

Overview of the Special Issue on Contextual Search and Recommendation

PAUL N. BENNETT, Microsoft Research
KEVYN COLLINS-THOMPSON, University of Michigan
DIANE KELLY, University of North Carolina at Chapel Hill
RYEN W. WHITE, Microsoft Research
YI ZHANG, University of California Santa Cruz

Information systems that leverage contextual knowledge about their users and their search situations – such as histories, demographics, surroundings, constraints or devices – can provide tailored search experiences and higher-quality task outcomes. Within information retrieval, there is a growing focus on how knowledge of user interests, intentions, and context can improve aspects of search and recommendation such as ranking and query suggestion, especially for exploratory and/or complex tasks that can span multiple queries or search sessions. The interactions that occur during these complex tasks provide context that can be leveraged by search systems to support users' broader information-seeking activities. Next-generation recommender systems face analogous challenges, including integrating signals from user exploration to update recommendations in real time.

Within the space of search, much of the work on modeling context and search personalization has focused on constructing topical profiles of the user's short- and long-term search history [Gauch et al. 2004; Chirita et al. 2005; Speretta and Gauch 2005; Ma et al. 2007; Bennett et al. 2010; White et al. 2010; Xiang et al. 2010; Sontag et al. 2012] or more generally, models of their query and result-click sequences [Cao et al. 2008; Cao et al. 2009; Mihalkova and Mooney 2009]. Related research has also considered a more content-driven representation such as language-model based approaches [Tan et al. 2006] or weighted term vectors derived from long-term desktop search activities [Teevan et al. 2005; Matthijs and Radlinski 2011]. However, a variety of recent investigations to contextualize search include a broader set of factors based on: a user's location [Bennett et al. 2011], a user's task-based search activity [Jones and Klinkner 2008; Kanoulas et al. 2011b; 2011a; Kanoulas et al. 2012; Sontag et al. 2012; Melucci 2012; Raman et al. 2014], the long-term vs. short-term interests of the user [Sugiyama et al. 2004; Li et al. 2007; Bennett et al. 2012], the ability of users to consume information at differing levels of complexity [Collins-Thompson et al. 2011], and patterns of re-finding the same search result over time [Teevan et al. 2011; Shokouhi et al. 2013].

The growth in the types of context explored and the information available to search systems derives from the timely convergence of several factors. The rapid growth in the use of different devices – most notably smartphones and tablets, but also including stationary devices such as game consoles, smart televisions, and augmented conference rooms – provides opportunities to obtain both raw and derived contextual signals that could power next-generation search and recommendation systems. The use of such signals in search and recommendation tasks has been recently explored in such venues as the Context-awareness in Retrieval and Recommendation workshops at IUI 2011-2012 [Luca et al. 2011; Luca et al. 2012], WSDM 2013 [Böhmer et al. 2013], and ECIR 2014 [Said et al. 2014].

Furthermore, a variety of recent work and venues have noted that much information retrieval research on web search has focused on optimizing and evaluating single queries, even though a significant fraction of queries are associated with more complex tasks [Jones and Klinkner 2008; Kanoulas et al. 2011b; 2011a; Belkin et al. 2012a;

Kanoulas et al. 2012; Sontag et al. 2012; Bailey et al. 2012; Dean-Hall et al. 2013; Raman et al. 2014], which span one or more search sessions [Liu and Belkin 2010; Kotov et al. 2011], and often consist of multiple queries. Such complex tasks have additionally been the focus of a recent National Institute of Informatics sponsored workshop (Whole-Session Evaluation of Interactive Information Retrieval Systems [Belkin et al. 2012b]) as well as a recent National Science Foundation sponsored workshop (Task-Based Information Search [Kelly et al. 2013]). The interactions that take place during these complex tasks provide context that can be leveraged not simply in a query-myopic manner but in a way that supports users' broader information-seeking activities [Morris et al. 2008; Agichtein et al. 2012].

