Report on the WSDM 2018 Workshop on Learning from User Interactions

Rishabh Mehrotra
Spotify
London, UK
rishabhm@spotify.com

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

Emine Yilmaz
University College London
The Alan Turing Institute, London
emine.yilmaz@ucl.ac.uk

Abstract

The WSDM 2018 Workshop on Learning from User Interactions (Learn-IR'18)was a full day workshop which took place on February 9, 2018 at Los Angeles, CA. The workshop wasa highly interactive that provided a forum for academic and industrial researchers working at the intersection of user understanding, search tasks, conversational IR and user interactions. The workshop was meant 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. The workshop drew contributions from both industry and academia, and included six accepted research contributions. Additionally, the workshop had three keynote speakers from well established industry as well as academic researchers. The workshop witnessed lively discussion and interaction among participants, and brought together a wide variety of related research under the broader umbrella of learning from user interactions. This report outlines the events of the workshop and summarizes the major outcomes. More information about the workshop is available at https://task-ir.github.io/wsdm2018-learnIR-workshop/

1 Introduction

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 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 learnig from past interactions and incorporating implicit feedback signals into their models for rapid development.

2 Theme and Topics of Interest

The objective of the workshop was to provide a forum for academic and industrial researchers working at the intersection of user understanding, search tasks, conversational IR and user interactions. Topics of interest included but were not limited to:

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 Modeling; 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 Keynote Talks

The workshop witnessed three keynote talks by eminent researchers from industry and academia. The workshop started by a keynote talk by **Marc Najork** from Google Research, on *Learning from User Interactions in Personal Search*. Marc's talk focused on personal search scenario and elaborated techniques for training ranking functions on result clicks in an unbiased and scalable fashion, most of which are being used in various Google products, such as Gmail, Inbox, Drive and Calendar. While in web search, labels may either be assigned explicitly (through crowd-sourced assessors) or based on implicit user feedback (result clicks). In personal (e.g. email) search, obtaining labels is more difficult: document-query pairs cannot be given to assessors due to privacy constraints, and clicks on query-document pairs are extremely sparse since each user has a separate corpus. Marc's talk outlined in detail different approaches to tackle such challenges, and ignited a healthy discussion on the novel research needed to address personal and privacy aware ranking algorithms.

The second keynote was given by **Prof. Eugene Agichtein** who is an Associate Professor of Computer Science at Emory University. Eugene's talk on "*Improving Search By Learning from (Evolving) User Interactions*" focused on adapting user modeling techniques for the conversational setting, where users expect more from search while doing less. Voice-based assistants such as Alexa, Cortana and Siri, have rekindled the dream of a true conversation with a computer search engine. Hence, finding and presenting useful answers for the searchers' information needs is more important than ever. User interactions have been previously shown helpful for improving ranking, passage retrieval, and result summary generation. By viewing and interacting with the results and the underlying content, users implicitly and explicitly indicate quality, interestingness, and relevance of content to their information needs. Eugene's talk motivated the audience to consider fine grained user interaction signals to improve user's experience and resulted in a in-depth discussion on the different useful signals, and metrics like attention time and variants of dwell time.

The final keynote was given by **Alex Beutel** who is a Senior Research Scientist at Google Research. Alex's talk "User Dynamics and Context in Neural Recommender Systems" focused on some of his recent work on neural recommendation systems. Neural networks have become increasingly successful throughout many machine learning applications, and in the past few years, have become the state-of-the-art approach recommender systems. While DNNs have been well explored for applications in computer vision and natural language processing, their application to collaborative filtering-style recommendation creates new opportunities to understand user behavior and opens interesting questions on how to design these models. Alex's talk presented recent research on user and item dynamics modeled by RNNs, challenges in successfully capturing long-range dynamics, and how contextual information factors into modeling user behavior with neural networks.

4 Contributed Papers

The workshop invited research contributions in the form of full papers spanning 4-6 pages. All submissions were peer-reviewed by atleast 3 different reviewers and a total of six papers were accepted for inclusion in the workshop program. Each accepted paper was presented as an oral presentation. The papers covered a wide range of topics, including user representations, counterfactual learning to rank, user satisfaction and personal user experience. We next briefly describe the contributions.

- **Mixture-of-tastes Models for Representing Users with Diverse Interests [1]**: Presented by Maciej Kula, this paper highlighted the fact that most existing recommendation approaches implicitly treat user tastes as unimodal, resulting in an average-of-tastes representations when multiple distinct interests are present. The paper showed that appropriately modelling the multi-faceted nature of user tastes through a mixture-of-tastes model leads to large increases in recommendation quality.

