

Snippet-based Relevance Predictions for Federated Web Search

Thomas Demeester¹, Dong Nguyen²,
Dolf Trieschnigg², Chris Develder¹, and Djoerd Hiemstra²

¹ Ghent University, Ghent, Belgium

{tdmeeste, cdvelder}@intec.ugent.be

² University of Twente, Enschede, The Netherlands

{d.nguyen, d.trieschnigg, d.hiemstra}@utwente.nl

Abstract. How well can the relevance of a page be predicted, purely based on snippets? This would be highly useful in a Federated Web Search setting where caching large amounts of result snippets is more feasible than caching entire pages. The experiments reported in this paper make use of result snippets and pages from a diverse set of actual Web search engines. A linear classifier is trained to predict the snippet-based user estimate of page relevance, but also, to predict the actual page relevance, again based on snippets alone. The presented results confirm the validity of the proposed approach and provide promising insights into future result merging strategies for a Federated Web Search setting.

Keywords: Federated Web search, snippets, classification, relevance judgments.

1 Introduction

The actual influence of result snippets on the overall efficiency of search engines has been largely understudied until now. The reason is, that until recently no dedicated test collection containing actual result snippets as well as the resulting web pages from a wide variety of sources, was available to researchers. The collection used for the current paper attempts to fill this gap. In fact, it has been created to stimulate research in Federated Web Search (FWS), and contains a large amount of sampled data from over a hundred diverse online search engines, and relevance judgments of snippets and pages for the results from these resources in response to the 50 topics from the 2010 TREC Web Track [3]. A description of the collection can be found in [5], and an analysis of the relevance judgments in [4]. The full collection and further details are available on <http://snipdex.org/datasets/>.

The goal of this paper is to analyze how well result snippets from a specific origin can be used to predict the relevance of the corresponding page, given a query, and how well these can be cast into a single merged list. This research objective is motivated in the first place by the question whether in FWS, merging the results from various resources into a single ranked list, could be done based

on the snippets alone. Another issue is whether caching large amounts of (small) result snippets, as opposed to a limited number of full pages for the same storage capacity, would improve the accuracy of fast FWS systems, e.g., organized in a peer-to-peer setting.

One of the main problems in Federated Information Retrieval (FIR) and related to the goal of this paper, is result merging, or ordering the results retrieved from several resources into a single ranked list [7]. A related recent trend in FWS is Aggregated Search, the study of how best to integrate results from multiple specialized search services (or verticals) into the Web search results [2]. Classification-based methods have already been used in FIR research, but mostly focusing on predicting the relevant resources, for example in [1] for verticals. The snippets in our collection allow applying classification techniques for relevance predictions on the result level. In this paper we focus on predicting binary relevance and use a maximum entropy classifier [6].

We will describe a linear classifier designed to make snippet-based relevance predictions, shortly covering the extracted snippet features, and reporting classification efficiency for snippet and page predictions. We will also report the precision of the top-30 results, ranked according to the classifier output, and show that even with simple features the results clearly outperform a round-robin merging of results from three of the strongest search engines.

2 Inferring Relevance from Snippets

The FWS collection presented in [5] contains several levels of relevance judgments, but here we will only consider strict binary relevance levels, for which snippets were judged to be definitely relevant³, and pages were indicated as highly relevant or better. In [4] it was shown that even for the best resources only about 2/3 of all ‘definitely relevant’ snippets corresponded with a highly relevant result page. A classifier trained on snippet labels would hence not be able to overcome this intrinsic gap in relevance when used for the prediction of the page labels. Instead, we need to directly focus on page relevance, but still based on the information contained in the snippets alone, as motivated above. We will show that a binary linear classifier, trained on simple snippet features, is able to predict page relevance with a comparable accuracy as for snippet predictions, despite the apparent mismatch in relevance.

In [4] it was shown that a number of the used resources contain very few highly relevant pages, due to the nature of the query set, which has been designed for the 2010 TREC Web Track [3]. Therefore, we will only consider the resources from the following verticals: General Web Search, Multimedia, Encyclopedia/Dictionaries, Blogs, Books Libraries, News, and Shopping. In total, we retain 50 out of the original 108 resources, which together provided almost 90% of all highly relevant result pages from the collection. No further resource selection is performed.

³ With *snippet relevance*, we actually mean the snippet-based estimate of the page relevance (or, in annotation terms, ‘would you click this snippet?’).

For our classification experiments using the maximum entropy classifier, we applied 10-fold cross validation, for each run training on the results from 45 out of the total of 50 topics, testing on the remaining 5. Our presented results were averaged over each of the test topics.

