity is related to several target entities. These relations can be also seen as specifying (possibly multi-valued) attributes of the main entities.

Entity relation search concerns tuples of type <query, category, relation-query, target-category>. The query and category are defined in the same way as in the entity ranking task. The relation-query, given as free text, describes the desired relation between main and target entities. The relation query consists of a relation title, relation description, and relation narrative fields. The target-category specifies which category (entity type) is desired for the target entity.

The results of an entity relation search topic consist of pairs of main and target entities. For each pair of entities to be judged as a correct pair, the main entity must be judged as relevant to the original query, the main entity has to be of its correct category, the target entity is of its correct category, and the relation between them matches the relation topic.

For example, given 'Art museums and galleries' as the category and 'Impressionist art in the Netherlands' as the query topic for the main entities, 'cities' as the category for the target entities and relation query topic 'located in', we expect answer pairs like 'Van Gogh museum' and 'Amsterdam', representing the fact that the 'Van Gogh museum' is located in 'Amsterdam'.

Like in the entity ranking task, the entity types for both the main and target entities are only loosely defined by their categories 'art museums and galleries' and 'cities', respectively. Correct answers may belong to other categories close to these two categories in the Wikipedia category graph, respectively, or may not have been categorized at all by the Wikipedia contributors.

We formulate the example topic as below:

<title>Impressionist art in the Netherlands</title>
<description>I want a list of art galleries and museums in the Netherlands that have impressionist art.</description>
<narrative>Each answer should be the article about a specific art gallery or museum that contain impressionist or post-impressionist art works.</narrative>
<categories>
<category id="10855">art museums and galleries </category>
</categories>
<entity-relation>
<relation-title>located in</relation-title >
<relation-description>I want the cities where these art galleries and museums are located. </relation-description>
<relation-narrative>Each answer should be a city where a specific art gallery or museum that contain impressionist or post-impressionist art works is located. </relation-narrative>
<target-categories>
<category id="2917">cities</category>
</target-categories>
</entity-relation>

In evaluating entity relation search results, Wikipedia pages for both main and target entities are returned. The evaluator may need to read both pages in order to find evidence for judging whether their relations match the relation topic. Therefore, entity relation judgment is more complex than entity ranking judgment. After an initial pilot experiment developing some topics with assessments, we believe however that modeling relation search as an exploratory search scenario (that extends an initial ranking of the main entities to explore their attributes) alleviates this complexity sufficiently.

For evaluating the effectiveness of systems performing this task, we need to check whether they correctly identified the main entity (as in Entity Ranking), the relation, and the target entity. It is possible to extract out of the human judgments the correct (i.e., relevant) triples of the form (main entity, relation ,target entity). Similarity, it is possible to extract out of the system results the proposed (i.e., retrieved) triples. In this way we can compare relevant and retrieved results and traditional evaluation measures (such as MAP and average R-precision) can be used to measure performance of systems on entity relation search.

5. RELATION RETRIEVAL FOR EXPERTS

Our proposed entity relation search task can be applied to expert finding as well, since we may be interested in exploring further details of an expert on a search topic. We can similarly divide expert relation search into an expert search stage followed by a relation search stage. Further details about an expert are in the form of relations between the expert and other entities.

In that case, expert relation search would concern tuples of type <query, relation-query, target-category>. The query describes an expertise request, such as find experts on “semantic web”. The relation-query in form of free text describes the relation between an expert and an entity, and consists of a relation title, relation description, and relation narrative fields. The target-category specifies which category (entity type) is desired for the entity.

The results of an expert relation search topic consist of pairs of experts and entities, e.g., the relations between experts and entities of type “projects”, “organizations”, and “academic departments” can be defined as “projects-works-on”, “clients-consulted”, and “department-works-for”, respectively. Similar to entity relation search, the relations between experts and entities can be either 1 to 1, i.e., one expert is related to one entity, or 1 to many, i.e., one expert is related to several entities. For each pair of entities to be judged as a correct pair, the expert must be judged as relevant to the query, the target entity be of the correct type, and the relation between them matching the relation topic; e.g., to find “software engineering” experts and the projects they work on, for each pair consisting of person X and entity Y, we need to judge in three steps: Firstly, is person X an expert on “software engineering”? Secondly, is Y a project name? Finally, does X work on project Y?

For example, given a query topic ‘semantic web’, ‘scientific journals’ as the category for the target entities, and a relation query topic ‘published in’, the correct answers to this relation search topic will be pairs of experts and journals where each pair consists of an expert on “semantic web”, and a journal where the expert has published at least one paper.

We can formulate the example topic as below:

<title>semantic web</title>
<description>I want a list of people who are knowledge in semantic web research in my organization.</description>
<narrative>Each answer should be a person who is an expert on semantic web in my organization.</narrative>
<entity-relation>
<relation-title>published in</relation-title >
<relation-description>I want to find the journals where experts on semantic web in my organization have published papers. </relation-description>
<relation-narrative>Each answer should be a journal where an expert in my organization has published a paper.
</relation-narrative>
<target-categories>
<category id="112">scientific journals</category>
</target-categories>
</entity-relation>

In TREC2005 and 2006 Expert Search task, a crawl of the W3C website was used for expert finding [6, 13]. A predefined list of W3C related people consisting of their names and email addresses was given. Participants employed named entity recognition techniques to annotate the dataset for occurrences of these candidates. The domain for TREC2007 expert finding is the CSIRO website [1], and there was not a predefined list of candidates. Therefore, participants need to employ effective named entity recognition techniques for annotation of CSIRO related people.

