Living Labs for Online Evaluation: From Theory to Practice

Anne Schuth[1] and Krisztian Balog[2]

[1] University of Amsterdam, The Netherlands
[2] University of Stavanger, Stavanger, Norway
anne.schuth@uva.nl, krisztian.balog@uis.no

Abstract. Experimental evaluation has always been central to Information Retrieval research. The field is increasingly moving towards *online* evaluation, which involves experimenting with real, unsuspecting users in their natural task environments, a so-called *living lab*. Specifically, with the recent introduction of the Living Labs for IR Evaluation initiative at CLEF and the OpenSearch track at TREC, researchers can now have direct access to such labs. With these benchmarking platforms in place, we believe that online evaluation will be an exciting area to work on in the future. This half-day tutorial aims to provide a comprehensive overview of the underlying theory and complement it with practical guidance.

1 Motivation and Overview

Experimental evaluation has always been a key component in Information Retrieval research. Most commonly, systems are evaluated following the Cranfield methodology [22,4]. Using this approach, systems are evaluated in terms of document relevance for given queries, which is assessed by trained experts. While the Cranfield methodology ensures high internal validity and repeatability of experiments, it has been shown that the users' search success and satisfaction with an IR system are not always accurately reflected by standard IR metrics [31,29]. One reason is that the relevance judges typically do not assess queries and documents that reflect their own information needs, and have to make assumptions about relevance from an assumed user's point of view. Because the true information need can be difficult to assess, this can cause substantial biases [34,11,30]. To address these shortcomings, the field is increasingly moving towards *online* evaluation, which involves experimenting with real, unsuspecting users in their natural task environments. Essentially, the production search engine operates as a "living lab." For a long time, this type of evaluation was only available to those working within organizations that operate a search engine. But this is about to change. For one thing, the need to involve real users is know openly and widely acknowledged in our community (as witnessed, e.g., by the panel discussion at ECIR'15 and the Salton Award keynote lecture of Nicholas J. Belkin at SIGIR'15 [2]). For another thing, pioneering efforts to realize the idea of living labs in practice are now in place and are available to the community. Specifically

the Living Labs for IR Evaluation (LL4IR)[3] initiative runs as a benchmarking campaign at CLEF, but also operates monthly challenges so that people do not have to wait for a yearly evaluation cycle. The most recent initiative is the OpenSearch track at TREC[4], which focuses on *academic literature search*.

Understanding the differences between online and offline evaluation is still a largely unexplored area of research. There is a lot of fundamental research to happen in this space that has not happened yet because of the lack availability of experimental resources to the academic community. With recent developments, we believe that online evaluation will be an exciting area to work on in the future. The motivation for this tutorial is twofold: (1) to raise awareness and promote this form of evaluation (i.e., online evaluation with living labs) in the community, and (2) to help people get started by working through all the steps of the development and deployment process, using the LL4IR evaluation platform.

This half-day tutorial aims to provide a comprehensive overview of the underlying theory and complement it with practical guidance. The tutorial is organized in two 1,5 hours sessions with a break in between. Each session interleaves theoretical, practical, and interactive elements to keep the audience engaged. For the practical parts, we break with the traditional format by using hands-on instructional techniques. We will make use of an online tool, called DataJoy,[5] that proved invaluable in our previous classroom experience. This allows participants to (1) run Python code in a browser window without having to install anything locally, (2) follow the presenter's screen on their own laptop and, (3) at the same time, have their own private copy of the project on a different browser tab.

2 Target Audience and Learning Objectives

The primary target audience are graduate students and lecturers/professors teaching IR classes. Engineers from companies operating search engines might also find the tutorial useful. Our learning objectives include the following topics.

We will start our tutorial with an extensive overview of online evaluation methods. We begin with *A/B Testing* [16], which compares two systems by showing system A to one group of users and system B to another group. A/B testing then tries to infer a difference between the systems from differences in observed behavior. We describe many ways of measuring observed behavior: (1) click through rate (CTR) [14]; (2) dwell time [34]; (3) satisfied clicks [15]; (4) tabbed browsing [13]; (5) abandonment [18,28]; (6) query reformulation [8]; (7) skips [32]; (8) mouse movement [6,33,7,10,5]; and (9) in-view time [17].

