DOI: http://dx.doi.org/10.14236/ewic/FDIA2015.5

Heterogeneous Information Access Through Result Composition

Horațiu Bota
School of Computing Science
University of Glasgow
18 Lilybank Gardens
Glasgow, UK
www.horatiubota.com
h.bota.1@research.gla.ac.uk

Modern search engines aggregate information from a variety of sources (e.g. images, videos) and return this information to users, merged into a single results page. Current aggregation techniques are limited to merging blocks of heterogeneous content into organic result rankings. We propose a new approach to search aggregation that takes into account result semantics and explicit searcher preferences in the form of *result composition*. Our findings suggest that result composition can be an effective search paradigm and can positively impact search behaviour in certain contexts.

Keywords: Web search, Search interfaces, User behaviour

1. INTRODUCTION

Over the past decades, we have seen an explosion of immediately accessible information. The first Internet search engine, Aliweb (Koster 1994), was released in 1994 and marked the beginning of a continuously expanding area of on-line activity: Web search. Today, with more than 80% of Internet surfers using search engines to satisfy their information needs (Shokouhi and Si 2011), search technology attracts notable attention from industry and research alike.

Modern search engines have started displaying results aggregated from heterogeneous information sources — also known as **verticals**, for example images, videos, news — in return to user queries. This search paradigm is primarily driven by the dynamics and volume of heterogeneous content on the Web, and is intended to support ambiguous information needs and varying user behaviour patterns (Zhou et al. 2012). Different approaches to aggregating information on the Web have been proposed and studied in prior work on *federated search*, *aggregated search* or *semantic web search*. In general, these approaches have been limited to searching multiple text collections simultaneously or merging results from different information sources within standard search results pages.

In light of these limitations, we argue that current approaches to aggregating results across different sources are insufficient and propose a new method of presenting users with a structured assembly of heterogeneous documents: *composite retrieval*. Instead of a ranked list of documents or a set of heterogeneous blocks, we want to provide users with semantically structured assemblies of documents, incorporated into a set of **composite information objects**.

Consider the following user information need "finding all information to plan a trip to Greece". Answering this information need typically involves submitting several queries to gather information about airports and visa policies, to read on-line reviews about hotels, and to check the geographic proximity of places to visit. Current search engines aggregate results from multiple verticals, however, the presentation of search results is limited to heterogeneous blocks. As the web has made available a large variety of verticals, it is becoming important to return to users "organised" results, containing information extracted from different sources of information. Doing so will not only support users in complex search tasks, but also allow them to understand the diversity of the information space and select what matters to them most. Furthermore, composite retrieval on the Web can potentially promote exploratory user behaviour and move the

focus of search away from providing results to users' queries, towards providing answers to users' questions.

Our work focuses on returning to users results "organised" into what we refer to as a *composite information objects*, where individual objects contain results extracted from various sources (e.g. videos, images and blogs), each object focusing on different aspects of a user's query.

In broad terms, our research aims to answer the following research questions:

1. Can composition of results lead to better performance, in terms of traditional IR relevance, than ranked lists of results?
2. From the users' perspective, what results should composite objects contain? What functions do individual results play in a composite information object?
3. How does the presence of a composite information object on a traditional search results page influence user behaviour and perceived task workload? Can composite objects make search tasks easier for the user?

The following sections provide a brief overview of the existing literature on aggregation based Web search (Section 2), a summary of our current work (Section 3), and an outline of future directions for our efforts (Section 4).

2. BACKGROUND

The concept of responding to search queries by presenting a composition of items has been proposed and investigated in a number of recent papers (Guo et al. 2012; Tran et al. 2011; Zhao et al. 2011). Many of the above papers have provided contributions on the theoretical side, studying the complexity of evaluating queries with constraints, and proposing different algorithmic formulations. Other works have focused on building systems that perform composite retrieval under a number of different semantics and targeting specific application domains, such as on-line shopping or travel planning (Basu Roy et al. 2010; Xie et al. 2010).

In the context of heterogeneous information access, aggregated search is widely used by modern search engines. Aggregated search is the task of retrieving information from a variety of resources (or verticals) and merging it into a single interface (Arguello et al. 2009; Zhou et al. 2012). Aggregated search can be compared to federated search (Shokouhi and Si 2011) (also known as distributed information retrieval), which deals with merging result rankings from different search engines into one single ranking list. The main challenges in aggregated search and federated search are resource selection and result merging. The former deals with deciding which sources of information contain the most relevant results to a given query and the latter deals with selecting a subset of items from relevant sources and presenting them as results. In aggregated search, the most common result presentation strategy consists of merging blocks of heterogeneous results into ranked lists of organic results. Similar to aggregated search, selecting and organising results from heterogeneous sources is the main focus of composite retrieval. However, rather than presenting the results of each selected vertical as a block of homogeneous items, composite retrieval aims to present results into cohesive information objects, where each object contains heterogeneous items (retrieved from several verticals).

