UvA-DARE (Digital Academic Repository)

CLEF 2017 dynamic search evaluation lab overview

Kanoulas, E.; Azzopardi, L.

DOI

[10.1007/978-3-319-65813-1_31](https://doi.org/10.1007/978-3-319-65813-1_31)

Publication date

2017

Document Version

Final published version

Published in

Experimental IR Meets Multilinguality, Multimodality, and Interaction

License

Article 25fa Dutch Copyright Act

[Link to publication](#)

Citation for published version (APA):

Kanoulas, E., & Azzopardi, L. (2017). CLEF 2017 dynamic search evaluation lab overview. In G. J. F. Jones, S. Lawless, J. Gonzalo, L. Kelly, L. Goeuriot, T. Mandl, L. Cappellato, & N. Ferro (Eds.), *Experimental IR Meets Multilinguality, Multimodality, and Interaction: 8th International Conference of the CLEF Association, CLEF 2017, Dublin, Ireland, September 11–14, 2017 : proceedings* (pp. 361–366). (Lecture Notes in Computer Science; Vol. 10456). Springer. https://doi.org/10.1007/978-3-319-65813-1_31

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

CLEF 2017 Dynamic Search Evaluation Lab Overview

Evangelos Kanoulas¹(✉) and Leif Azzopardi²

¹ Informatics Institute, University of Amsterdam, Amsterdam, Netherlands
E.Kanoulas@uva.nl

² Computer and Information Sciences, University of Strathclyde, Glasgow, UK
leif.azzopardi@strath.ac.uk

Abstract. In this paper we provide an overview of the first edition of the CLEF Dynamic Search Lab. The CLEF Dynamic Search lab ran in the form of a workshop with the goal of approaching one key question: how can we evaluate dynamic search algorithms? Unlike static search algorithms, which essentially consider user request's independently, and which do not adapt the ranking w.r.t the user's sequence of interactions, dynamic search algorithms try to infer from the user's intentions from their interactions and then adapt the ranking accordingly. Personalized session search, contextual search, and dialog systems often adopt such algorithms. This lab provides an opportunity for researchers to discuss the challenges faced when trying to measure and evaluate the performance of dynamic search algorithms, given the context of available corpora, simulations methods, and current evaluation metrics. To seed the discussion, a pilot task was run with the goal of producing search agents that could simulate the process of a user, interacting with a search system over the course of a search session. Herein, we describe the overall objectives of the CLEF 2017 Dynamic Search Lab, the resources created for the pilot task and the evaluation methodology adopted.

1 Introduction

Information Retrieval (IR) research has traditionally focused on serving the best results for a single query – so-called ad-hoc retrieval. However, users typically search iteratively, refining and reformulating their queries during a session. IR systems can still respond to each query in a session independently of the history of user interactions, or alternatively adopt their model of relevance in the context of these interactions. A key challenge in the study of algorithms and models that dynamically adapt their response to a user's query on the basis of prior interactions is the creation of suitable evaluation resources and the definition of suitable evaluation metrics to assess the effectiveness of such IR algorithms. Over the years various initiatives have been proposed which have tried to make progress on this long standing challenge.

The TREC Interactive track [7], which ran between 1994 and 2002, investigated the evaluation of interactive IR systems and resulted in an early standardization of the experimental design. However, it did not lead to a reusable test

collection methodology. The High Accuracy Retrieval of Documents (HARD) track [1] followed the Interactive track, with the primary focus on single-cycle user-system interactions. These interactions were embodied in clarification forms which could be used by retrieval algorithms to elicit feedback from assessors. The track attempted to further standardize the retrieval of interactive algorithms, however it also did not lead to a reusable collection that supports adaptive and dynamic search algorithms. The TREC Session Track [3], which ran from 2010 through 2014, made some headway in this direction. The track produced test collections, where included with the topic description was the history of user interactions with a system, that could be used to improve the performance of a given query. While, this mean adaptive and dynamic algorithms could be evaluated for one iteration of the search process, the collection's are not suitable for assessing the quality of retrieval over an entire session. In 2015, the TREC Tasks track [11, 12], a different direction was taken, where the test collection provides queries for which all possible sub-tasks did to be inferred, and the documents relevant to those sub-tasks identified. Even though the produced test collections could be used in testing whether a system can help the user to perform a task end-to-end, the focus was not on adapting and learning from the user's interactions as in the case of dynamic search algorithms.

