



Special issue on learning from user interactions

Rishabh Mehrotra¹ · Ahmed Hassan Awadallah² · Emine Yilmaz³

Received: 7 October 2020 / Accepted: 8 October 2020
© Springer Nature B.V. 2020

When users interact with online services (e.g. search engines, recommender systems, conversational agents), they leave behind traces of interaction patterns. The ability to record and interpret user interaction signals (Guo and Agichtein 2012) and understand user behavior (Mehrotra et al. 2016a) 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 such as improving user engagement, incorporating user feedback and gauging user satisfaction.

Understanding and learning from user interactions involves a number of different aspects (Mehrotra et al. 2018; Mehrotra 2018)—from understanding user intent and tasks (Mehrotra et al. 2016b, c; Mehrotra and Yilmaz 2017a; Santy et al. 2019; White et al. 2015), to developing user models and personalization services (Mehrotra and Yilmaz 2015, 2017; Liu et al. 2019). Beyond understanding user needs, learning from user interactions involves developing the right metrics for evaluation and experimentation systems (Mehrotra et al. 2017b, c, Verma et al. 2016), understanding user interaction processes (Liu et al. 2014), their usage context (Mehrotra and Yilmaz 2017b) and designing interfaces capable of helping users (Hassan Awadallah et al. 2014). As such, understanding user behavior could allow the system to support users at the various stages of their tasks. This could have implications on many aspects of the system design including user interface, understanding user intents in search and recommendation scenarios (Mehrotra et al. 2019), presentation of information, retrieving and ranking (Bendersky et al. 2017; Glowacka et al. 2013), unbiased learning (Jagerman et al. 2019), etc.

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 Kiseleva et al. (2016), Mehrotra et al. (2017a). Chatbots, personal assistants on mobile phones, smart speakers and other 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

✉ Rishabh Mehrotra
rishabhm@spotify.com

¹ Spotify Research, London, UK

² Microsoft Research, Redmond, USA

³ University College London, London, UK

users are increasingly relying on such digital assistants to fulfil 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.

1 Overview of articles in this special issue

Maistro et al. focus on boosting learning to rank with user dynamics and continuation methods. Specifically, this research explores ways of modeling user dynamics into learning to rank models, and integrating an explicit knowledge of the underlying user-interaction model in the training process. The possibility of targeting different objective functions is explored to effectively drive the algorithm towards promising areas of the search space. By considering complex query-based user dynamics instead of simply discounting the gain by the rank position, and by designing a learning path across different loss functions that can capture different signals in the training data, the learning process is enriched using user interaction information.

Gao et al. propose to leverage user interaction information for interactive multi-document summarisation. This research studies preference-based interactive summarisation, so as to reduce the number of interaction rounds, and proposes an active preference based reinforcement learning framework to learn a summary ranking function from the preferences, and neural reinforcement learning to efficiently search for the (near-)optimal summary. Findings from this research suggests that users can easily provide reliable preferences over summaries and therefore, provides grounds for developing NLP models which learn from user interactions.

2 Conclusion

As user-centric search and recommender systems evolve, and newer forms of platforms emerge which help users in their information seeking journeys, learning from user interactions will increasingly become a topic of greater importance. This special issue focuses on developing techniques in ranking and natural language processing which can benefit from learning from user interactions. We hope the selected papers and this special issue provides the community with a better understanding of the current directions, and that they inspire readers with possible areas to focus on in their future research.

Acknowledgements We thank all authors for submitting manuscripts to this special issue and choosing it as an outlet for their research in the area of learning from user interactions. We are also thankful to all the reviewers for their efforts in providing useful feedback to the authors. Our gratitude further extends to the editors-in-chief, for their kind support during the preparation and completion of this special issue.

References

- Bendersky, M., Wang, X., Metzler, D., & Najork, M. (2017). Learning from user interactions in personal search via attribute parameterization. In *Proceedings of the tenth ACM international conference on web search and data mining* (pp. 791–799).
- Glowacka, D., Ruotsalo, T., Konuyshkova, K., Athukorala, K., Kaski, S., & Jacucci, G. (2013). Directing exploratory search: Reinforcement learning from user interactions with keywords. In *Proceedings of the 2013 international conference on Intelligent user interfaces* (pp. 117–128).

