

Query Log Mining for Inferring User Tasks and Needs

August 1, 2016

Presenters & Affiliations

1. Emine Yilmaz
Associate Professor, University College London
Consultant, Microsoft Research, Cambridge
2. Rishabh Mehrotra
Final year PhD student
University College London

1 Contact Details

Email: emine.yilmaz@ucl.ac.uk; r.mehrotra@cs.ucl.ac.uk

Address: 5th Floor, One Euston Square, UCL, London Media Technology Campus (UCL/BBC), 40 Melton Street, London NW1 2FD

Contact: Tel: +44 (0)20 3549 5267

2 Tutorial Details

Type of tutorial: Half Day

Abstract:

Search behavior, and information behavior more generally, is often motivated by tasks that prompt search processes that are often lengthy, iterative, and intermittent, and are characterized by distinct stages, shifting goals and multitasking.

Developing a comprehensive understanding of user's needs and tasks would help in providing better support and recommendations to users based on their contextual information and as a result, help users accomplish their information needs. In this tutorial, we begin by discussing recent advancements towards inferring user needs and tasks from query logs and present analytical results which highlight the importance of considering tasks as the focal unit of modeling search behavior. We consider the challenge of extracting tasks from a given collection of search log data and present some recently proposed task extraction techniques which rely on recent advancements in bayesian non parametrics, word embeddings and memory networks. We further show how tasks could be useful for personalisation of search results.

Target Audience:

Intermediate researchers, industry technologists and practitioners. We are expecting a general background in information retrieval.

Learning outcomes:

The tutorial is aimed at introducing practitioners to the methods that can be used to infer needs of users by identifying the type of tasks people tend to use search engines for. The main focus of the tutorial is (i) discussing novel algorithmic approaches to extract tasks from log data, and (ii) imparting knowledge on how to leverage this task information for various applications. The tutorial would help practitioners from various different subfields including search, user modeling, personalization and recommender systems develop a comprehensive understanding of user tasks and needs and equip them with the necessary mathematical and analytical tools required to extract task information to be used in their domain of choice.

2.1 Outline of Tutorial

1. Introduction

- (a) Evolution of search
 - (b) Functionality Levels of Search offerings
2. From Sessions to Tasks
 - (a) Understanding Sessions
 - i. Session detection
 - ii. Session based retrieval
 - iii. Session based personalization
 - (b) Understanding Tasks
 - i. Multitasking
 - ii. User groups & behavioral differences
 - iii. Topical heterogeneity
 3. Task extraction algorithms
 - (a) Latent task extraction
 - (b) Deconstructing Complex Tasks
 4. Hierarchies of Tasks & Subtasks
 5. Applications
 - (a) User modelling & personalization
 - (b) Digital Assistants
 6. Conclusion & Open Research questions

2.2 Tutorial Description

Search behavior, and more generally, information-seeking behavior is often motivated by tasks, that prompt search processes which are lengthy, iterative and intermittent. Such search tasks are also characterized by distinct stages, shifting goals and multitasking behaviors. Most research in information retrieval has focused on identifying documents relevant to the query submitted by the user, as opposed to focusing on the underlying information need or task that have led the user issue the query.

In this tutorial, we present the work that has been done on inferring tasks and needs of users given queries previously issues by users. We present an overview of the work that has been done on (i) understanding searcher’s behaviors, (ii) developing task extraction techniques and (iii) the benefits of task information via improved personalization.

We start with discussing recent research based on using search sessions and using session information for devising improved retrieval systems. We further argue that sessions are not enough to capture user needs and describe methods that can identify and extract search tasks. There has been some recent work on task extraction; however, most existing task-extraction methods treat search tasks as structure-less clusters which inherently lack insights about the presence or demarcation of subtasks associated with individual search tasks. A more naturalistic viewpoint would involve considering complex search tasks as being decomposed into more focused subtasks. We further discuss some initial techniques that aim at representing tasks as a hierarchy. Finally, we focus on personalization and show how sessions and tasks can be useful for personalizing search results.

