
WORKSHOP REPORT

HCIR 2012: The Sixth International Symposium on Human-Computer Interaction and Information Retrieval

Robert Capra¹
University of North Carolina
at Chapel Hill
rcapra@unc.edu

Gene Golovchinsky
FX Palo Alto
Laboratory
gene@fxpal.com

Bill Kules
The Catholic University
of America
kules@cua.edu

Catherine L. Smith
Kent State University
csmit141@kent.edu

Daniel Tunkelang
LinkedIn
dtunkelang@linkedin.com

Ryen W. White
Microsoft Research
ryenw@microsoft.com

¹ Symposium organizers are listed in alphabetical order.

For more information about the symposium series, please visit <http://hcir.info>.

Abstract

This report describes the 2012 Symposium on Human-Computer Interaction and Information Retrieval. Now in its sixth year, the two-day symposium (formerly a one-day workshop) was held in October in Cambridge, MA. The event brought together researchers and practitioners from academia, industry, and government and a range of disciplines for in-depth discussions in an informal atmosphere. The symposium attracted 75 attendees, over a third of which were from industry. New for this year, we accepted full papers that will be archived and published in the ACM Digital Library. We continued the HCIR Challenge, this year focusing on the problem of people and expertise finding, five in-depth system demonstrations, and audience selection of a challenge winner.

1 Introduction

Human-computer information retrieval (HCIR) is the study of information retrieval (IR) techniques that integrate human intelligence and algorithmic search to help people explore, understand and use large information spaces better. The fields of human-computer interaction (HCI) and IR have both developed innovative techniques to address the challenge of navigating complex information spaces, but their insights frequently fail to cross disciplinary boundaries. Human-computer information retrieval has emerged in academic research and industrial practice to bring together research in the fields of IR and HCI, in order to create new kinds of search systems that depend on continuous human control of the search process (Marchionini, 2006).

The HCIR Symposium provides a venue for in-depth discussion of models, tools, and evaluation methods at the intersection of human factors and search. The event has run annually since 2007 and has attracted growing interest from HCI and IR researchers, serving as bridge between the two communities. Its success inspired the EuroHCIR workshop, first held in July 2011 and now in its second year. The sixth HCIR symposium (HCIR 2012) was held October 4-5, 2012 at IBM Research in Cambridge, MA. This year's symposium attracted 75 attendees, and featured presentations and posters, with lots of opportunity for discussion among attendees.

2 The Symposium

Symposium activities included a keynote, full- and short- paper presentations, a highly-interactive poster session, and the HCIR Challenge in which participating groups built systems to help address the people and expertise finding problem using data provided by Mendeley.

2.1 Keynote

Professor Marti Hearst delivered an interesting and entertaining keynote address titled “Halloween Cauldron of Ideas for Research (HCIR)” in which she tried to outline some potential directions for future HCIR research. Her talk was divided into three themes: the seams of sense-making, multiple cognitive channels, and radical collaboration, and included several periods of audience discussion.

With respect to sense-making, Marti identified saving search results and how to stay in the “flow” (Csikszentmihalyi, 1975) as being under-represented in the existing literature and systems. Saving search results is an important aspect of exploratory, recall-oriented search. She illustrated some of the design challenges with an example from an existing information retrieval system that supports this functionality, but does not offer a smooth interface for doing so.

Flow that characterizes engaged, uninterrupted, intrinsically-motivated behavior is also important for information seeking. Marti talked about the need to reduce friction of interaction, and the need to integrate note-taking into exploration. She also mentioned some of the advantages of spatial hypertext for organizing found information, but wanted to see a more light-weight interaction that did not detract from the flow. During discussion, several other challenges were raised including, how to keep track of new query ideas, how to implement transclusion (Nelson, 1999), and how to manage interactions with standing queries and alerts.

In addition to text, Marti suggested that designers need to think about other modalities such as audio to augment certain kinds of interactions with information seeking systems. Integrated audio interaction is something of a missed opportunity because it is difficult to work with in current programming frameworks. Yet there are advantages to audio both as input and as output: Marti cited work by Soudian and Fels (2002) that showed that programmers are more likely to leave detailed and informative comments in audio than in text; audio interfaces for textbooks, designed specifically for people with visual impairments, have also been shown to be useful to people with normal sight. Despite some advantages of audio, significant challenges exist in integrating it into interfaces in cognitively-appropriate ways.