In the related domain of dialogue systems, the advancement of deep learning methods has led to a new generation of data-driven dialog systems. Broadly-speaking, dialog systems can be categorized along two dimensions, (a) goal-driven vs. non-goal-driven, and (b) open-domain vs. closed domain dialog systems. Goal-driven open-domain dialog systems are in par with dynamic search engines: as they seek to provide assistance, advice and answers to a user over unrestricted and diverse topics, helping them complete their task, by taking into account the conversation history. While, a variety of corpora is available for training such dialog systems [10], when it comes to the evaluation, the existing corpora are inappropriate. This is because they only contain a static set of dialogues and any dialog that does not develop in a way similar to the static set cannot be evaluated. Often, the evaluation of goal-driven dialogue systems focuses on goal-related performance criteria, such as goal completion rate, dialogue length, and user satisfaction. Automatically determining whether a task has been solved however is an open problem, while task-completion is not the only quality criterion of interest in the development of dialog systems. Thus, simulated data is often generated by a simulated user [4, 5, 9]. Given a sufficiently accurate model of how user's converse, the interaction between the dialog system and the user can be simulated over a large space of possible topics. Using such data, it is then possible to deduce the desired metrics. This suggests that a similar approach could be taken in the context of interactive IR. However, while significant effort has been made to render the simulated data as realistic as possible [6, 8], generating realistic user simulation models remains an open problem.

2 Lab Overview

Essentially, the CLEF Dynamic Tasks Lab attempts focus attention towards building a bridge between batch TREC-style evaluation methodology and the Interactive Information Retrieval evaluation methodology - so that dynamic search algorithms can be evaluated using re-usable test collections.

The objectives of the lab is threefold:

1. to devise a methodology for evaluating dynamic search algorithms by exploring the role of simulation as a means to create re-usable test collections
2. to develop evaluation metrics that measure the quality during the session (and at different stages) and at the end of the session (overall measures of quality).
3. to develop algorithms that can provide an optimal response in an interactive retrieval setup.

The focus of the CLEF 2017 Dynamic Tasks Lab is to provide a forum that can help foster research around the evaluation of dynamic retrieval algorithms. The Lab, in the form of a Workshop, solicits the submission of two types of papers: (a) position papers, and (b) data papers. Position papers focus on evaluation methodologies for assessing the quality of search algorithms with the user in the loop, under two constraints: any evaluation framework proposed should allow the (statistical) reproducibility of results, and lead to a reusable benchmark collection. Data Papers describe test collections or data sets suitable for guiding the construction of dynamic test collections, tasks and evaluation metrics.

3 Pilot Task

Towards the aforementioned goals of generating simulation data the CLEF 2017 Dynamic Tasks Lab ran a pilot task in the context of developing Task Completion Engines [2] and Intelligent Search Agents [6]. Task Completion Engines and Autonomous Search Agents are being developed to help users in acquire information in order to make a decision and complete a search task. At the same time such Intelligent Search Agents, encode a model of a user, and so present the potential to simulate users submitting queries, which can enable the evaluation of dynamic search algorithms. Such engines/agents need to work with a user to ascertain their information needs, then perform their own searches to dynamically identify relevant material, which will be useful in completing a particular task. For example, consider the task of organizing a wedding. There are many different things that need to be arranged and ordered, e.g. a venue, flowers, catering, gift list, dresses, car hire, hotels, etc. Finding relevant sites and resources requires numerous searches and filtering through many documents/sites. A search agent could help to expedite the process by finding the relevant sites to visit, while a task completion engine would provide a structured interface to help complete the process.