2.2.1 Understanding Sessions

Search behavior can be analyzed over time to identify queries that express the same underlying information need. Most previous work has focused on search behavior analysis and prediction within a single search session. Related queries within a session have been referred to as being part of a query chain [8] or search goal [1]. Raman *et al.* [9] investigated whole session relevance by considering intrinsic diversity in search results. He *et al.* [2] proposed an algorithm to segment a query stream into sessions by detecting topical shifts between the queries. Hassan *et al.* [] modeled session-level search goals using hidden Markov models. Single-session analyses have also been used for various search-related tasks such as query suggestion, interactive feedback, and query disambiguation. Another related research direction concerns the construction of tours and trails to guide searchers resource selection decisions during the search process. Trails have also been proposed as a way to guide users through the specific steps required to accomplish search tasks. Singla *et al.* [10] proposed trailfinding methods to support Web search by identifying query-relevant trails from logs that could be shown to complement or replace traditional search result lists.

2.2.2 Understanding Searcher’s Task Behavior

While a major share of prior work have considered search sessions as the focal unit of analysis for seeking behavioral insights, search tasks are emerging as a competing perspective in this space. In a recent work ([3]), we quantify multi-tasking behavior of web search users and show that over 50% of search sessions have more than 2 tasks. Further, we provide a method to categorize users into focused, multi-taskers or supertaskers depending on their level of task-multiplicity and show that the search effort expended by these users varies across the groups. Additionally, in a follow up work ([4]) we relate user’s multitasking propensities to tasks and topics. Specifically, we analyze user-disposition, topic and user-interest level heterogeneities that are prevalent in search task behavior. We find that not only do users have varying propensities to multi-task, they also search for distinct topics across single-task and multi-task sessions. The findings from our analysis provide useful insights about task-multiplicity in an online search environment and hold potential value for search engines that wish to personalize and support search experiences of users based on their task behavior.

2.2.3 Extracting Hierarchies of Tasks & Subtasks

When engaged in complex search tasks, users require holistic support in determining the required task activities. Most research in information retrieval has focused on identifying documents relevant to the query submitted by the user, as opposed to focusing on the underlying information need or task that have led the user issue the query. In this part of the tutorial, we discuss the methods that have been devised in inferring search tasks from query logs. Most existing task extraction systems focus on representing tasks as flat structures. However, tasks often tend to have multiple subtasks associated with them and a more naturalistic viewpoint would involve viewing query logs as hierarchies of tasks with complex search tasks being decomposed into more focused sub-tasks. To this end, we discuss an efficient Bayesian nonparametric model for extracting hierarchies of such tasks & subtasks. The approach makes use of the multi-relational aspect of query associations which are important in identifying query-task associations. We describe a greedy agglomerative model selection algorithm based on the Gamma-Poisson conjugate mixture that take just one pass through the data to learn a fully probabilistic, hierarchical model of trees that is capable of learning trees with arbitrary branching structures as opposed to the more common binary structured trees. Further details of the work are available in our research paper [6].

2.2.4 Task Based Personalization

In order to demonstrate the usefulness of using task information in retrieval, in a recent work [5], [7] we presented a novel approach to couple user’s topical interest information with their search task information & their term usage behavior to learn a joint user representation technique. We demonstrated that coupling user’s task information with their topical interests indeed helps us build better user models. We show through extensive experimentation that our task based method outperforms existing query term based and topical interest based user representation methods.

2.2.5 Applications

Inferring user needs through search sessions and tasks can have various different applications and in this part of the tutorial we will describe some possible applications of inferring user needs and tasks in devising digital assistants.

2.3 Tutorial logistics/materials

We intend to release the slides of the tutorial to the participants. We would also make available all related material, including compiled references and datasets on a dedicated website.