While providing flexibility and control, A/B comparisons typically require a large number of observations. *Interleaved comparison methods* reduce the variance of measurement by presenting users with a result list that combines the rankings of systems A and B. We provide a comprehensive overview of the following interleaving methods: (1) balanced interleave (BI) [14]; (2) team draft

[3] http://living-labs.net
[4] http://trec-open-search.org/
[5] http://getdatajoy.com

interleave (TDI) [21]; (3) document constraints (DC) [9]; (4) probabilistic interleave (PI) [12]; (5) optimized interleave (OI) [20]; (6) team draft multileave (TDM) [27]; and (7) probabilistic multileave (PM) [24].

Next, we discuss a comparison of interleaving and A/B metrics [25]. We then turn to simulating user interactions [26] using *click models* [3]. Finally, we touch on *learning to rank* in two variants: *offline* learning to rank [19] and *online* learning to rank [35], of which the latter requires the aforementioned evaluation methods.

Having provided the necessary theoretical background, we introduce the *living labs for IR* (LL4IR) [1] evaluation platform in depth. We will focus on two specific use-cases [23] from the CLEF lab: product search and web search. During the practical sessions, participants will gain hands-on experience with the LL4IR platform [1], which includes: (1) registering and obtaining an API key; (2) getting queries and candidate items; (3) generating and uploading a ranking; and (4) obtaining feedback and outcomes. The API documentation and course material are available at `http://living-labs.net`.

References

1. Krisztian Balog, Liadh Kelly, and Anne Schuth. Head First: Living Labs for Ad-hoc Search Evaluation. In *CIKM '14*, pages 1815–1818, New York, New York, USA, nov 2014. ACM Press.
2. Nicholas J. Belkin. Salton award lecture: People, interacting with information. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '15, pages 1–2. ACM, 2015.
3. Aleksandr Chuklin, Ilya Markov, and Maarten de Rijke. *Click Models for Web Search*. Synthesis Lectures on Information Concepts, Retrieval, and Services. Morgan & Claypool Publishers, August 2015.
4. CW Cleverdon and M Keen. Aslib Cranfield research project-Factors determining the performance of indexing systems; Volume 2, Test results. 1966.
5. Fernando Diaz, Ryen White, Georg Buscher, and Dan Liebling. Robust models of mouse movement on dynamic web search results pages. In *CIKM*, pages 1451–1460. ACM Press, October 2013.
6. Qi Guo and Eugene Agichtein. Understanding "Abandoned" Ads: Towards Personalized Commercial Intent Inference via Mouse Movement Analysis. *SIGIR-IRA*, 2008.
7. Qi Guo and Eugene Agichtein. Towards predicting web searcher gaze position from mouse movements. In *CHI EA*, page 3601, April 2010.
8. Ahmed Hassan, Xiaolin Shi, Nick Craswell, and Bill Ramsey. Beyond clicks: query reformulation as a predictor of search satisfaction. In *CIKM*, 2013.
9. Jing He, Chengxiang Zhai, and Xiaoming Li. Evaluation of methods for relative comparison of retrieval systems based on clickthroughs. In *CIKM '09*. ACM, 2009.
10. Yin He and Kuansan Wang. Inferring search behaviors using partially observable markov model with duration (POMD). In *WSDM*, 2011.
11. William Hersh, Andrew H. Turpin, Susan Price, Benjamin Chan, Dale Kramer, Lynetta Sacherek, and Daniel Olson. Do batch and user evaluations give the same results? In *SIGIR*, pages 17–24, 2000.