Marti’s third theme was titled “radical collaboration.” She described several variants on collaborative search, which she classified into four categories: “My results affect your results (and vice versa),” “my search is independent of yours (until I am ready to share),” “I’ll search more if I see you are (even if you are a fakebot),” and “I’ll search with 100 other people (crowdsourcing complex finding tasks).” These scenarios reflect some of the interesting and significant use cases for collaboration in information seeking, an important emerging area of HCIR.

Marti raised some interesting questions and challenges, some of which we will surely see discussed in next year’s meeting and in publications in related conferences such as SIGCHI and SIGIR.

2.2 Full Papers

For the first time we accepted full papers for the symposium. We made this addition to provide attendees with the opportunity to present complete descriptions of their research beyond the position papers or posters that have been offered at previous HCIR events. Full papers were 10 pages in the SIGCHI two-column format. They were peer-reviewed for novelty, rigor, and the quality of presentation according to the standards of first-tier conferences. Accepted papers appear in the ACM Digital Library. These papers were presented as 20-minute talks, with 10 minutes for questions.

Four high-quality papers were accepted following peer-review and discussion. Two papers focused on systems (supporting search by children and provenance-based search) and two papers focused on evaluation (metrics and task development).

Knowledge Journey: A Web Search Interface for Young Users

Marcus Nitsche described *Knowledge Journey*, a new user interface for Web search by young users. She explained challenges of developing an interface for this user group and discussed the rationale behind the design decisions that were made in its development. Knowledge Journey afforded search by querying and navigating menus, as well as offering an anthropomorphic guidance figure to provide emotional support to its young users and a result storage functionality (cutely represented in the interface by a treasure chest) to support cognitive recall. Gossen described a comparative study of the search interface against a traditional text search user interface provided by most current search engines. The findings of the study showed that participants preferred Knowledge Journey, and allowed the designers to learn which features participants liked or disliked.

Leyline: Provenance-Based Search Using a Graphical Sketchpad

Carlos Jensen described research on developing new models of information organization and techniques to improve the efficiency of existing search tools within organizations. The presented research focused on the use of file *provenance* in retrieval operations, providing a record of the information about the creation, use, and sharing of documents and their context, including collaborators, source documents or references. Jensen argued that the addition of this information could enhance the effectiveness of regular keywords and he described a system to investigate the potential of provenance in search. Jensen outlined the findings of a usability study that showed that provenance data is useful and desirable in search, and that an interface based on a graphical sketchpad met participants' usability and efficiency expectations.

Modeling User Variance in Time-Biased Gain

Mark Smucker presented research on modeling user variance (i.e., the differences between users for each topic used in evaluation) to more accurately estimate the effect size of differences between retrieval systems. He argued that modeling user variance is critical to understanding the impact of effectiveness differences on the actual user experience. If the variance of a difference is high, the effect on user experience will be low. Smucker described an extension to the stochastic simulation of time-biased gain (described originally by Smucker and Clarke (2012)) to model the variation between users. He described a validation of this new version of time-biased gain by showing that it produces distributions of gain that agree well with actual distributions produced by real users. Smucker argued that measuring the effect size of differences allows researchers to understand the extent to which predicted performance improvements matter to real users.

Assigning Search Tasks Designed to Elicit Exploratory Search Behaviors

Barbara Wildemuth and Luanne Freund described research on the search tasks assigned to participants during studies of exploratory search behaviors and exploratory search systems. She demonstrated through a review of past studies that exploratory search tasks have a number of properties: they focus on learning and investigative search goals; they are general (rather than specific), open-ended, and often target multiple items/documents; they involve uncertainty and are motivated by ill-defined or ill-structured problems; they are dynamic and evolve over time; they are multi-faceted and may be procedurally complex; and they are often accompanied by other information or cognitive behaviors, such as sense making. Wildemuth concluded by proposing a set of recommendations for the design of exploratory search task descriptions. The presentation included an "audience participation" section during which Wildemuth solicited audience responses to examine how two example exploratory search tasks met the proposed properties.

2.3 Short Papers

Five short papers were selected for presentation in a ten-minute format, which included five minutes for presentation and five minutes for discussion.