2.1 Snippet Feature Extraction

The features used for classification, extracted from each of the snippets, can be divided into three groups. The first group, indicated in the following as ‘constant’, contains an overall constant feature and a binary feature for each of the considered resources. Considering only these features corresponds to assigning a prior probability of relevance to each resource, according to the number of highly relevant results found in the training set. These vary widely, with the largest general web search engines on top. The second group contains a number of features that can be directly calculated from each resource’s snippets, without the need for a large amount of snippets from other resources. This distinction is important, e.g., for result merging purposes, where snippets from only a few resources would be available. These features include binary features for the considered snippet’s rank in the result list (we only crawled the top 10 results for each resource), features indicating the presence of query terms in the title, summary and location fields of the query (see [5]), and the length of these fields. In our reported results, the combination of these features together with the constant features is indicated as ‘local’. Finally, a number of global features are extracted, calculated from the whole set of snippets for the considered topic. These include further length-based features for the different snippet fields, but relative to the average length in the considered set, and tf-idf scores for the snippet title and summary in response to the query (without stemming). All features together are indicated with ‘all’.

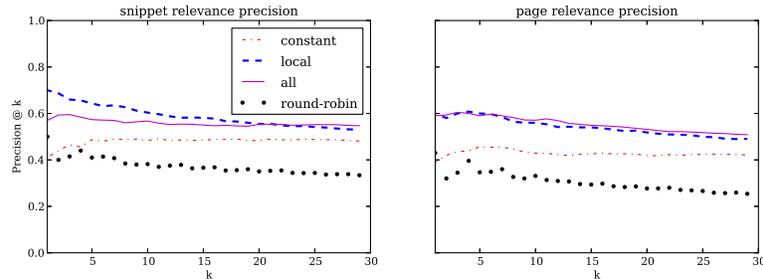
2.2 Results

In web search, precision is typically considered more relevant than recall, we thus set the relevance cut-off for our classifier on a predicted probability of relevance of 0.5, with relatively high precision, and sufficient recall. The mean classifier precision and recall, both for the prediction of snippet and page relevance, are shown in Table 1. Using all features, the prediction of page relevance performs only slightly lower than for snippets. Note that no resource had on average more than 5 out of 10 highly relevant pages (as seen from the zero scores using only the constant features), whereas the snippet judgments were more optimistic. Also, calculating the global features did not yield much improvement in terms of precision, but allowed retrieving more relevant results.

Ranking the results according to the predicted probability of relevance, we can create a merged result list. As a baseline, we created a round-robin ranking of the three major search engines, i.e., Google, Yahoo!, and Bing. The precision in the top-30 of this ranked list, together with the baseline, are shown in Fig. 1. The precision for the predicted results is considerably higher than for the round-robin baseline, even for the constant features (i.e., the mere ranking of resources

Table 1: Snippet-based prediction accuracy for snippet and page relevance.

features	snippet relevance		page relevance	
	precision	recall	precision	recall
constant	0.52	0.04	0.00	0.00
local	0.65	0.27	0.61	0.15
all	0.67	0.37	0.62	0.23

Fig. 1: Precision as a function of the rank k for the 30 highest ranked results as predicted for all resources, compared with the round-robin result for Google, Yahoo!, and Bing.

that best explains the training labels), and the precision in page predictions is again not much lower than for snippets.

3 Conclusions

We demonstrated that in a Federated Web Search setting, result page relevance predictions are possible, based on snippet features alone. A detailed analysis in terms of the number of resources that provide relevant results, which features contributed the most to the results, etc., is proposed as future work, together with full-fledged result merging research for the current setting, based on snippet classification.

References

1. Arguello, J., Callan, J., Diaz, F.: Classification-based resource selection. In: Proceeding of the 18th ACM conference on Information and knowledge management (CIKM '09). ACM Press, New York, New York, USA (2009)
2. Arguello, J., Diaz, F., Callan, J.: Learning to Aggregate Vertical Results into Web Search Results Categories and Subject Descriptors. In: Proceeding of the 20th ACM conference on Information and knowledge management - CIKM 2011. pp. 201–210 (2011)
3. Clarke, C.L.A., Craswell, N., Soboroff, I., Cormack, G.V.: Overview of the TREC 2010 Web Track. TREC pp. 1–9 (2010)

4. Demeester, T., Nguyen, D., Trieschnigg, D., Develder, C., Hiemstra, D.: What Snippets Say about Pages in Federated Web Search. In: Proceedings of the 8th Asia Information Retrieval Society Conference (AIRS) (2012)
5. Nguyen, D., Demeester, T., Trieschnigg, D., Hiemstra, D.: Federated Search in the Wild: the Combined Power of over a Hundred Search Engines. Proceedings of the 21st ACM Conference on Information and Knowledge Management (CIKM) (2012)
6. Nigam, K., Lafferty, J., McCallum, A.: Using Maximum Entropy for Text Classification. IJCAI'99 Workshop on Information Filtering (1999)
7. Shokouhi, M., Li, L.: Federated Search. Foundations and Trends in Information Retrieval 5(1), 1–102 (2011)