3 Speaker Bio

Emine Yilmaz:

Emine Yilmaz is an associate professor in the Department of Computer Science University College London, a faculty fellow of the Alan Turing Institute on Data Science and a research consultant for Microsoft Research Cambridge. She is the recipient of the Karen Sparck Jones Award in 2015 and the Google Faculty Award in 2014/15. Her main interests are evaluating quality of retrieval systems, modelling user behaviour, learning to rank, and inferring user needs while using search engines. She has published research papers extensively at major information retrieval venues such as SIGIR, CIKM, WWW and WSDM. She has previously given keynote talks at various conferences & workshops including at ECIR 2016 along with several tutorials on various Information Retrieval topics at the CHIIR 2016, SIGIR 2015, SIGIR 2012 and SIGIR 2010 Conferences and at the RuSSIR/EDBT Summer School in 2011. She has also organized several workshops on Crowdsourcing (WSDM2011, SIGIR 2011 and SIGIR 2010) and User Modelling for Retrieval Evaluation (SIGIR 2013). She has served as one of the organizers of the ICTIR Conference in 2009, as the demo chair for the ECIR Conference in 2013, and as the PC chair for the SPIRE 2015 conference. She is also a co-coordinator of the Tasks Track in TREC 2015 and 2016. .

Rishabh Mehrotra:

Rishabh Mehrotra is a final year PhD student at University College London partially supported by a Google Faculty Research Award. His current research focuses on inference of search tasks from query logs and their applications. Some of his recent work has been published at top conferences including NAACL, SIGIR, WWW, CIKM, RecSys and WSDM. During his PhD research, Rishabh has received various grants including the Y. Ali grant and the Donald B Crouch Grant. He has given over 10 guest and invited talks at various Machine Learning meetups, research groups and companies. Over the past few years, Rishabh has been working closely with leading industrial researchers at Microsoft Research, LinkedIn & NICTA on interesting machine learning & data science projects. He has supervised over 5 Masters thesis and has served as a reviewer for top tier conferences and workshops. He is also a co-coordinator of the Tasks Track in TREC 2015 and 2016.

References

- [1] A. Hassan. A semi-supervised approach to modeling web search satisfaction. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, pages 275–284. ACM, 2012.
- [2] D. He, A. Göker, and D. J. Harper. Combining evidence for automatic web session identification. *Information Processing & Management*, 38(5):727–742, 2002.
- [3] R. Mehrotra, P. Bhattacharya, and E. Yilmaz. Characterizing users’ multi-tasking behavior in web search. In *Proceedings of the 2016 ACM on Conference on Human Information Interaction and Retrieval*, pages 297–300. ACM, 2016.
- [4] R. Mehrotra, P. Bhattacharya, and E. Yilmaz. Sessions; tasks & topics - uncovering behavioral heterogeneities in online search behavior. In *Proceedings of the 2016 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2016.
- [5] R. Mehrotra and E. Yilmaz. Terms, topics & tasks: Enhanced user modelling for better personalization. In *Proceedings of the 2015 International Conference on The Theory of Information Retrieval*, pages 131–140. ACM, 2015.
- [6] R. Mehrotra and E. Yilmaz. Towards hierarchies of search tasks & subtasks. In *Proceedings of the 24th International Conference on World Wide Web Companion*, pages 73–74. International World Wide Web Conferences Steering Committee, 2015.
- [7] R. Mehrotra, E. Yilmaz, and M. Verma. Task-based user modelling for personalization via probabilistic matrix factorization. In *RecSys Posters*, 2014.
- [8] F. Radlinski and T. Joachims. Query chains: learning to rank from implicit feedback. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*, pages 239–248. ACM, 2005.
- [9] K. Raman, P. N. Bennett, and K. Collins-Thompson. Toward whole-session relevance: exploring intrinsic diversity in web search. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*, pages 463–472. ACM, 2013.
- [10] A. Singla, R. White, and J. Huang. Studying trailfinding algorithms for enhanced web search. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, pages 443–450. ACM, 2010.