From Republicans to Teenagers – Group Membership and Search (GRUMPS)

Ingmar Weber[1], Djoerd Hiemstra[2], and Pavel Serdyukov[3]

[1] Qatar Computing Research Institute, Doha, Qatar
`ingmarweber@acm.org`
[2] University of Twente, Twente, The Netherlands
`d.hiemstra@utwente.nl`
[3] Yandex, Moscow, Russia
`pavser@yandex-team.ru`

Abstract. In the early years of information retrieval, the focus of research was on systems aspects such as crawling, indexing, and relevancy ranking. Over the years, more and more user-related information such as click information or search history has entered the equation creating more and more personalized search experiences, though still within the scope of the same overall system. Though fully personalized search is probably desirable, this individualistic perspective does not exploit the fact that a lot of a users behavior can be explained through their group membership. Children, despite individual differences, share many challenges and needs; as do men, Republicans, Chinese or any user group. This workshop takes a group-centric approach to IR and invites contributions that either (i) propose and evaluate IR systems for a particular user group or that (ii) describe how the search behavior of specific groups differ, potentially requiring a different way of addressing their needs.

Keywords: information retrieval, user groups, user modeling.

1 Background

IR systems traditionally had a global, one-size-fits-all approach to retrieval with the result presentation and the search process ignoring the users identity, their history and attributes. Over the last years, an increasing amount of work has been devoted to personalizing search (including sponsored search). This personalization typically happens in a manner transparent to the user based on the users current and past interaction with the system. Such an approach, despite its targeting potential and accuracy, (i) often comes with problems related to cold-start, (ii) generally leads to few high-level insights to why a certain resource was surfaced, and (iii) is based on the assumption that the differences among users are small enough such that all users still fit under the same general umbrella underlying the mathematical models. For certain user groups these disadvantages have led to the creation of alternative IR systems. Examples include search systems for the visually impaired or systems designed for children. This workshop

wants to motivate more work on group-oriented IR and give researchers a platform to discuss advances in this area. This research area will be of interest to many ECIR attendees and related work has been presented at past instances of the conference.

The workshop will welcome contributions on any issue concerning adapting search for specific user groups and, in particular, on:

- Understanding of search behavior and relevance criteria of specific user groups [1–3]
- Identifying labeled user groups from search behavior [4–7]
- Using specific user features (e.g., demographics) for personalized advertising/sponsored search [8, 9]
- Understanding the effects of domain expertise, age, user experience and cognitive abilities on search goals and results evaluation [10–13]
- Development of test collections for evaluation of personalized search systems focused on users with specific features [14–16]
- Search interfaces and result representation for people with specific needs and specific interests [17–20]

2   Workshop Program

The full-day workshop will have oral presentations of accepted research papers. There will also be a 45 minute impromptu featherweight presentation session where attendees can present topics of their choice for 2-5 minutes. This is done in an effort to improve awareness of related work and to foster personal connections between the participants.

Additionally, there will be two invited talks. Daria Sorokina, senior data scientist at LinkedIn, will talk about "Recruiters, Job Seekers and Spammers: Innovations in Job Search at LinkedIn", and Nicholas Belkin, professor at Rutgers University, will talk about "How to Know About the Information Needs and Behaviors of User Groups".

3   Workshop Organizers

Ingmar Weber is a research scientist at Yahoo! Research in Barcelona, Spain. His research covers a wide subject area, often with a focus on demographic factors and their influence on search behavior and with a lot of his current work involving interdisciplinary studies in web science. In August 2012 he taught the course "An Introduction to Web Science" at the RuSSIR summer school. Together with Ana-Maria Popescu and Marco Pennacchiotti he organized the PLEAD (Politics, Elections and Data) workshop at CIKM 2012. He did his Ph.D. at the Max-Planck Institute for Informatics working on CompleteSearch, which is now the default search engine for DBLP.

Djoerd Hiemstra is associate professor at the database group of the University of Twente. He wrote an often cited Ph.D. thesis on language models for

information retrieval and contributed to over 150 research papers in the field of information retrieval. His research interests include formal models of information retrieval, XML search and multimedia search, peer-to-peer search, and search systems for children. Djoerd was involved in the local organization of ACM SIGIR 2007 conference in Amsterdam, and in the organization of several workshops including several editions of the Dutch-Belgian Information Retrieval Workshop series, and the Accessible Search Workshop at SIGIR 2010.