The first paper, presented by video, proposed a model of consumer-oriented information behavior and search. Tony Russell-Rose described nine search *modes*, which reflect existing models of information seeking behavior from the research literature. The modes expand Marchionini's (2006) framework for exploratory search, and comprise behaviors that support the fundamental goals of *lookup*, *learn*, and *investigate*. Tony provided examples of design implications for several modes, emphasizing the need to bridge between models from the research community and design practice. In the next paper, David Karger presented a talk on structured data on the web and the need for emerging standards for data-bearing web pages. David argued that ubiquitous user interface patterns and affordances, such as faceted searching and display, present an opportunity to standardize the implementation of interaction and data visualization. As an extension to *html*, a vocabulary of common interactive methods could be implemented in the browser. These ideas have been piloted in the Exhibit framework, a set of prototype html extensions and an attendant JavaScript library. Third, Elena Agapie presented results from an experiment in persuasive computing. In order to motivate users to produce longer queries, in an experimental system, the search box changed as searchers typed. As a query grew longer, a halo around the box shifted color and size. In a pilot study of the system, searchers who saw the halo generated significantly longer queries than those who did not. In the next paper, Roberto González-Ibáñez presented results from a pilot study on a novel approach to collaboration. *Pseudo-collaboration* seeks to recognize and facilitate opportunities for productive collaboration among searchers with common information needs, but who might not otherwise be affiliated. While results showed benefits from pseudo-collaboration, as searches progressed over time the probability of the benefit was lower. Finally, Jeff Huang presented work on search interaction on touch-screen mobile devices and the challenges of obtaining meaningful behavioral data equivalent to mouse clicks. Jeff described the use of *viewport* coordinates (the visible portion of a Web page) as an indicator of searcher attention. He concluded by discussing of the need for better data collection methods across the non-standard mobile environment.

2.4 Posters

Sixteen posters were chosen for presentation, including position papers and work-in-progress reports.

Erik Choi presented results from a study that analyzed questions in four types of online Q&A services along four dimensions, finding patterns in the types of questions posed in the different Q&A services. Erik also presented related work to examine differences in users' motivations and expectations for using online Q&A services and virtual reference services in relation to general web search. Roberto González-Ibáñez described work to examine how emotions and affective processes influence how individuals search for information both individually and in collaborative teams. Jingjing Liu reported on results from a study to examine task difficulty from a user perspective, outlining reasons for task difficulty and categorizing them as related to the user, task, resource, and system. Youssef Meguebli had a poster that described an architecture for personalized information retrieval that uses a Dung argumentation framework and incorporates aspects of opinions, trust, search history, and social media. Marcus Nitsche presented an innovative system called CARSA that provides several search views that a user can switch among, along with features to use customized re-ranking algorithms, in-frame browsing, and integrated bookmark management. Marcus also described Trailblazer, which provides a number of features to support exploratory search, including an orthogonal overview layout of search result thumbnails to show search paths. Zhen Yue talked about a study conducted to investigate the search tactics used in two collaborative search tasks, finding differences in the need for awareness and explicit communication between the two tasks. Xiangmin Zhang described work to examine differences in document selection and relevance ratings between users with high and low

domain knowledge, reporting that users with high domain knowledge viewed more documents and assigned those documents higher ratings. Vladimir Zelevinsky described a process of conducting data mining using correlation statistics in combination with a faceted navigation and dynamic queries. Jingjing Liu and Xiaojun Yuan described the results of an experiment that found searchers with varying levels of domain knowledge chose different documents. Michael Zarro presented a developing model of exploratory search behavior and cognitive processes, set in the consumer health domain, where searchers are often motivated but lacking in knowledge and skills. Christopher Sasarak, David Stalnaker, Lei Hu, Robert LiVolsi, Siyu Zhu and Richard Zanibbi presented and demonstrated a system for entering math formula queries using a combination of finger or mouse-drawn symbols and keyboard input. Jacek Gwizdka described a model of user assessments of task difficulty, which draws on short-term eye movements to infer cognitive effort during search. Their results suggest that task difficulty is influenced more by early phases of the search session. Xiaojun Yuan presented the results of a comparative user study that found beneficial outcomes for the CiteSpace system, which provides visualizations of search results, over a text-based system, with shorter search times and higher subjective ratings. She also presented results of a related study that found searchers with more domain knowledge found the CiteSpace system more useful than those with low domain knowledge.