At the same time as e-commerce and mobile device usage have expanded, recommender systems have become increasingly important and received much research and commercial interest [Resnick and Varian 1997; Herlocker et al. 2004; Adomavicius and Tuzhilin 2011]. Recommender systems gather information from a given user, create and update the user's profile using implicit and explicit feedback, predict which items the user might prefer and, without requiring explicit user queries, recommend items tailored to the user's personal tastes. Recently, people have recognized the potential value of contextual information for recommendation in e-commerce, travel, and mobile applications. There has been an increasing research interest on contextual recommendation and different contextual information such as user intent, time, companion, weather, location, objects around, season, and temperature have all been studied in various recommender system settings [Palmisano et al. 2008; Karatzoglou et al. 2010; Li et al. 2010; Rendle et al. 2011; Baltrunas et al. 2012]. Like contextual search, the next generation of recommender systems face many of the same challenges of incorporating heterogeneous contexts into recommendations as well as an analogous challenge of incorporating the interactive process of the users exploration in a single session to contextually update recommendations on the fly.

We solicited articles for this special issue that describe the state-of-the-art and emerging trends in contextual search and recommendation. While manuscripts focused on all areas of contextual search and recommendation were considered, we especially encouraged submissions that targeted exploratory and/or complex tasks. In particular, representations and approaches to context that enable task-oriented search including tasks that persist longitudinally. From 20 submissions, we selected four high-quality articles that represent current themes of research on contextual search and recommendation. We now briefly summarize each of these accepted articles in turn.

One critical aspect of context for search activity is the overall task goal that a user has in seeking information [Ingwersen and Järvelin 2005]. In their article, *Task-Based Information Interaction Evaluation: The Viewpoint Of Program Theory*, Järvelin and colleagues target the task-based evaluation of information systems, user behaviors, and their interactions. They define task-based information interaction (TBII), and examine a number of interrelated behavioral and cognitive activities related to task planning, searching for information items, deciding between items, working with items, and synthesizing and reporting content and outcomes. To ground their research on TBII, the authors focus on learning tasks [Marchionini 2006], i.e., those that include an intentional or unintentional learning component. They develop an evaluation framework based on program theory (PT) [Rossi et al. 2004]. PT is useful since it supports the comprehensive evaluation of complex programs, which have several interacting factors affecting their performance. This is clearly the case in information seeking scenarios, especially during multi-faceted complex search tasks. PT facilitates an enhanced understanding of whether a program comprising TBII activities and tools will be effective, the salient processes involved, and importantly, the causal factors that determine search effectiveness. Järvelin and colleagues describe how the evaluation

of TBII could be accomplished using PT. Although they do not offer fully developed or verified program theories for the activities or the entire task process, their research charts a course toward a more complete evaluation of information interactions and their supporting tools. The application of PT to TBII places an increased focus on human factors within IR evaluation. It also enables a holistic assessment of the effectiveness of search systems designed to help people tackle complex search tasks, such as those with an emphasis on learning.

Continuing the focus on search task as context, the article *User Activity Patterns During Information Search* by Cole and colleagues explores the connection between a user's interaction behavior and the nature and difficulty of their task goal. The authors study information-finding tasks in two different domains – journalism and genomic research – by capturing interaction signals that reflect both high-level activity in the form of page interaction sequences from log files, and low-level reading activity based on eye movement patterns. The authors contribute techniques for representing a user's activity patterns during extended search sessions, and present results showing that these techniques can be used to accurately differentiate task types using complexity-based measures of activity patterns. Applications of this research include improved contextual retrieval that can adapt ranking algorithms or other aspects of the search interface to better match the properties of the inferred task goal.