Understanding user search behaviour is a key component of modelling and evaluating search engine performance. In the context of aggregated search, user behaviour has been shown to differ significantly compared to the more traditional *ten blue links* environment. For example, in a study analysing click-through rates in an aggregated search scenario, Sushmita et al. (2010) found users click more on vertical results that are relevant to the task, shown higher in ranking and more visually salient. Diaz et al. (2013) mined users' mouse movement interactions from a commercial search engine log and found that different results presentation strategies create different biases with respect to user attention and browsing sequence. All previously mentioned studies investigate aggregated search scenarios in which heterogeneous content is displayed in blocks of items, embedded into the organic Web results list.

Unlike aggregated search, result composition merges results from different sources into singular information objects to be presented on the results page. From a presentation perspective, composite information objects are similar to entity cards (Navalpakkam et al. 2013; Lagun et al. 2014), which are shown on existing search results pages in response to ambiguous or entity-specific user queries. Entity cards are related to our work because they are instances of composite information objects: they contain heterogeneous results, extracted from different sources, assembled using various semantic retrieval techniques, and shown in response to ambiguous user queries. There is limited understanding of user behaviour in entity-card search: to our knowledge, only two prior studies examined the effects of entity cards on user

interaction with search interfaces. With an eye and mouse tracking study, Navalpakkam et al. (2013) found that the flow of user attention on non-linear page layouts (with knowledge cards shown at the top-right corner of the SERP) is different from the widely believed, top-down linear examination of search results. In the context of mobile search, Lagun et al. (2014) performed a similar study in understanding how user attention is distributed between knowledge cards and web results in a mobile context.

3. CURRENT WORK

3.1. Retrieval Performance

To address our first research question, we adapted an existing composite retrieval framework (Amer-Yahia et al. 2013) to Web search. Due to the heterogeneous nature of the multi-vertical environment we explored, novel ways to model and estimate the various components of our proposed framework were developed. We used named entities (textual spots mapped to Wikipedia articles) to bridge the semantic gap between documents across verticals and developed algorithms to construct composite information objects. We applied our algorithms on a federated search test collection (Nguyen et al. 2012), which contained results from 108 search engines categorised into 11 different verticals.

Our results indicate that composite retrieval can significantly improve the performance over various current search paradigms, such as traditional "general web only" ranking, federated search ranking and aggregated search. The composite retrieval search paradigm we propose aims to promote a diverse information space for users to explore. For an exploratory task, rather than requiring searchers to issue multiple queries related to different aspects of their information need, issued to several vertical search engines, composite retrieval provides a unified page that consists of *relevant* objects focused on different aspects of a searcher's query. Our results have implications for work in heterogeneous information access and diversity in IR. An in-depth analysis of results is presented in Bota et al. (2014).

3.2. User Perspectives

To address our second research question, we designed and ran a user study to analyse user-generated composite objects. Our main objective was to determine how composite objects are manually generated by searchers. In particular, our interest was to analyse composite objects with respect to their topical focus, content and user-assessed characteristics (i.e. relevance, cohesion, diversity).

Our results show that, firstly, there is an agreement between users on the topical focus of composite objects — namely that different users construct composite objects which focus on similar aspects of a given topic. Secondly, we observe that composite objects contain documents that play different roles. For instance, central documents (or pivots), are assessed by users as being more relevant than other document within the composite object, and reflect the object title, whereas ornament documents are less relevant but provide value to searchers through composition with pivots. Finally, our results suggest that no clear hierarchy of user-assessed object characteristics can be determined and that, although explicit relevance is crucial in search, composition of diverse results can generate additional value to users. An in-depth analysis of the results is available in Bota et al. (2015).

3.3. Effects on User Behaviour

In addition to aggregating results from different sources (e.g. images, video, news), modern search engines have started displaying complex information objects, or entity cards (ECs), on the results page. As mentioned in section 2, entity search cards are instances of composite information objects. Entity cards are intended to enhance search experience in several ways: ***(i)*** they help searchers navigate diversified results, ***(ii)*** provide a summary of relevant content directly on the results page and ***(iii)*** support exploratory search by highlighting relevant entities associated with a given user query. Because we wanted to understand the effect of result composition on user search behaviour, we designed and ran a large-scale crowd-sourced user study, with more than one thousand unique searchers, in which we studied the effect of entity cards and their properties (relevance, cohesion and diversity) on search behaviour and perceived task workload.

Our results suggest that the presence of ECs has a strong effect on both the way users interact with search results and their perceived task workload. The results of our investigation indicate that ECs have significantly different effects in simple versus complex tasks. Furthermore, by manipulating EC properties (*content*, *coherence* and *diversity*), we uncover different effects and interactions between card properties on measures of search behaviour and workload.

4. FUTURE WORK

In terms of future work, many open questions remain. With regard to our first research question, more rigorous evaluation metrics, tailored to non-linear search environments need to be developed

in order to reliably investigate result composition performance.

Secondly, with regard to our second and third research questions, our work so far provides an extensive analysis of user behaviour in both a result composition scenario, and in an entity-card search scenario. However, because there is limited understanding of presentation strategies for composite information objects, we aim to investigate presentation optimisation strategies for result composition in Web search.