2.5 The HCIR Challenge

Following its successful introduction at the 2010 workshop, the HCIR Challenge is now in its third year. This year it focused on the problem of people and expertise finding. Mendeley (mendeley.com) provided this year's corpus: a database of over a million researcher profiles with associated metadata including published papers, academic status, disciplines, awards, and more taken from Mendeley's network of 1.6M+ researchers and 180M+ academic documents.

Five participating teams built systems that could perform three kinds of tasks: hiring based on a job description, assembling a conference program, and finding people to deliver patent research or expert testimony. Each of the teams was given 30 minutes to present at the symposium.

Magnifico: A Platform for Expert Mining Using Metadata

École Polytechnique Fédérale de Lausanne student Na Li presented joint work with Lei Zhou and Denis Gillet. Magnifico used a modified TF-IDF approach—where the IDF is an inverse discipline frequency—to match search queries to topic experts. It also assigned a multi-disciplinary reputation metric based on the expertise distribution of an author's readers.

Social Network Based Search for Experts

Ben-Gurion University student Dima Kagan presented joint work with Yehonatan Bitton, Michael Fire, Bracha Shapira, Lior Rokach, and Judit Bar-Ilan. Their system made use of additional publicly available data, cross-referencing the Mendeley user profiles with data from Academia.edu and using Microsoft Academic Search to categorize publication and journals.

IRIS-IPS: An Interactive People Search System for HCIR Challenge

University of Pittsburgh student Shuguang Han presented joint work with Daqing He, Zhen Yue, Jiepu Jiang, and Wei Jeng. The system used three different types of evidence to suggest candidates: expertise relevance, authority based on a PageRank algorithm applied to the co-authorship network, and social similarity using the Jaccard similarity between co-authors.

Exposing and Exploring Academic Expertise with Virtu

Luanne Freund and Kristof Kessler, both from the University of British Columbia (UBC), presented joint work with Michael Huggett and Edie Rasmussen. Virtu took a task-based approach exposing and giving the user control over dimensions of expertise that are more or less desirable depending on the type of expert-finding task. The search interface supported information interaction and exploration through a number of browsing and filtering tools, including facets and sliders.

'iF' People Search System

UCLA student Fei Liu presented a mobile application, designed for the iPad and supporting swipe and multi-touch gestures. A professional-grade application, iF offered a novel approach to exploring the corpus of documents and people via analyses of their reputations and social network relationships.

Although competition was fierce, Virtu stood out for the compelling approach it took to offering users control over the expert-finding process. Congratulations to the team from UBC!

3 Concluding Remarks

HCIR 2012 was an extraordinary success. Expanding the one-day workshop into a two-day symposium was particularly valuable: attendees appreciated the additional opportunity for discussion without being rushed by schedule constraints. The diversity of attendees was also notable: the record number of industry attendees brought a welcome perspective to the event. The feedback we have received from attendees has been extremely positive, and we are grateful to everyone who contributed to the event's success. Our special thanks goes to IBM Research for hosting the event, to our other sponsors, FXPAL, LinkedIn, Mendeley, Microsoft Research, MIT CSAIL, Oracle, and ACM SIGIR, who provided student travel awards. We are particularly grateful to all of the authors and presenters of papers, posters, and challenge entries, who did an outstanding job in both conducting such great research and in presenting it in an amazingly accessible manner.

We are already discussing HCIR 2013. In the meantime, look out for our special issue of Information Processing and Management on HCIR, likely to appear in 2013!

4 References

- Csikszentmihalyi, M. (1975). *Beyond Boredom and Anxiety*. San Francisco: Jossey-Bass.
- Marchionini, G. (2006). Exploratory search: from finding to understanding. *CACM*, 49(4): 41-46.
- Marchionini, G. (2006). Toward human-computer information retrieval. *Bulletin of the American Society for Information Science*, June/July.
- Nelson, T.H. (1999). Xanalogical structure, needed now more than ever: parallel documents, deep links to content, deep versioning, and deep re-use. *ACM Computing Surveys*, 31(4), 33.
- Smucker, M.D. and Clarke, C.L.A. (2012). Time-based calibration of effectiveness measures. *Proc. SIGIR*, 95-104.
- Soudian, S. and Fels, D.I. (2002). Verbal source code descriptor. *Proc. IEEE WESS'*.
-