Report on EuroHCIR 2013: The 3rd European Workshop on Human-Computer Interaction and Information Retrieval

Preben Hansen
Stockholm University, Sweden

Max L. Wilson
University of Nottingham, UK

Birger Larsen
Aalborg University Copenhagen, Denmark

Kristian Norling
Norling & Co, Sweden

Tony Russell-Rose
UXLabs, UK

Abstract

This paper reports on the 3rd European Workshop on Human-Computer Interaction and Information Retrieval co-located with the 36th international ACM SIGIR conference on Research and development in information retrieval, held in Dublin, Ireland, August 1, 2013. The workshop gathered 32 participants with 6 papers for oral presentations and 8 short papers with poster presentations covering a good range of topics and research questions in the intersection of HCI and IR. Group exercises in connection to the oral presentations resulted in a set of important aspects to be considered in future research in HCI and IR.

1 Introduction

HCIR is a phrase coined by Marchionini in 2005 [4] and is representative of the growing interest in uniting both those who are interested in how information systems are built (IR community) and those who are interested in how humans search for information and how to interact and represent IR features and functionalities (HCI and Information Seeking). For a long time, both these fields have both developed innovative techniques to support search and search strategies. However, insights from both communities often do not bridge disciplinary borders and therefore, important knowledge may not be utilized. Five increasingly popular HCIR workshops brought focus to this multi-disciplinary issue in the USA. Alongside the American HCIR series, EuroHCIR aims to stimulate this focus in Europe.

The EuroHCIR workshops have three main goals:

1. Present and discuss novel HCIR designs, systems, and findings.

2. Identify and unite European researchers and industry professionals working in this area.

3. Facilitate and encourage collaboration and joint academic and industry ventures.

In addition to the focus of investigating Human-Computer Interaction (HCI) issues relating to Information Retrieval (IR), the 4th Goal of EuroHCIR 2013 was to better engage with the needs of the Information Retrieval community at SIGIR. The previous two workshops have taken place in other settings: at BCS-HCI [13], a HCI focused conference, and co-located with IIiX2012 [14], an Interactive IR symposium. The previous two workshops engaged the human-focused communities well, and led to many design and human-evaluation papers, but fewer IR measure and systems papers. One notable exception was Smucker & Clarke's introduction to Time-Based Gain measure [8], which is the kind of paper that we wanted to encourage in the EuroHCIR 2013 workshop. We hoped that the community at SIGIR2013, within the EuroHCIR workshop would be able to focus on extended search scenarios like holiday planning, entertainment, and casual browsing, especially with the growing interest in Search Sessions [2, 6]. The increased focus on users at SIGIR2013, with several user-oriented paper sessions, was evidence of this interest outside of our workshop too. Consequently, we believe that our proposal to organise a EuroHCIR workshop at SIGIR was a crucial step towards creating a space to encourage more IR and HCI integration and was well timed.

2 Workshop Structure

The workshop began by examining the spread of participants in the room. All participants tagged themselves with varying sized blue and red dots, where systems focused people had larger mostly blue dots, perhaps with a small human-focused red dot. Conversely, more human-focused participants had larger red dots, and smaller blues. We were pleased to have at least one participant label themselves as entirely red, and another entirely blue, indicating that we had attracted participants with both polar research interests, rather than just those with overlapping research interests. This categorisation process allowed us to create groups in the room that would represent different areas of expertise. Although this appeared to create a divide between different expertise, the rest of the workshop involved feeding back to the room, which gave a voice to each research perspective without a unmanagable 32-way discussion. This process meant that everyone in the room began to understand the issues identified and considered by people with differing expertise.

2.1 Accepted Papers

The accepted papers were divided into two different formats: 6 working sessions each focused on 1 of 6 oral presentations, and another session with poster presentations preceded by lightning talks to announce the poster topics. In the working sessions, each oral presentation lasted for about 10 minutes followed by a 15 minutes of group discussions. Finally each group had around 1 minute to report back to the whole workshop. During the discussions, each group had access to a Google document to make note of important issues identified by their group.

2.2 Working Sessions: Paper 1-6

Paper 1: Thomas, Scholer and Moffat [9] Fading Away: Dilution and User Behaviour
This paper was a focused example of when it is difficult to determine how system changes will affect

search behaviour. The authors purposefully degraded a system, using jargon terms, to investigate relevance decisions, but failed to see the reduced performance seen in previous literature. They discussed possible causes of this difference.

The user groups were primarily focused on a) the reactions of users, rather than logged behaviours, and b) the result of different task types. For a) they concluded that measures such as confidence in their results or satisfaction might have led to differences, where behaviour is approximately the same or not measurably different. For b) they considered that learning tasks might be more affected than single-result finding tasks where they trust the ranking. Consequently, they considered the possible behaviour change when results were ranked differently.

The user-system groups focused on a) the behavioural differences created by experiments rather than systems, and b) other system changes that might lead to behaviour changes instead. They also wondered if elements such as task type or user interface changes would have seen a different effect. Finally, they also considered ways to tuning features like degrading and measuring behaviour differences at a larger scale.