The two articles described thus far adopted an individual task-centered view of context. In *Profile-Based Summarisation For Web Site Navigation*, Alhindi and colleagues approach the topic of contextual search from the perspective of personalizing search to a community of users, rather than to an individual. They present a series of laboratory user studies to evaluate the effectiveness of different approaches to creating profile-based summaries in scenarios where users navigate a common information space, in their case, a university website. These summaries are created using the accumulated search histories of those who have used the website in the past. The basic idea is that members of the community have information needs that distinguish them from non-members, and furthermore, that many of these needs occur with some regularity. The community profile is acquired by using an ant colony optimization analogy, which takes temporal context into consideration. This allows the model to associate concepts with certain times of the academic year and anticipate the occurrence of certain needs. The community model is then used to generate customized summaries of webpages that appear as community members mouse-over hyperlinks. This work represents a unique perspective on contextual search by focusing on search support that is tailored to groups, or communities of users, and their information needs. In this way, it is similar to what Teevan et al. [2009] refer to as 'groupization,' or using group information to personalize retrieval and also extends past work by Smyth, Freyne and colleagues who explored search and recommendation techniques tailored to communities of users [Freyne et al. 2007; Smyth et al. 2005]. This work also exemplifies the successful translation of information retrieval research into local practice; Alhindi and colleagues use their own university website logs in this research and design techniques specifically to help their university community. Others might use these ideas and techniques to tailor search to their own local communities.

Finally, when people access content via smart phones, tablets and wearable devices, contextual information such as location, time, and recorded activities are important aspects that can be used to improve search and recommendation performance. In their article, *Who, Where, When and What: a Non-parametric Bayesian Approach to Context-aware Recommendation and Search for Twitter Users*, Yuan and colleagues consider how to leverage such types of context. In particular, they focus on four context variables: user (who), spatial (Where), temporal (When), and activity (What). They propose a new topic model to jointly model the four components for information in micro-blogs.

In their experiments with two data sets, the authors show that the proposed model can discover users' spatial-temporal topics, predict the location for a particular tweet, and make location recommendations given a user requirement. Their work serves as an example of how to integrate complex contextual information in social media and how to develop new models for a particular contextual recommendation problem based on good intuitions. In addition, this article also represents a unique perspective on how to evaluate context based recommendations – a very challenging problem due to the lack of good evaluation data sets or methodologies. We hope the discussion in this article prompts future research in this direction. In particular, while some researchers might follow the authors' ideas of using an existing data set and a reasonable amount of manual effort to build a data set tailored to their own contextual recommendation problem, others might borrow ideas from search engine evaluations or propose new evaluation methodologies.

Thus, the recent growth in work on complex task-oriented search and recommendation combined with the interest in context derived from mobile and situated devices – as well as across devices – make this an opportune time for a special issue in this area. Given the timeliness and breadth of the topic, and the level of interest in events such as related workshops, we believe that readers will find these articles both informative and thought-provoking.

Guest Editors

Paul N. Bennett,
Kevyn Collins-Thompson,
Diane Kelly,
Ryen W. White,
Yi Zhang

REFERENCES

- Gediminas Adomavicius and Alexander Tuzhilin. 2011. Context-Aware Recommender Systems. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul B. Kantor (Eds.). Springer US, 217–253. DOI: http://dx.doi.org/10.1007/978-0-387-85820-3_7
- Eugene Agichtein, Ryen W. White, Susan T. Dumais, and Paul N. Bennett. 2012. Search, interrupted: understanding and predicting search task continuation. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval (SIGIR '12)*. ACM, New York, NY, USA, 315–324. DOI: <http://dx.doi.org/10.1145/2348283.2348328>
- Peter Bailey, Liwei Chen, Scott Groesenick, Li Jiang, Yan Li, Paul Reinholdtsen, Charles Salada, Haidong Wang, and Sandy Wong. 2012. User task understanding: a web search engine perspective. (2012). <http://research.microsoft.com/apps/pubs/default.aspx?id=180594>.
- Linas Baltrunas, Bernd Ludwig, Stefan Peer, and Francesco Ricci. 2012. Context Relevance Assessment and Exploitation in Mobile Recommender Systems. *Personal Ubiquitous Comput.* 16, 5 (June 2012), 507–526. DOI: <http://dx.doi.org/10.1007/s00779-011-0417-x>
- Nicholas J. Belkin, Charles L.A. Clarke, Ning Gao, Jaap Kamps, and Jussi Karlgren. 2012a. Report on the SIGIR Workshop on "Entertain Me": Supporting Complex Search Tasks. *SIGIR Forum* 45, 2 (Jan. 2012), 51–59. DOI: <http://dx.doi.org/10.1145/2093346.2093354>
- Nicholas J. Belkin, Susan Dumais, Noriko Kando, and Mark Sanderson. 2012b. Whole-Session Evaluation of Interactive Information Retrieval Systems. <http://www.nii.ac.jp/shonan/seminar020/>. (October 2012).
- Paul N. Bennett, Filip Radlinski, Ryen W. White, and Emine Yilmaz. 2011. Inferring and Using Location Metadata to Personalize Web Search. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '11)*. ACM, New York, NY, USA, 135–144. DOI: <http://dx.doi.org/10.1145/2009916.2009938>