Given that mobile devices have become ubiquitous, and that Web search is increasingly prevalent on mobile devices, we intend to investigate the role result composition can play in mobile Web search. In particular, our objective is to understand whether user context (e.g. time, location, device size) can predict and explain the usefulness of composite information objects, as reflected by user engagement metrics, in mobile Web search.

5. CONCLUSION

As the Web has made available an enormous variety of textual and multimedia resources, people have started performing increasingly more complex search tasks, aimed at finding rich answers that require information extracted from various sources. To satisfy these complex information needs, modern search systems need to build solutions that aggregate information, taking into account users' intents and preferences. We argue that composition of results can provide users with a more structured approach to Web search. Our work so far suggests that returning composite information objects to users' search queries can not only provide a better search experience, in terms of traditional IR metrics, but can also positively impact user search behaviour and perceived task workload in certain contexts.

Many avenues for future work remain open. Fully understanding result composition requires the development of comprehensive evaluation metrics that take into account both the content and the presentation of composite objects. In addition, given the increasingly growing usage of mobile Web search, understanding the role composite objects play in various mobile contexts is crucial. We aim to address both these aspects in our future work.

Acknowledgements

The author would like to thank Dr. Ke Zhou and Prof. Joemon Jose for their invaluable support, and ELIAS (Evaluating Information Access Systems) for the scholarship to participate at ESSIR 2015 and the FDIA Symposium.

REFERENCES

Amer-Yahia, S., F. Bonchi, C. Castillo, E. Feuerstein, I. Méndez-Díaz, and P. Zabala (2013). Complexity and algorithms for composite retrieval. In *Proceedings of the 22nd international conference on World Wide Web companion*, pp. 79–80. International World Wide Web Conferences Steering Committee.

Arguello, J., F. Diaz, J. Callan, and J.-F. Crespo (2009). Sources of evidence for vertical selection. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, SIGIR '09, pp. 315–322.

Basu Roy, S., S. Amer-Yahia, A. Chawla, G. Das, and C. Yu (2010). Constructing and exploring composite items. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*, SIGMOD '10, New York, NY, USA, pp. 843–854. ACM.

Bota, H., K. Zhou, and J. Jose (2015). Exploring composite retrieval from the users' perspective. In A. Hanbury, G. Kazai, A. Rauber, and N. Fuhr (Eds.), *Advances in Information Retrieval*, Volume 9022 of *Lecture Notes in Computer Science*, pp. 13–24. Springer International Publishing.

Bota, H., K. Zhou, J. M. Jose, and M. Lalmas (2014). Composite retrieval of heterogeneous web search. In *Proceedings of the 23rd International Conference on World Wide Web*, WWW '14, New York, NY, USA, pp. 119–130. ACM.

Diaz, F., R. White, G. Buscher, and D. Liebling (2013). Robust models of mouse movement on dynamic web search results pages. ACM CIKM '13, pp. 1451–1460.

Guo, X., C. Xiao, and Y. Ishikawa (2012). Transactions on large-scale data- and knowledge-centered systems vi. Chapter Combination Skyline Queries, pp. 1–30. Berlin, Heidelberg: Springer-Verlag.

Koster, M. (1994). Aliweb - archie-like indexing in the web. *Computer Networks and ISDN Systems 27* (2), 175 – 182. Selected Papers of the First World-Wide Web Conference.

Lagun, D., C.-H. Hsieh, D. Webster, and V. Navalpakkam (2014). Towards better measurement of attention and satisfaction in mobile search. In *ACM SIGIR '14*, pp. 113–122.

Navalpakkam, V., L. Jentzsch, R. Sayres, S. Ravi, A. Ahmed, and A. Smola (2013). Measurement and modeling of eye-mouse behavior in the presence of nonlinear page layouts. In *ACM WWW '13*, pp. 953–964.

Nguyen, D., T. Demeester, D. Trieschnigg, and D. Hiemstra (2012).
Federated search in the wild: the combined power of over a hundred search engines. In *Proceedings of the 21st ACM international conference on Information and knowledge management*, pp. 1874–
1878.
ACM.

Shokouhi, M. and L. Si (2011). Federated search. *Foundations and Trends in Information Retrieval 5*(1), 1–102.

Sushmita, S., H. Joho, M. Lalmas, and R. Villa (2010). Factors affecting click-through behavior in aggregated search interfaces. In *ACM CIKM '10*, pp. 519–
528. ACM.

Tran, Q. T., C.-Y. Chan, and G. Wang (2011). Evaluation of set-based queries with aggregation constraints. In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management*, CIKM '11, New York, NY, USA, pp. 1495–1504. ACM.

Xie, M., L. V. Lakshmanan, and P. T. Wood (2010). Breaking out of the box of recommendations: from items to packages. In *Proceedings of the fourth ACM conference on Recommender systems*, RecSys '10, New York, NY, USA, pp. 151–158. ACM.

Zhao, B., X. Lin, B. Ding, and J. Han (2011). Texplorer: Keyword-based object search and exploration in multidimensional text databases. In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management*, CIKM '11, New York, NY, USA, pp. 1709–
1718.
ACM.

Zhou, K., R. Cummins, M. Lalmas, and J. M. Jose (2012). Evaluating aggregated search pages. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, pp. 115–124. ACM.