- Guo, Q., & Agichtein, E. (2012). Beyond dwell time: Estimating document relevance from cursor movements and other post-click searcher behavior. In *Proceedings of the 21st international conference on World Wide Web* (pp. 569–578).
- Hassan Awadallah, A., White, R. W., Pantel, P., Dumais, S. T., & Wang, Y. M. (2014). Supporting complex search tasks. In *Proceedings of the 23rd ACM international conference on conference on information and knowledge management* (pp. 829–838).
- Jagerman, R., Oosterhuis, H., & de Rijke, M. (2019). To model or to intervene: A comparison of counterfactual and online learning to rank from user interactions. In *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval* (pp. 15–24).
- Kiseleva, J., Williams, K., Hassan Awadallah, A., Crook, A. C., Zitouni, I., & Anastasakos, T. (2016). Predicting user satisfaction with intelligent assistants. In *Proceedings of the 39th international ACM SIGIR conference on research and development in information retrieval* (pp. 45–54).
- Liu, C., Liu, J., & Belkin, N. J. (2014). Predicting search task difficulty at different search stages. In *Proceedings of the 23rd ACM international conference on conference on information and knowledge management* (pp. 569–578).
- Liu, J., Mitsui, M., Belkin, N. J., & Shah, C. (2019). Task, information seeking intentions, and user behavior: Toward a multi-level understanding of Web search. In *Proceedings of the 2019 conference on human information interaction and retrieval* (pp. 123–132).
- Mehrotra, R. (2018). *Inferring user needs and tasks from user interactions*. Doctoral dissertation, UCL (University College London).
- Mehrotra, R., Awadallah, A. H., Kholy, A. E., & Zitouni, I. (2017a). Hey Cortana! Exploring the use cases of a desktop based digital assistant. In *SIGIR 1st international workshop on conversational approaches to information retrieval (CAIR'17)* (Vol. 4).
- Mehrotra, R., Awadallah, A. H., Shokouhi, M., Yilmaz, E., Zitouni, I., El Kholy, A., & Khabsa, M. (2017b). Deep sequential models for task satisfaction prediction. In *Proceedings of the 2017 ACM on conference on information and knowledge management* (pp. 737–746).
- Mehrotra, R., Awadallah, A. H., & Yilmaz, E. (2018). Learnir: Wsdm 2018 workshop on learning from user interactions. In *Proceedings of the eleventh ACM international conference on web search and data mining* (pp. 797–798).
- Mehrotra, R., Bhattacharya, P., & Yilmaz, E. (2016a). Uncovering task based behavioral heterogeneities in online search behavior. In *Proceedings of the 39th international ACM SIGIR conference on research and development in information retrieval* (pp. 1049–1052).
- Mehrotra, R., Bhattacharya, P., & Yilmaz, E. (2016b). Deconstructing complex search tasks: A bayesian non-parametric approach for extracting sub-tasks. In *Proceedings of the 2016 conference of the north american chapter of the association for computational linguistics: Human language technologies* (pp. 599–605).
- Mehrotra, R., Bhattacharya, P., & Yilmaz, E. (2016c). Characterizing users' multi-tasking behavior in web search. In *Proceedings of the 2016 ACM on conference on human information interaction and retrieval* (pp. 297–300).
- Mehrotra, R., Lalmas, M., Kenney, D., Lim-Meng, T., & Hashemian, G. (2019). Jointly leveraging intent and interaction signals to predict user satisfaction with slate recommendations. In *The World Wide Web conference* (pp. 1256–1267).
- Mehrotra, R., & Yilmaz, E. (2015). Terms, topics & tasks: Enhanced user modelling for better personalization. In *Proceedings of the 2015 international conference on the theory of information retrieval* (pp. 131–140).
- Mehrotra, R., & Yilmaz, E. (2017a). Extracting hierarchies of search tasks & subtasks via a bayesian nonparametric approach. In *Proceedings of the 40th international ACM SIGIR conference on research and development in information retrieval* (pp. 285–294).
- Mehrotra, R., & Yilmaz, E. (2017b). Task embeddings: Learning query embeddings using task context. In *Proceedings of the 2017 ACM on conference on information and knowledge management* (pp. 2199–2202).
- Mehrotra, R., Zitouni, I., Hassan Awadallah, A., Kholy, A. E., & Khabsa, M. (2017c). User interaction sequences for search satisfaction prediction. In *Proceedings of the 40th International ACM SIGIR conference on research and development in information retrieval* (pp. 165–174).
- Santy, S., Zulfikar, W., Mehrotra, R., & Yilmaz, E. (2019). Towards task understanding in visual settings. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 33, pp. 10027–10028).
- Verma, M., Yilmaz, E., Mehrotra, R., Kanoulas, E., Carterette, B., Craswell, N., & Bailey, P. (2016). Overview of the TREC tasks track 2016. In *TREC*.
- White, R. W., Richardson, M., & Yih, W. T. (2015). Questions vs. queries in informational search tasks. In *Proceedings of the 24th international conference on World Wide Web* (pp. 135–136).