- Paul N. Bennett, Krysta Svore, and Susan T. Dumais. 2010. Classification-enhanced ranking. In *Proceedings of the 19th international conference on World wide web (WWW '10)*. ACM, New York, NY, USA, 111–120. DOI: <http://dx.doi.org/10.1145/1772690.1772703>
- Paul N. Bennett, Ryen W. White, Wei Chu, Susan T. Dumais, Peter Bailey, Fedor Borisjuk, and Xiaoyuan Cui. 2012. Modeling the Impact of Short- and Long-term Behavior on Search Personalization. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '12)*. ACM, New York, NY, USA, 185–194. DOI: <http://dx.doi.org/10.1145/2348283.2348312>
- Matthias Böhmer, Ernesto W. De Luca, Alan Said, and Jaime Teevan (Eds.). 2013. *CaRR '13: Proceedings of the 3rd Workshop on Context-awareness in Retrieval and Recommendation*. ACM, New York, NY, USA.
- Huanhuan Cao, Derek Hao Hu, Dou Shen, Daxin Jiang, Jian-Tao Sun, Enhong Chen, and Qiang Yang. 2009. Context-aware Query Classification. In *Proceedings of the 32Nd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '09)*. ACM, New York, NY, USA, 3–10. DOI: <http://dx.doi.org/10.1145/1571941.1571945>
- Huanhuan Cao, Daxin Jiang, Jian Pei, Qi He, Zhen Liao, Enhong Chen, and Hang Li. 2008. Context-aware query suggestion by mining click-through and session data. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '08)*. ACM, New York, NY, USA, 875–883. DOI: <http://dx.doi.org/10.1145/1401890.1401995>
- Paul Alexandru Chirita, Wolfgang Nejdl, Raluca Paiu, and Christian Kohlschütter. 2005. Using ODP Metadata to Personalize Search. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '05)*. ACM, New York, NY, USA, 178–185. DOI: <http://dx.doi.org/10.1145/1076034.1076067>
- Kevyn Collins-Thompson, Paul N. Bennett, Ryen W. White, Sebastian de la Chica, and David Sontag. 2011. Personalizing Web Search Results by Reading Level. In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management (CIKM '11)*. ACM, New York, NY, USA, 403–412. DOI: <http://dx.doi.org/10.1145/2063576.2063639>
- Adriel Dean-Hall, Charles L.A. Clarke, Jaap Kamps and Paul Thomas, and Ellen Vorhees. 2013. Overview of the TREC 2012 Contextual Suggestion Track. In *Proceedings of the Twenty-First Text Retrieval Conference (TREC '12)*.