In the expert relation search scenario, we envisage that either a predefined list of entities of different types would have to be provided for annotating a dataset similar to the TREC2005 and 2006 Expert Search tasks, or named entity recognition techniques should be employed to recognize entities of different categories from text like in the TREC2007 expert search task.

Of course, the cooperation of the enterprise for which to develop the collection is required. A possibly attractive alternative would be to carry out expert finding and expert relation search tasks on the Wikipedia dataset, where a list of people and category information are readily available; an example would be to find experts on "big bang theory" in the Wikipedia dataset and find the country where each of these experts was born. Like in the entity relation search task, the category for the entities is only loosely defined, and correct answers may belong to other categories close to this category.

In evaluating expert relation search results, there are two kinds of approaches we may choose. First, ask domain experts to judge like in the TREC2007 expert search task. However, this may depend on the nature of relations and type of entities, e.g., country-of-origin of the expert does not really relate to the domain experts' domain knowledge. Second, return supporting documents for expert relations. Human evaluators judge the relations based on evidence contained in these supporting documents like in the TREC2006 expert search task [13].

Traditional evaluation measures, e.g., MAP and R-precision etc. can be used to measure performance expert relation search.

6. CONCLUSIONS

Substantial advances in terms of retrieval models and user evaluations etc. have been made in expert finding research. On the other hand, organization of the Entity Ranking track in INEX opens the door to the study of effective approaches for retrieval of entities of different types. In this paper, we explore how we can let research in expert finding and entity ranking complement each other, in particular, via our proposed relation search task for both entities and experts. We propose tentative guidelines for both entity and expert relation search tasks. We think that organization of the proposed task will help advance the research in both entity and expert retrieval by providing a platform for comparing and experimenting effective approaches for both entity and expert retrieval. A number of groups will participate in our entity relation search task in 2008. Their results on the task will provide insight into entity relation retrieval.

In the first step we design the relation retrieval task for both entities and experts as a two-stage process due to the following two reasons. Firstly, since relation retrieval task is based on entity/expert retrieval task, topic creation and user evaluation can be integrated for the two tasks. Thus topic creation and assessment can be greatly simplified. Secondly, the relationships between the two tasks can be more easily studied.

Our proposed two-stage relation retrieval task opens the door to exploring other types of relation retrieval task. One way is to focus on the relationships between main entities, e.g., finding all pairs of impressionist artists who have influenced each other or all experts in the organization who have worked together on a project etc. The challenges in this type of task can be how to formally define the relations between entities, how to evaluate the relations, and how to define the scope of such relations, e.g., how to define the "influence" relation, how to evaluate the relation, and how to know how many people "influence" each other etc.

7. ACKNOWLEDGEMENTS

We would like to thank the anonymous reviewers' comments, which have helped improve our paper.

8. REFERENCES

[1] Bailey, P., Craswell, N., de Vries, A. P., Soboroff, I. (2007) Overview of the TREC 2007 Enterprise Track (DRAFT). In Proc. of TREC 2007 Notebook.

[2] Balog, K., Azzopardi, L. and de Rijke, M. (2006) Formal models for expert finding in enterprise corpora. In Proc. of SIGIR 2006: 43-50.

[3] Balog, K., & de Rijke, M. (2007) Determining Expert Profiles (With an Application to Expert Finding). In Proc. of IJCAI 2007: 2657-2662

[4] Cao, Y., Liu, J., Bao, S. and Li, H. (2006) Research on Expert Search at Enterprise Track of TREC 2005. In Proc. of TREC 2005.

[5] Cheng, T., Yan, X., and Chang, K. C-C (2007) EntityRank: Searching Entities Directly and Holistically. In Proc. of VLDB 2007: 387-398

[6] Craswell, N., de Vries, A.P., Soboroff, I. (2006) Overview of the TREC-2005 Enterprise Track. In Proc. of TREC 2005.

[7] de Vries, A.P., Thom, J.A., Vercoustre, A-M, Craswell N., and Lalmas, M. (2007) INEX 2007 Entity Ranking Track Guidelines - V1. In Pre-Proc. of INitiative for the Evaluation of XML Retrieval.

[8] Fang, H., and Zhai, C. (2007) Probabilistic Models for Expert Finding. In Proc. of ECIR 2007: 418-430

[9] Macdonald, C. & Ounis, I. (2006) Voting for candidates: adapting data fusion techniques for an expert search task. In Proc. of CIKM 2006: 387-396

[10] Petkova, D., and Croft, W. B. (2007) Proximity-based document representation for named entity retrieval. In Proc. of CIKM: 731-740.

[11] Serdyukov, P., Rode, H., and Hiemstra, D. (2007) University of Twente at the TREC 2007 Enterprise Track: Modeling relevance propagation for the expert search task. In Proc. of TREC 2007.

[12] Serdyukov, P., and Hiemstra, D. (2008) Modeling documents as mixtures of persons for expert finding. In Proc. of ECIR 2008.

[13] Soboroff, I., de Vries, A.P. and Craswell, N. (2007) Overview of the TREC 2006 Enterprise Track. In TREC 2006.

[14] Tsikrika, T., Serdyukov, P., Rode, H., Westerveld, T., Aly, R., Hiemstra, D., and de Vries, A. (2007) Structured Document Retrieval, Multimedia Retrieval, and Entity Ranking Using PF/Tijah. In Proc. of the 6th Initiative on the Evaluation of XML Retrieval (INEX-2007), Dagstuhl, Germany.