Finally, the system group focused on design choices that would make users move to a second page of results. They considered whether it was possible to degrade systems enough that participants move to a second page, or whether degradation would lead instead to a switch of search engine.

Paper 2: Potthast et al. [5]. Exploratory Search Missions for TREC topics.
This paper focused on the development of a corpus of user behaviour data produced by participants of a user study. The participants were professional writers, researching topics and writing essays; a considerable sum of money was invested in this recruitment. The corpus of behaviour data is freely available to all.

The user groups discussed the translation of writing tasks to querying behaviour, and whether this was determinable within the dataset. They also considered whether topic difficulty was determinable in the dataset, as this would allow them to examine different behaviours. They concluded that it might be most interesting to investigate search evolution within the dataset, such as topic switching, task time, and how aware users are of time versus effort.

The user-system group listed the following questions about the data: does the data contain retrieval and document use information? And is demographic information available on authors? They hoped to cluster people by behaviour and approach based on cultural background. They concluded that more description of the task was needed and they needed to consider the bias of the single UI used. Finally, they discussed how to visualize the process of project stages e.g. parallel searches. The system group mentioned that log data could be used for modelling exploratory tasks, but considered that the quality of the retrieval mechanism would have a big effect user interaction.

Paper 3: Hall, Katsaris and Toms [11]. A Pluggable Work-bench for Creating Interactive IR Interfaces
This paper described the development of a framework for creating quick and dirty, plug and play, user interfaces. The utility is described for making quick comparisons of user interface ideas, and could be easily used for student projects. More detailed and custom uses are available for more detailed research projects. The framework is freely available onlne, but some discussion in the room focused on the cost/benefits of using existing frameworks to make customisations.

For the group with participants representing the user perspective, they thought that a plug-

gable work-bench for creating interactive IR interfaces could be good. This kind of platform could be good for student projects, replication or teaching. It facilitates rapid UI component development but there is a balance to be found between simplicity and flexibility.

The user-system group noted that the results collected through variants of the system need to be comparable; that one possibly could be to replicate the same experiment across multiple groups to test consistency/replicability of experiments. They also wondered if the same system could be used to recruit, group, and allocate participants to conditions of the experiments, while managing their tasks.

Finally, the main focus of the system group was that it would be nice if the system better supported easy extension of customized widgets. **Paper 4: Chirag Shah [7]. A Proposal for User-Focused Evaluation and Prediction of Information Seeking Process.**
Shah proposed a machine-learning based method to dynamically evaluate and predict search performance several time-steps ahead at each given time point of the search process during an exploratory search task.

The user group discussed the question of prediction, and the value of it, within their research. They considered ways to build prediction into user interfaces, search suggestions, and adaptive changes to support the user.

As the presented work was a description of work that will soon be undertaken, the user-system group considered the feasiblity of reliably knowing what users 'know' and what they are thinking, or whether they think at all. They also discussed the ambiguity of actions in relation to optimal behaviour: how do you know whether re-reading a document is desirable behavior or undesirable? Further, the group discussed how it would be possible to estimate 'expected value' of actions to a user? Sometimes the unexpected is desirable, and individual differences and behavioural patterns can be important.

Finally, the system group discussed that it was surprising that the authors considered more data would lead to lower prediction accuracy. They also considered how frequently you would see sessions with sufficient number of queries and data points to make a prediction.

Paper 5: Gossen, Nitsche, and Nürnberger [1]. Evolving Search User Interfaces
The authors introduced the concept of an evolving search user interface (ESUI). It adapts the UI dynamically based on the derived capabilities of the user interacting with it. The presentation focused on user characteristics that change over time, and discussed how each of them can influence the SUI design using an example of a girl growing from six to fourteen.

The user groups discussed priming effects and individualism, such as familiarity with a technique like coverflow. They discussed the importance of age and context appropriate results, and thus why only adapt the front-end and not the back-end? Finally, they discussed the impact of learning to adapt to adapting user interfaces, and the preference users show for familiarity.

The user-system groups mentioned that one should be able to change the interface based on task complexity. One should give children the opportunity to choose their own set-up, rather than adapt to expectations about them. They too discussed the issues of handling individual differences, time and experience. They also suggested that behavioural cues from the work by Shah could be used to adapt the interface.

The system groups focused on the cautions with personalisation and changing entire interfaces, while discussing the options to gradually and more subtly make changes to a system.

Paper 6: Simon Tretter, Gene Golovchinsky, Pernilla Qvarfordt [10]. SearchPanel:

A browser extension for managing search activity
The authors, including the now sadly missed Golovchinsky, designed and built a Chrome extension called SearchPanel that helps people manage their ongoing information seeking. The extension combines document and process metadata into an interactive representation of the retrieved documents that can be used for sensemaking, for navigation, and for re-finding documents.