- Jill Freyne, Rosta Farzan, Peter Brusilovsky, Barry Smyth, and Maurice Coyle. 2007. Collecting Community Wisdom: Integrating Social Search & Social Navigation. In *Proceedings of the 12th International Conference on Intelligent User Interfaces (IUI '07)*. ACM, New York, NY, USA, 52–61. DOI: <http://dx.doi.org/10.1145/1216295.1216312>
- Susan Gauch, Jason Chaffee, and Alexander Pretschner. 2004. Ontology-Based User Profiles for Search and Browsing. *Web Intelligence and Agent Systems* 1, 3-4 (2004), 219–234.
- Jonathan L. Herlocker, Joseph A. Konstan, Loren G. Terveen, and John T. Riedl. 2004. Evaluating Collaborative Filtering Recommender Systems. *ACM Transactions on Information Systems (TOIS)* 22, 1 (Jan. 2004), 5–53. DOI: <http://dx.doi.org/10.1145/963770.963772>
- Peter Ingwersen and Kalervo Järvelin. 2005. *The Turn: Integration of Information Seeking and Retrieval in Context (The Information Retrieval Series)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA.
- Rosie Jones and Kristina Lisa Klinkner. 2008. Beyond the session timeout: automatic hierarchical segmentation of search topics in query logs. In *Proceedings of the 17th ACM conference on Information and knowledge management (CIKM '08)*. ACM, New York, NY, USA, 699–708. DOI: <http://dx.doi.org/10.1145/1458082.1458176>
- Evangelos Kanoulas, Ben Carterette, Paul D. Clough, and Mark Sanderson. 2011a. Evaluating multi-query sessions. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval (SIGIR '11)*. ACM, New York, NY, USA, 1053–1062. DOI: <http://dx.doi.org/10.1145/2009916.2010056>
- Evangelos Kanoulas, Ben Carterette, Paul D. Clough, and Mark Sanderson. 2011b. Overview of the TREC 2010 Session Track. In *Proceedings of the 19th Text Retrieval Conference (TREC '10)*.
- Evangelos Kanoulas, Mark M. Hall, Paul Clough, Ben Carterette, and Mark Sanderson. 2012. Overview of the TREC 2011 Session Track. In *Proceedings of the 20th Text Retrieval Conference (TREC '11)*.
- Alexandros Karatzoglou, Xavier Amatriain, Linas Baltrunas, and Nuria Oliver. 2010. Multiverse Recommendation: N-dimensional Tensor Factorization for Context-aware Collaborative Filtering. In *Proceedings of the Fourth ACM Conference on Recommender Systems (RecSys '10)*. ACM, New York, NY, USA, 79–86. DOI: <http://dx.doi.org/10.1145/1864708.1864727>
- Diane Kelly, Jaime Arguello, and Robert Capra. 2013. NSF Workshop on Task-based Information Search Systems. *SIGIR Forum* 47, 2 (Jan. 2013), 116–127. DOI: <http://dx.doi.org/10.1145/2568388.2568407>

