

LearnIR: WSDM 2018 Workshop on Learning from User Interactions

Rishabh Mehrotra
Spotify Research
London, United Kingdom
rishabhm@spotify.com

Ahmed Hassan Awadallah
Microsoft Research
Redmond, WA
hassanam@microsoft.com

Emine Yilmaz
University College London
London, United Kingdom
emine.yilmaz@ucl.ac.uk

ACM Reference format:

Rishabh Mehrotra, Ahmed Hassan Awadallah, and Emine Yilmaz. 2018. LearnIR: WSDM 2018 Workshop on Learning from User Interactions. In *Proceedings of 11th ACM International Conference on Web Search and Data Mining, February 5–9, 2018, Marina Del Rey, CA, USA, (WSDM 2018)*, 2 pages. DOI: <https://doi.org/10.1145/3159652.3160598>

Website: <https://task-ir.github.io/wsdm2018-learnIR-workshop/>

1 MOTIVATION

While users interact with online services (e.g. search engines, recommender systems, conversational agents), they leave behind fine grained traces of interaction patterns. The ability to understand user behavior, record and interpret user interaction signals, gauge user satisfaction and incorporate user feedback gives online systems a vast treasure trove of insights for improvement and experimentation. More generally, the ability to learn from user interactions promises pathways for solving a number of problems and improving user engagement and satisfaction.

Understanding and learning from user interactions involves a number of different aspects - from understanding user intent and tasks, to developing user models and personalization services. A user's understanding of their need and the overall task develop as they interact with the system. Supporting the various stages of the task involves many aspects of the system, e.g. interface features, presentation of information, retrieving and ranking. Often, online systems are not specifically designed to support users in successfully accomplishing the tasks which motivated them to interact with the system in the first place. Beyond understanding user needs, learning from user interactions involves developing the right metrics and experimentation systems, understanding user interaction processes, their usage context and designing interfaces capable of helping users.

Learning from user interactions becomes more important as new and novel ways of user interactions surface. There is a gradual shift towards searching and presenting the information in a conversational form. Chatbots, personal assistants in our phones and eyes-free devices are being used increasingly more for different purposes, including information retrieval and exploration. With improved speech recognition and information retrieval systems, more

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

WSDM 2018, February 5–9, 2018, Marina Del Rey, CA, USA

© 2018 Copyright held by the owner/author(s). ISBN 978-1-4503-5581-0/18/02.

DOI: <https://doi.org/10.1145/3159652.3160598>

and more users are increasingly relying on such digital assistants to fulfill their information needs and complete their tasks. Such systems rely heavily on quickly learning from past interactions and incorporating implicit feedback signals into their models for rapid development.

2 SCOPE

Learning from User Interactions will be a highly interactive full day workshop that will provide a forum for academic and industrial researchers working at the intersection of user understanding, search tasks, conversational IR and user interactions. The purpose is to provide an opportunity for people to present new work and early results, brainstorm different use cases, share best practices, and discuss the main challenges facing this line of research.

- (1) **User Needs & Tasks Understanding:** User intent analysis/prediction; User goals & missions; Task identification; Task aware suggestions & recommendations
- (2) **User Modeling & Personalization:** Short and Long-term User Modelling; Personalization; Diversification; Coherence
- (3) **Metrics and Evaluation :** Metrics based on user interactions; User engagement metrics design; Evaluation mechanisms; User satisfaction prediction; Controlled laboratory study Online metrics Test collection
- (4) **User Interaction Processes & Context :** User Journey Optimization; Evolution of search process; Stages of user interactions; User journey through the system; Leveraging contextual signals; Learning for user interaction optimization: algorithms; frameworks & system designs
- (5) **Intelligent interface designs:** Adaptive personal digital assistants; Tailored decision support; Adaptive collaboration support
- (6) **Applications:** Conversational search, chatbots, digital assistants; Contextual Advertising; E-commerce recommendations; Customer Support; Intelligent interfaces; Personal search; Case studies of real world implementations

3 EXPECTED OUTCOMES

This is a workshop proper where discussion is central, and all attendees are active participants. We intend to host a number of keynote and invited talks from experts in related fields to provide workshop participants with a crisp overview of work on tasks and other related domains. Given the apt industrial utility of the workshop topic, we intend to keep a balanced mix of speakers from industry and academia alike.

Contributed papers will either have a short presentation or be showcased in an interactive poster session during the coffee breaks.

Beyond traditional paper & poster presentations, we will organize breakout sessions with predefined themes to foster a more productive discussion on what we learned, concrete plans for the next year, and a roadmap for the longer term.

The results will be disseminated in various ways:

- A high quality, peer reviewed workshop proceedings, published in the CEUR workshop proceedings series.
- A report on the results of the workshop in the ACM SIGIR Forum.
- Depending on the outcome, we will consider a special issue in an appropriate journal.
- Last, but not least, the results can be fed into the running tracks at TREC, CLEF, and other evaluation campaigns.