Unlike the previous sessions, for this orally presented paper, teams were reorganised to create mixed expertise groups. The groups were asked for focus on how such a system might be evaluated using mixed methods.

The questions that emerged during this discussion included: Can we find task types reliably? Could the data be good for learning to rank? How will SerchPanel change user behaviour? Are users able to explore more results? Does showing non-clicked results help - do users care? Is length of a query a factor? What kinds of tasks is this tool useful for? Finally, Re-find information - is it more efficient to influence re-finding behavior in a naturalistic context?

For user-focused methods, group discussions ranged from evaluating engagement through surveys, diary studies and long term studies. Another take on the evaluation was to elicit insights by using two tasks: one exploratory-oriented and one fact-finding task. A controlled study (with or without the search panel) could measure the completion time and provide the users with a questionnaire. One group suggested that an eye-tracking study could be added to see what people actually do. Others within the workshop suggested it was best to run an uncontrolled experiment and analyse the interactions with different parts of the UI with users in order to log actiual use.

3 Discussion and Conclusions

The third European Workshop on Human-Computer Interaction and Information Retrieval brought together highly motivated researchers interested the overlapping area of Information Retrieval and Human Computer Interaction. A factor that resulted in the successful outcome was the creation of distinct groups with differing expertise and perspectives. Each group represented one aspect of either the system-oriented side and/or the user-oriented side. Each group had to find arguments, criticize and discuss based on their 'side' they represented. This stimulated discussion and exchange of interesting ideas between the groups of researchers and the authors in different areas, ranging from human-oriented information seeking and task-based IR, to user interface design and to retrieval systems technology. In this way a range of ideas, suggestions, proposals and statements emerged as a contribution not only to the authors of ALL papers accepted, but also to all workshop participants. These aspects covered both system-side as well as user side of the issue of human-computer interaction and Information Retrieval. Details about the workshop including the papers, presentations and slides can be found at: http://ceur-ws.org/Vol-1033/

4 Acknowledgement

We would like to thank ACM and SIGIR for hosting this workshop, especially the SIGIR team of local organizers. We dedicate this workshop report to the memory of Gene Golovchinsky, a great scientist as well as a keen supporter and participant of HCIR and EuroHCIR, who attended the workshop only a few weeks before his untimely death.

References

[1] T. Gossen, M. Nitsche, and A. Nürnberger. Evolving search user interfaces. In Wilson et al. [12], pages 31–34.

[2] D. Guan, S. Zhang, and H. Yang. Utilizing query change for session search. In Jones et al. [3], pages 453–462.

[3] G. J. F. Jones, P. Sheridan, D. Kelly, M. de Rijke, and T. Sakai, editors. *The 36th International ACM SIGIR conference on research and development in Information Retrieval, SIGIR '13, Dublin, Ireland - July 28 - August 01, 2013*. ACM, 2013.

[4] G. Marchionini. Toward human-computer information retrieval. *Bulletin of the American Society for Information Science and Technology*, 32(5):20–22, 2006.

[5] M. Potthast, M. Hagen, M. Völske, and B. Stein. Exploratory search missions for trec topics. In Wilson et al. [12], pages 7–10.

[6] K. Raman, P. N. Bennett, and K. Collins-Thompson. Toward whole-session relevance: exploring intrinsic diversity in web search. In Jones et al. [3], pages 463–472.

[7] C. Shah. A proposal for user-focused evaluation and prediction of information seeking process. In Wilson et al. [12], pages 39–42.

[8] M. D. Smucker and C. L. A. Clarke. The fault, dear researchers, is not in cranfield, but in our metrics, that they are unrealistic. In Wilson et al. [14], pages 11–12.

[9] P. Thomas, F. Scholer, and A. Moffat. Fading away: Dilution and user behaviour. In Wilson et al. [12], pages 3–6.

[10] S. Tretter, G. Golovchinsky, and P. Qvarfordt. Searchpanel: A browser extension for managing search activity. In Wilson et al. [12], pages 51–54.

[11] R. Villa, P. Clough, M. M. Hall, and S. Rutter. Search or browse? casual information access to a cultural heritage collection. In Wilson et al. [12], pages 19–22.

[12] M. L. Wilson, T. Russell-Rose, B. Larsen, P. Hansen, and K. Norling, editors. *Proceedings of the 3rd European Workshop on Human-Computer Interaction and Information Retrieval, Dublin, Ireland, August 1, 2013*, volume 1033 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2013.

[13] M. L. Wilson, T. Russell-Rose, B. Larsen, and J. Kalbach, editors. *Proceedings of the 1st European Workshop on Human-Computer Interaction and Information Retrieval, Newcastle, UK, July 4th, 2011*, volume 763 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2011.

[14] M. L. Wilson, T. Russell-Rose, B. Larsen, and J. Kalbach, editors. *Proceedings of the 2nd European Workshop on Human-Computer Interaction and Information Retrieval, Nijmegen, The Netherlands, August 25, 2012*, volume 909 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2012.