- **Counterfactual Learning-to-Rank for Optimizing DCG [2]**: Contibuted by Aman Agarwal and Thorsten Joachims from Cornell University, this paper shows how the counterfactual framework can be generalized to directly optimize a broad class of information retrieval metrics. This class includes the Discounted Cumulative Gain (DCG) metric, which is considered to be a more user-centric metric than average rank. Based on the generalized counterfactual framework, the authors developed the PropDCG method for optimizing DCG and empirically find that PropDCG performs significantly better than PropRank in terms of DCG, and that PropDCG is robust to varying severity of presentation bias.

- **Understanding and Predicting User Satisfaction in Image Search [3]**: Contibuted by Zhijing Wu, Yiqun Liu, Min Zhang and Shaoping Ma; this paper studied the differences of user behavior under satisfied and unsatisfied occasions in image search. The paper presents three approaches to predict image search satisfaction using both hand crafted features and interaction sequences and sheds light on the importance of understanding user behavior to better predict user satisfaction on image search platforms.

- **Fully Automated QA System for Large Scale Search and Recommendation Engines Leveraging Implicit User Feedback [4]**: Contibuted by Khalifeh Aljadda, Mohammed Korayem and Trey Grainger, this paper introduces a fully-automated quality assurance (QA) system for search and recommendation engines that does not require participation of end users in the process of evaluating any changes in the existing relevancy algorithms. The proposed system has been used successfully on CareerBuilders web-scale search engine, where it has been demonstrated to accurately simulate nDCG scores offline for previously untested algorithms.

- **Personal User Experience: The Needs of the Elderly and Intelligent Interfaces Using Data Analysis Techniques [5]**: Contributed by Seungho Chae, Yoonsik Yang, Hyocheol Ro and Tack-Don Han; this paper studies personal UX for elderly people. Personal UX was applied to real-life elderly people using the pervasive assistive robot system, which is composed of a projection system capable of pan/tilt control. Elderly people have different degrees of discomfort due to aging, and we provided appropriate functions to users using data collected from real life.

- **Deriving Tourist Mobility Patterns from Check-in Data [6]**: Contributed by Linus W. Dietz, Daniel Herzog and Wolfgang Wrndl; this paper presented a data-mining approach to improve destination recommender systems with learned travel patterns. Specifically, a methodology to mine trips from location-based social networks was proposed to improve recommendations for the duration of stay at a destination. The authors demonstrate the utility of the approach using a Foursquare data set from which we extract 23,418 trips in 77 countries, and presented a number of interesting analysis which could guide the development of recommender systems for user assistance in planning a trip.

5 Conclusion and Future

The WSDM 2018 Workshop on Learning from User Interactions (Learn-IR) was successful in its goal of bringing together researchers to discuss the various ways in which machines could learn from user interactions. The first edition of the workshop witnessed participation from a diverse crowd interested in different aspects of user interactions and attracted a diverse set of contributions and discussions. WSDM was an excellent venue for the workshop owing to the rich and diverse mix of industrial and academic researchers and participants.

This workshop was a first and much needed forum to gather different research problems under the common umbrella of learning from user interactions. The need to have an active community and a common forum to discuss problems and research opportunities came to the forefront in the various discussions at the workshop. The workshop attendees were enthusiastic in their agreement about the need for future workshops on this topic. We also intend to propose and publicize a special issue of the Information Retrieval Journal on Learn-IR.

6 Acknowledgements

References

1. Maciej Kula. *Mixture-of-tastes Models for Representing Users with Diverse Interests*. In *Proceedings* of the *WSDM 2018 Workshop on Learning from User Interactions (Learn-IR)*, 2018.

2. Aman Agarwal and Thorsten Joachims. *Counterfactual Learning-to-Rank for Optimizing DCG*. In *Proceedings* of the *WSDM 2018 Workshop on Learning from User Interactions (Learn-IR)*, 2018.

3. Zhijing Wu, Yiqun Liu, Min Zhang and Shaoping Ma. *Understanding and Predicting User Satisfaction in Image Search*. In *Proceedings* of the *WSDM 2018 Workshop on Learning from User Interactions (Learn-IR)*, 2018.

4. Khalifeh Aljadda, Mohammed Korayem and Trey Grainger. *Fully Automated QA System for Large Scale Search and Recommendation Engines Leveraging Implicit User Feedback*. In *Proceedings* of the *WSDM 2018 Workshop on Learning from User Interactions (Learn-IR)*, 2018.

5 Seungho Chae, Yoonsik Yang, Hyocheol Ro and Tack-Don Han. *Personal User Experience: The Needs of the Elderly and Intelligent Interfaces Using Data Analysis Techniques.* In *Proceedings* of the *WSDM 2018 Workshop on Learning from User Interactions (Learn-IR)*, 2018.

6 Linus W. Dietz, Daniel Herzog and Wolfgang Wrndl. *Deriving Tourist Mobility Patterns from Check-in Data.* In *Proceedings* of the *WSDM 2018 Workshop on Learning from User Interactions (Learn-IR)*, 2018.