4 RELATED WORKSHOPS

There have been several workshops focused on user interactions, user intents & task, including CHI 2012 workshop on End-user Interactions with Intelligent Systems [6], Larsen, Lioma and de Vries's Task-based and Aggregated Search Workshop [5] and the Second Strategic Workshop on Information Retrieval in Lorne (SWIRL) [1]. Task-based search was discussed within the context of several larger themes identified by SWIRL participants, and was also presented briefly as a mini-theme. A NII Shonan Meeting, which was held in Japan in October 2012, focused on whole-session evaluation of interactive information retrieval systems. Some of the organizers were involved in a closely related workshop on Supporting Complex Search Tasks [2, 4] at CHIIR 2017 and SIGIR 2011. The NSF Task-Based Information Search Systems Workshop [3] was an invite-only workshop which fostered initial discussions around task understanding. Also related are the recent TREC Tasks Track [7, 8].

5 LIST OF ORGANIZERS

The following organizers are involved in organizing the workshop:

Rishabh Mehrotra

Spotify Research

Rishabh Mehrotra is a Research Scientist at Spotify Research and recently finished his PhD at University College London partially supported by a Google Faculty Research Award. His PhD research focused on inference of search tasks from query logs and their applications. Beyond tasks, his research focuses on user modelling & personalization, counterfactual analysis and deep learning for modelling user satisfaction. Some of his recent work has been published at top conferences including WWW, SIGIR, NAACL, CIKM, RecSys and WSDM. He has given many guest and invited talks at various conferences, Machine Learning meetups, research group seminars and industrial research events. He has supervised over 10 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, 2016 and 2017 and co-tutored tutorials at Search Solutions 2017 and CIKM 2017.

Ahmed Hassan Awadallah

Microsoft Research

Ahmed leads a team of scientists/engineers in Microsoft Research

Technologies; a new R&D organization created to focus on large-scale company technology initiatives. His work is focused on creating new technologies and experiences with end-to-end integration with current and future products in the areas of intelligence assistance, search and productivity improvement. A key part of his work involves using Machine Learning to model large scale text data and user behavior data with applications to user modeling, quality evaluation, ranking, recommendation, personalization and intelligent task completion assistance. In the past, Ahmed has organized a number of workshops including ACL workshop on Graph based methods (TexGraphs-7).

Emine Yilmaz

University College London

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 is also a co-coordinator of the Tasks Track in TREC 2015 and 2016.

Steering Committee

We're thankful to the members of our steering committee:

- (1) Milad Shokouhi (Microsoft)
- (2) Fernando Diaz (Spotify)
- (3) Filip Radlinski (Google Research)
- (4) Evangelos Kanoulas (University of Amsterdam)

REFERENCES

- [1] James Allan, Bruce Croft, Alistair Moffat, and Mark Sanderson. 2012. Frontiers, challenges, and opportunities for information retrieval: Report from SWIRL 2012 the second strategic workshop on information retrieval in Lorne. In *ACM SIGIR Forum*, Vol. 46. ACM, 2–32.
- [2] Nicholas J Belkin, Charles LA Clarke, Ning Gao, Jaap Kamps, and Jussi Karlgren. 2012. Report on the sigir workshop on entertain me: supporting complex search tasks. In *ACM SIGIR Forum*, Vol. 45. ACM, 51–59.
- [3] Diane Kelly, Jaime Arguello, and Robert Capra. 2013. NSF workshop on task-based information search systems. In *ACM SIGIR Forum*, Vol. 47. ACM, 116–127.
- [4] Marijn Koolen, Jaap Kamps, Toine Bogers, Nicholas J Belkin, Diane E Kelly, and Emine Yilmaz. 2017. Report on the Second Workshop on Supporting Complex Search Tasks. (2017).
- [5] Birger Larsen, Christina Lioma, and Arjen de Vries. 2012. Report on TBAS 2012: workshop on task-based and aggregated search. In *ACM SIGIR Forum*, Vol. 46. ACM, 71–77.
- [6] Simone Stumpf, Margaret Burnett, Volkmar Pipek, and Weng-Keen Wong. 2012. End-user interactions with intelligent and autonomous systems. In *CHI '12 Extended Abstracts on Human Factors in Computing Systems*. ACM, 2755–2758.
- [7] Manisha Verma, Emine Yilmaz, Rishabh Mehrotra, Evangelos Kanoulas, Ben Carterette, Nick Craswell, and Peter Bailey. 2016. Overview of the TREC Tasks Track 2016. In *TREC*.
- [8] Emine Yilmaz, Manisha Verma, Rishabh Mehrotra, Evangelos Kanoulas, Ben Carterette, and Nick Craswell. 2015. Overview of the TREC 2015 Tasks Track. In *TREC*.