Pavel Serdyukov is the head of research projects at Yandex. He has published a number of papers relating to web search (including adapting search for children and other socio-demographic groups), enterprise/entity search, query log analysis, location-specific retrieval and recommendation. He co-organized several workshops at SIGIR: on Accessible Search in 2010 and on Entity-oriented search in 2011 and 2012. He is a co-organizer of the "Web Search Click Data" workshop at WSDM 2012 and 2013. He was a co-organizer of the Entity track at TREC 2009-2011. Currently, he serves as a General Co-Chair of ECIR 2013.

References

1. Maglaughlin, K., Sonnenwald, D.: User perspectives on relevance criteria: A comparison among relevant, partially relevant, and not-relevant judgments. Journal of the American Society for Information Science and Technology 53(5), 327–342 (2002)
2. Tombros, A., Ruthven, I., Jose, J.M.: How users assess web pages for information seeking. Journal of the American Society for Information Science and Technology 56(4), 327–344 (2005)
3. Barry, C.L.: User-defined relevance criteria: an exploratory study. J. Am. Soc. Inf. Sci. 45(3), 149–159 (1994)
4. Jones, R., Kumar, R., Pang, B., Tomkins, A.: "i know what you did last summer": query logs and user privacy. In: Proceedings of the Sixteenth ACM Conference on Information and Knowledge Management, pp. 909–914 (2007)
5. Teevan, J., Morris, M.R., Bush, S.: Discovering and using groups to improve personalized search. In: Proceedings of the Second ACM International Conference on Web Search and Data Mining, pp. 15–24 (2009)
6. Weber, I., Castillo, C.: The demographics of web search. In: SIGIR, pp. 523–530 (2010)
7. Weber, I., Jaimes, A.: Who uses web search for what: and how. In: Proceedings of the Fourth ACM International Conference on Web Search and Data Mining, WSDM 2011, pp. 15–24 (2011)
8. Cheng, H., Cantú-Paz, E.: Personalized click prediction in sponsored search. In: Proceedings of the Third ACM International Conference on Web Search and Data Mining, pp. 351–360 (2010)
9. Jansen, B.J., Moore, K., Carman, S.: Evaluating the performance of demographic targeting using gender in sponsored search. Inf. Process. Manage. 49(1), 286–302 (2013)
10. Sharit, J., Hernández, M.A., Czaja, S.J., Pirolli, P.: Investigating the roles of knowledge and cognitive abilities in older adult information seeking on the web. ACM Trans. Comput.-Hum. Interact. 15(1), 3:1–3:25 (2008)

11. Duarte Torres, S., Weber, I.: What and how children search on the web. In: Proceedings of the 20th ACM International Conference on Information and Knowledge Management, pp. 393–402 (2011)
12. Thatcher, A.: Web search strategies: The influence of web experience and task type. Inf. Process. Manage. 44(3), 1308–1329 (2008)
13. Gossen, T., Low, T., Nürnberger, A.: What are the real differences of children's and adults' web search. In: SIGIR, pp. 1115–1116 (2011)
14. Yue, Z., Harplale, A., He, D., Grady, J., Lin, Y., Walker, J., Gopal, S., Yang, Y.: Citeeval for evaluating personalized social web search. In: SIGIR Workshop on the Future of IR Evaluation, pp. 23–24 (2009)
15. Matthijs, N., Radlinski, F.: Personalizing web search using long term browsing history. In: Proceedings of the Fourth ACM International Conference on Web Search and Data Mining, pp. 25–34 (2011)
16. Chernov, S., Serdyukov, P., Chirita, P.-A., Demartini, G., Nejdl, W.: Building a Desktop Search Test-Bed. In: Amati, G., Carpineto, C., Romano, G. (eds.) ECIR 2007. LNCS, vol. 4425, pp. 686–690. Springer, Heidelberg (2007)
17. Leporini, B., Andronico, P., Buzzi, M.: Designing search engine user interfaces for the visually impaired. In: Proceedings of the 2004 International Cross-Disciplinary Workshop on Web Accessibility (W4A), pp. 57–66 (2004)
18. Chevalier, A., Rozencwajg, P., Desjours, B.: Impact of Prior Knowledge and Computer Interface Organization in Information Searching Performances: A Study Comparing Younger and Older Web Users. In: Stephanidis, C. (ed.) Posters, HCII 2011, Part I. CCIS, vol. 173, pp. 373–377. Springer, Heidelberg (2011)
19. Druin, A., Foss, E., Hatley, L., Golub, E., Guha, M.L., Fails, J., Hutchinson, H.: How children search the internet with keyword interfaces. In: Proceedings of the 8th International Conference on Interaction Design and Children, pp. 89–96 (2009)
20. Kerkmann, F., Lewandowski, D.: Accessibility of web search engines: Towards a deeper understanding of barriers for people with disabilities. Library Review 61, 608–621 (2012)