- Alexander Kotov, Paul N. Bennett, Ryen W. White, Susan T. Dumais, and Jaime Teevan. 2011. Modeling and analysis of cross-session search tasks. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval (SIGIR '11)*. ACM, New York, NY, USA, 5–14. DOI: <http://dx.doi.org/10.1145/2009916.2009922>
- Lin Li, Zhenglu Yang, Botao Wang, and Masaru Kitsuregawa. 2007. Dynamic Adaptation Strategies for Long-term and Short-term User Profile to Personalize Search. In *Proceedings of the Joint 9th Asia-Pacific Web and 8th International Conference on Web-age Information Management Conference on Advances in Data and Web Management (APWeb / WAIM'07)*. Springer-Verlag, Berlin, Heidelberg, 228–240. <http://dl.acm.org/citation.cfm?id=1769708.1769742>
- Yize Li, Jiazhong Nie, Yi Zhang, Bingqing Wang, Baoshi Yan, and Fuliang Weng. 2010. Contextual Recommendation Based on Text Mining. In *Proceedings of the 23rd International Conference on Computational Linguistics: Posters (COLING '10)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 692–700. <http://dl.acm.org/citation.cfm?id=1944566.1944645>
- Jingjing Liu and Nicholas J. Belkin. 2010. Personalizing information retrieval for multi-session tasks: the roles of task stage and task type. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval (SIGIR '10)*. ACM, New York, NY, USA, 26–33. DOI: <http://dx.doi.org/10.1145/1835449.1835457>
- Ernesto W. De Luca, Matthias Böhmer, Alan Said, and Ed H. Chi (Eds.). 2012. *CaRR '12: Proceedings of the 2Nd Workshop on Context-awareness in Retrieval and Recommendation*. ACM, New York, NY, USA.
- Ernesto W. De Luca, Alan Said, Matthias Böhmer, and Florian Michahelles (Eds.). 2011. *CaRR '11: Proceedings of the 2011 Workshop on Context-awareness in Retrieval and Recommendation*. ACM, New York, NY, USA.
- Zhongming Ma, Gautam Pant, and Olivia R. Liu Sheng. 2007. Interest-based Personalized Search. *ACM Transactions on Information Systems (TOIS)* 25, 1, Article 5 (Feb. 2007). DOI: <http://dx.doi.org/10.1145/1198296.1198301>
- Gary Marchionini. 2006. Exploratory search: from finding to understanding. *Communications of the ACM (CACM)* 49, 4 (April 2006), 41–46. DOI: <http://dx.doi.org/10.1145/1121949.1121979>
- Nicolaas Matthijs and Filip Radlinski. 2011. Personalizing Web Search Using Long Term Browsing History. In *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining (WSDM '11)*. ACM, New York, NY, USA, 25–34. DOI: <http://dx.doi.org/10.1145/1935826.1935840>
- Massimo Melucci. 2012. Contextual Search: A Computational Framework. In *Foundations and Trends in Information Retrieval*, Douglas W. Oard and Mark Sanderson (Eds.). Vol. 6. Now Publishers Inc., Hanover, MA, USA.
- Lilyana Mihalkova and Raymond Mooney. 2009. Learning to Disambiguate Search Queries from Short Sessions. In *Proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases: Part II (ECML PKDD '09)*. Springer-Verlag, Berlin, Heidelberg, 111–127. DOI: http://dx.doi.org/10.1007/978-3-642-04174-7_8
- Dan Morris, Meredith Ringel Morris, and Gina Venolia. 2008. SearchBar: a search-centric web history for task resumption and information re-finding. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 1207–1216. DOI: <http://dx.doi.org/10.1145/1357054.1357242>
- Cosimo Palmisano, Alexander Tuzhilin, and Michele Gorgoglione. 2008. Using Context to Improve Predictive Modeling of Customers in Personalization Applications. *IEEE Transactions on Knowledge and Data Engineering (TKDE)* 20, 11 (Nov. 2008), 1535–1549. DOI: <http://dx.doi.org/10.1109/TKDE.2008.110>
- Karthik Raman, Paul N. Bennett, and Kevyn Collins-Thompson. 2014. Understanding Intrinsic Diversity in Web Search: Improving Whole-Session Relevance. *ACM Transactions on Information Systems (TOIS)* 32, 4, Article 20 (Oct. 2014), 45 pages. DOI: <http://dx.doi.org/10.1145/2629553>
- Steffen Rendle, Zeno Gantner, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2011. Fast Context-aware Recommendations with Factorization Machines. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '11)*. ACM, New York, NY, USA, 635–644. DOI: <http://dx.doi.org/10.1145/2009916.2010002>
- Paul Resnick and Hal R. Varian. 1997. Recommender Systems. *Communications of the ACM (CACM)* 40, 3 (March 1997), 56–58. DOI: <http://dx.doi.org/10.1145/245108.245121>
- Peter H. Rossi, Mark W. Lipse, and Howard E. Freeman. 2004. *Evaluation: A Systematic Approach*. Sage, Thousand Oak, CA, USA.
- Alan Said, Ernesto W. De Luca, Daniele Quercia, and Matthias Böhmer. 2014. 4th Workshop on Context-Awareness in Retrieval and Recommendation. In *Advances in Information Retrieval*, Maarten de Rijke, Tom Kenter, Arjen P. de Vries, ChengXiang Zhai, Franciska de Jong, Kira Radinsky, and Katja Hof-