In this year's Dynamic Search Task Track, the task can be interpreted in one of two ways:

1. to generate a series of queries that a search agent would issue to a search engine, in order to compile a set of links useful for the user. This set might be presented to the user or used for further processing by a task completion engine; or
2. to generate a series of query suggestions that a search engine would recommend to the user, and thus the suggested course of interaction.

As a starting point, for building a test collection, we first consider how people look for information required to complete various casual leisure and work tasks. The history of queries and interactions are then used as a reference point during the evaluation to determine if agents/dynamic search algorithms/query suggestion algorithms that can generate queries that are like those posed by people. And thus, see how well human like queries can be generated for suggestions)? Thus the focus of the track is on query generation and models of querying, based on task and interaction history.

3.1 Task Description and Data Sets

Starting with an initial query for a given topic, the task is to generate a series of subsequent or related queries. The data used is TREC ClueWeb Part B test collection and the topics used are sourced from the Session Track 2014 [3]. A fixed search engine was setup, where ClueWeb was indexed using ElasticSearch. The title, url, page contents were indexed, along with the spam rank and page rank of each document. The ElasticSearch API was then provided as the “search engine” that the agent or person is using to undertake each task/topic. From the Session Track 2014 topics, a subset of 26 topics were selected out of the original 50, based on the following criteria: there were four or more queries associated with the topic, where the subsequent interaction on each query lead to identifying at least one TREC relevant document. These were considered, good or useful, queries i.e. they helped identify relevant material. The set of “good” queries were with-held as the relevant set. The TREC topic title, was provided to participants as the initial seed query i.e. the first query in the session.

Interaction data was then provided to provide simulated interaction with queries issued to the search engine. It was anticipated that the simulated clicks could be used by the algorithms to help infer relevance of the documents. This data could be used as (a) a classifier providing relevance decisions regarding observed items in the result list, or (b) as clicks that the user performed when viewing the results of a query (i.e. given a query, assume that this is what the user clicks on, to help infer the next query). A set of judgments/click was generated based on the probability of Session Track users clicks data, conditioned by relevance (i.e. the probability of a click, if then document was TREC Relevant or TREC Non-Relevant).

The task, then, was to provide a list of query suggestions/recommendations along with a list of up to 50 documents.

3.2 Participants

Two teams participated in this Pilot Task: (1) Webis (Web Technology and Information Systems, Bauhaus-Universität Weimar), and (2) TUW (Vienna University of Technology).

Webis. The general research idea of the Webis team contribution is to evaluate query suggestions in form of key-queries for clicked documents. A key-query for a document set D is a query that returns the documents from D among the top- k ranks. Our query suggestion approach derives key-queries for pairs of documents previously clicked by the user. The assumption then is that the not-already-clicked documents in the top results of the key-queries could also be interesting to the user. The key-query suggestions thus focus on retrieving more documents similar to the ones already clicked.

TUW. The general research idea of the TUW team contribution is to leverage the structure of Wikipedia articles to understand search tasks. Their assumption is that human editors carefully choose meaningful section titles to cover the various aspects of an article. Their proposed search agent explores this fact, being responsible for two tasks: (1) identifying the key Wikipedia articles related to a complex search task, and (2) selecting section titles from those articles. TUW contributed 5 runs, by (a) only using the queries provided; (b) manually choosing which Wikipedia sections to use; (c) automatically choosing the top-5 Wikipedia sections to use; (d) automatically ranking Wikipedia section using word2vec; and (e) automatically ranking Wikipedia sections using word2vec and using a naive Bayes classifier trained on the past qrels to decide if a document was relevant or not.

