Dynamic Collective Entity Representations for Entity Ranking (Abstract)

David Graus
University of Amsterdam
d.p.graus@uva.nl

Manos Tsagkias
904Labs
manos@904labs.com

Wouter Weerkamp
904Labs
wouter@904labs.com

Edgar Meij
Yahoo Labs
emeij@yahoo-inc.com

Maarten de Rijke
University of Amsterdam
derijke@uva.nl

1. INTRODUCTION AND METHOD

We summarize the findings of Graus et al. [2]. Many queries issued to search engines are related to entities [4]. Entity ranking, where the goal is to position a relevant entity at the top of the ranking for a given query, is therefore becoming an ever more important task [1]. Entity ranking is inherently difficult due to the potential mismatch between the entity's description in a knowledge base and the way people refer to the same entity when searching for it.

We propose a method that aims to close this gap by leveraging collective intelligence as offered by external entity "description sources". We differentiate between dynamic description sources that are timestamped, and static ones that are not. We leverage five static description sources for expanding entity representations. First, from the knowledge base: (1) anchor text of inter-knowledge base hyperlinks, (2) redirects, (3) category titles, and (4) names of entities that are linked to or from an entity. From the web: (5) web anchors that link to entities. In addition, we leverage three dynamic description sources, whose content are added to entity representations in a streaming manner: (6) search engine queries that yield clicks on entities, (7) tweets, and (8) tags that mention entities.

We represent entities as fielded documents [5], where each field corresponds to content that comes from one description source. As external description sources continually come in, the content in the entity's fields changes, and previously learned feature weights may be sub-optimal. Hence, constructing a dynamic entity representation for optimal retrieval effectiveness boils down to dynamically learning to optimally weight the entity's fields. We exploit implicit user feedback (i.e., clicks) to retrain our model, and relearn the weights associated to the fields, much like online learning to rank [3]. As an entity ranker, we employ a random forest classifier, using its confidence scores as a ranking signal.

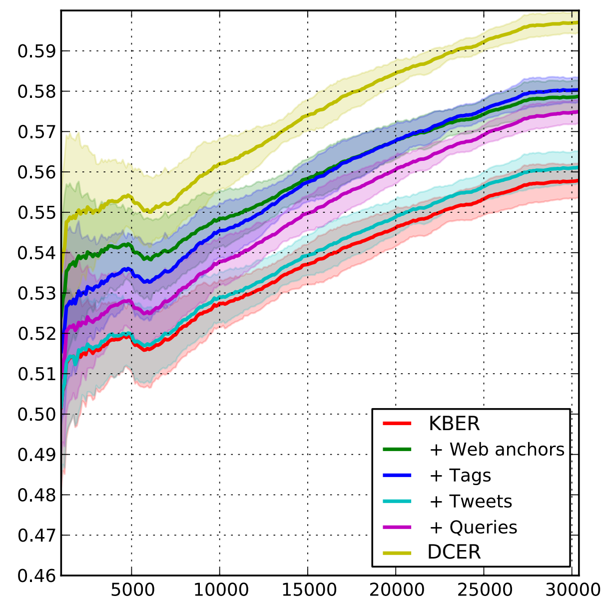

Figure 1: Impact on performance of individual description sources. MAP on the y-axis, number of queries on the x-axis. This plot is best viewed in color.

2. RESULTS AND CONCLUSIONS

To evaluate our method's performance over time, we treat users' clicks as ground truth, and the goal is to rank clicked entities at position 1. We split the query log into chunks, allocate the first chunk for training, and evaluate each succeeding query in the next chunk. Then, we add this chunk to the training set, retrain the classifier, and continue evaluating the next chunk. In Figure 1 we compare our Dynamic Collective Entity Representation method (`DCER`) to a static baseline, which only exploits the Knowledge Base description sources (`KBER`), i.e., sources 1–4 in Section 1. We see how each individual description source contributes to more effective ranking, with `KB+tags` narrowly outperforming `KB+web` as the best single source. We observe that after about 18,000 queries, `KB+tags` overtakes the (static) `KB+web` method, suggesting that newly incoming tags yield higher ranking effectiveness.

Our results demonstrate that incorporating dynamic description sources into dynamic collective entity representations enables a better matching of users' queries to entities. Furthermore, we show how continuously updating the ranker leads to improved ranking effectiveness in dynamic collective entity representations.

REFERENCES

[1] K. Balog, P. Serdyukov, and A. de Vries. Overview of the TREC 2010 entity track. In *TREC 2010*, 2011.

[2] D. Graus, M. Tsagkias, W. Weerkamp, E. Meij, and M. de Rijke. Dynamic collective entity representations for entity ranking. In *WSDM '16*, 2016.

[3] K. Hofmann, S. Whiteson, and M. de Rijke. A probabilistic method for inferring preferences from clicks. In *CIKM '11*, pages 249–258, 2011.

[4] R. Kumar and A. Tomkins. A characterization of online browsing behavior. In *WWW '10*, pages 561–570, 2010.

[5] C. Macdonald, R. L. Santos, I. Ounis, and B. He. About learning models with multiple query-dependent features. *ACM TOIS*, 31(3):11:1–11:39, 2013.