12. Katja Hofmann, Shimon Whiteson, and Maarten de Rijke. A probabilistic method for inferring preferences from clicks. In *CIKM '11*. ACM, 2011.
13. Jeff Huang, Thomas Lin, and Ryen White. No search result left behind. In *WSDM*, page 203, 2012.
14. Thorsten Joachims, Laura A. Granka, Bing Pan, Helene Hembrooke, Filip Radlinski, and Geri Gay. Evaluating the accuracy of implicit feedback from clicks and query reformulations in Web search. *ACM Trans. Inf. Syst.*, 25(2), 2007.
15. Y Kim, Ahmed Hassan, Ryen White, and I Zitouni. Modeling dwell time to predict click-level satisfaction. In *WSDM*, 2014.
16. Ronny Kohavi. Online Controlled Experiments, 2013.
17. Dmitry Lagun, CH Hsieh, D Webster, and V Navalpakkam. Towards Better Measurement of Attention and Satisfaction in Mobile Search. In *SIGIR*, 2014.
18. Jane Li, Scott Huffman, and Akihito Tokuda. Good abandonment in mobile and pc internet search. In *SIGIR '09*, pages 43–50, 2009.
19. Tie-Yan Liu. *Learning to rank for information retrieval*, volume 3. Now Publishers, 2009.
20. Filip Radlinski and Nick Craswell. Optimized interleaving for online retrieval evaluation. In *WSDM '13*. ACM, 2013.
21. Filip Radlinski, Madhu Kurup, and Thorsten Joachims. How does clickthrough data reflect retrieval quality? In *CIKM '08*. ACM, 2008.
22. Mark Sanderson. Test Collection Based Evaluation of Information Retrieval Systems. *Foundations and Trends in Information Retrieval*, 4(4):247–375, 2010.
23. Anne Schuth, Krisztian Balog, and Liadh Kelly. Overview of the living labs for information retrieval evaluation (ll4ir) clef lab 2015. In *CLEF 2015*. Springer, 2015.
24. Anne Schuth, Robert-Jan Bruintjes, Fritjof Büttner, Joost van Doorn, Carla Groenland, Harrie Oosterhuis, Cong-Nguyen Tran, Bas Veeling, Jos van der Velde, Roger Wechsler, David Woudenberg, and Maarten de Rijke. Probabilistic multileave for online retrieval evaluation. In *Proceedings of SIGIR*, 2015.
25. Anne Schuth, Katja Hofmann, and F Radlinski. Predicting Search Satisfaction Metrics with Interleaved Comparisons. In *SIGIR '15*, 2015.
26. Anne Schuth, Katja Hofmann, Shimon Whiteson, and Maarten de Rijke. Lerot: an Online Learning to Rank Framework. In *LivingLab '13*, pages 23–26. ACM Press, nov 2013.
27. Anne Schuth, Floor Sietsma, Shimon Whiteson, Damien Lefortier, and Maarten de Rijke. Multileaved comparisons for fast online evaluation. In *CIKM'14*, 2014.
28. Y Song, Xiaolin Shi, Ryen White, and Ahmed Hassan. Context-Aware Web Search Abandonment Prediction. In *SIGIR*, 2014.
29. Jamie Teevan, Susan Dumais, and Eric Horvitz. The potential value of personalizing search. In *SIGIR*, pages 756–757, 2007.
30. Andrew Turpin and William Hersh. Why batch and user evaluations do not give the same results. In *SIGIR*, pages 225–231, 2001.
31. Andrew Turpin and Frank Scholar. User performance versus precision measures for simple search tasks. In *SIGIR*, pages 11–18, 2006.
32. K Wang, T Walker, and Z Zheng. PSkip: Estimating relevance ranking quality from web search clickthrough data. In *KDD*, pages 1355–1364, 2009.
33. Kuansan Wang, Nikolas Gloy, and Xiaolong Li. Inferring search behaviors using partially observable Markov (POM) model. In *WSDM*, 2010.
34. Emine Yilmaz, Manisha Verma, Nick Craswell, Filip Radlinski, and Peter Bailey. Relevance and Effort: An Analysis of Document Utility. In *CIKM*, 2014.
35. Yisong Yue and Thorsten Joachims. Interactively optimizing information retrieval systems as a dueling bandits problem. In *ICML '09*, pages 1201–1208, 2009.