3.3 Evaluation

Given that evaluation is an open problem, the lab was also open to different ways in which to evaluate this task. Some basic measures that were employed are:

- Query term overlap: how well do the query terms in the suggestions match with the terms used
- Query likelihood: how likely are the queries suggested given the model of relevance.
- Precision and Recall based measures on the set of documents retrieved.
- Suggest your own measure: how to measure the usefulness and value of queries is a rather open question. So how can we evaluate how good a set of queries are in relation to completing a particular task?

During the lab, we will discuss the various challenges of constructing reusable test collections for evaluating dynamic search algorithms and how we can develop appropriate evaluation measures.

Acknowledgements. This work was partially supported by the Google Faculty Research Award program and the Microsoft Azure for Research Award program (CRM:0518163). All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors. We would also like to thank Dr. Guido Zuccon for setting up the Elasticsearch API.

References

1. Allan, J.: Hard track overview in TREC 2003 high accuracy retrieval from documents. Technical report, DTIC Document (2005)
2. Balog, K.: Task-completion engines: a vision with a plan. In: SCST@ECIR (2015)
3. Carterette, B., Clough, P.D., Hall, M.M., Kanoulas, E., Sanderson, M.: Evaluating retrieval over sessions: the TREC session track 2011–2014. In: Perego, R., Sebastiani, F., Aslam, J.A., Ruthven, I., Zobel, J. (eds.) *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2016, Pisa, Italy, 17–21 July 2016*, pp. 685–688. ACM (2016). <http://doi.acm.org/10.1145/2911451.2914675>
4. Georgila, K., Henderson, J., Lemon, O.: User simulation for spoken dialogue systems: learning and evaluation. In: *Interspeech* pp. 1065–1068 (2006)
5. Jung, S., Lee, C., Kim, K., Jeong, M., Lee, G.G.: Data-driven user simulation for automated evaluation of spoken dialog systems. *Comput. Speech Lang.* **23**(4), 479–509 (2009). <http://dx.doi.org/10.1016/j.csl.2009.03.002>
6. Maxwell, D., Azzopardi, L.: Agents, simulated users and humans: An analysis of performance and behaviour. In: *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management, CIKM 2016*, pp. 731–740 (2016)
7. Over, P.: The TREC interactive track: an annotated bibliography. *Inf. Process. Manage.* **37**(3), 369–381 (2001)
8. Pääkkönen, T., Kekäläinen, J., Keskustalo, H., Azzopardi, L., Maxwell, D., Järvelin, K.: Validating simulated interaction for retrieval evaluation. *Inf. Retrieval J.*, 1–25 (2017)
9. Pietquin, O., Hastie, H.: A survey on metrics for the evaluation of user simulations. *Knowl. Eng. Rev.* **28**(01), 59–73 (2013)
10. Serban, I.V., Lowe, R., Henderson, P., Charlin, L., Pineau, J.: A survey of available corpora for building data-driven dialogue systems. *CoRR* abs/1512.05742 (2015). <http://arxiv.org/abs/1512.05742>
11. Verma, M., Yilmaz, E., Mehrotra, R., Kanoulas, E., Carterette, B., Craswell, N., Bailey, P.: Overview of the TREC tasks track 2016. In: Voorhees, E.M., Ellis, A. (eds.) *Proceedings of The Twenty-Fifth Text REtrieval Conference, TREC 2016, Gaithersburg, Maryland, USA, 15–18 November 2016*, vol. Special Publication 500–321. National Institute of Standards and Technology (NIST) (2016). <http://trec.nist.gov/pubs/trec25/papers/Overview-T.pdf>
12. Yilmaz, E., Verma, M., Mehrotra, R., Kanoulas, E., Carterette, B., Craswell, N.: Overview of the TREC 2015 tasks track. In: Voorhees, E.M., Ellis, A. (eds.) *Proceedings of The Twenty-Fourth Text REtrieval Conference, TREC 2015, Gaithersburg, Maryland, USA, 17–20 November 2015*, vol. Special Publication 500–319. National Institute of Standards and Technology (NIST) (2015). <http://trec.nist.gov/pubs/trec24/papers/Overview-T.pdf>