- mann (Eds.). Lecture Notes in Computer Science, Vol. 8416. Springer International Publishing, 802–805. DOI: http://dx.doi.org/10.1007/978-3-319-06028-6_100
- Milad Shokouhi, Ryen W. White, Paul Bennett, and Filip Radlinski. 2013. Fighting Search Engine Amnesia: Reranking Repeated Results. In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '13)*. ACM, New York, NY, USA, 273–282. DOI: <http://dx.doi.org/10.1145/2484028.2484075>
- Barry Smyth, Evelyn Balfe, Jill Freyne, Peter Briggs, Maurice Coyle, and Oisín Boydell. 2005. Exploiting Query Repetition and Regularity in an Adaptive Community-Based Web Search Engine. *User Modeling and User-Adapted Interaction* 14, 5 (Jan. 2005), 383–423. DOI: <http://dx.doi.org/10.1007/s11257-004-5270-4>
- David Sontag, Kevyn Collins-Thompson, Paul N. Bennett, Ryen W. White, Susan Dumais, and Bodo Billerbeck. 2012. Probabilistic models for personalizing web search. In *Proceedings of the Fifth ACM international conference on Web search and data mining (WSDM '12)*. ACM, New York, NY, USA, 433–442. DOI: <http://dx.doi.org/10.1145/2124295.2124348>
- Mirco Speretta and Susan Gauch. 2005. Personalized search based on user search histories. In *Proceedings of the 2005 IEEE / WIC / ACM International Conference on Web Intelligence (WI '05)*. 622–628.
- Kazunari Sugiyama, Kenji Hatano, and Masatoshi Yoshikawa. 2004. Adaptive Web Search Based on User Profile Constructed Without Any Effort from Users. In *Proceedings of the 13th International Conference on World Wide Web (WWW '04)*. ACM, New York, NY, USA, 675–684. DOI: <http://dx.doi.org/10.1145/988672.988764>
- Bin Tan, Xuehua Shen, and ChengXiang Zhai. 2006. Mining Long-term Search History to Improve Search Accuracy. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '06)*. ACM, New York, NY, USA, 718–723. DOI: <http://dx.doi.org/10.1145/1150402.1150493>
- Jaime Teevan, Susan T. Dumais, and Eric Horvitz. 2005. Personalizing Search via Automated Analysis of Interests and Activities. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '05)*. ACM, New York, NY, USA, 449–456. DOI: <http://dx.doi.org/10.1145/1076034.1076111>
- Jaime Teevan, Daniel J. Liebling, and Gayathri Ravichandran Geetha. 2011. Understanding and Predicting Personal Navigation. In *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining (WSDM '11)*. ACM, New York, NY, USA, 85–94. DOI: <http://dx.doi.org/10.1145/1935826.1935848>
- Jaime Teevan, Meredith Ringel Morris, and Steve Bush. 2009. Discovering and Using Groups to Improve Personalized Search. In *Proceedings of the Second ACM International Conference on Web Search and Data Mining (WSDM '09)*. ACM, New York, NY, USA, 15–24. DOI: <http://dx.doi.org/10.1145/1498759.1498786>
- Ryen W. White, Paul N. Bennett, and Susan T. Dumais. 2010. Predicting short-term interests using activity-based search context. In *Proceedings of the 19th ACM international conference on Information and knowledge management (CIKM '10)*. ACM, New York, NY, USA, 1009–1018. DOI: <http://dx.doi.org/10.1145/1871437.1871565>
- Biao Xiang, Daxin Jiang, Jian Pei, Xiaohui Sun, Enhong Chen, and Hang Li. 2010. Context-aware Ranking in Web Search. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '10)*. ACM, New York, NY, USA, 451–458. DOI: <http://dx.doi.org/10.1145/1835449.1835525>