

HCIR 2010: The Fourth International Workshop on Human-Computer Interaction and Information Retrieval

Robert Capra

University of North Carolina at Chapel-Hill
rcapra3@unc.edu

Bill Kules

The Catholic University of America
kules@cua.edu

Catherine L. Smith

Kent State University
csmit141@kent.edu

Daniel Tunkelang

Google
dt@google.com

Ryen W. White

Microsoft Research
ryenw@microsoft.com

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

Abstract

This report describes the 2010 workshop on Human-Computer Interaction and Information Retrieval. Now in its fourth year, the event was held in August 2010 in conjunction with the Information Interaction in Context Symposium. The workshop brought together researchers from academia, industry, and government and a range of disciplines to present and discuss their research. We had a record 70 attendees, making this the largest of our workshops to date. New for this year, we ran a challenge, and six research groups participated.

1 Introduction

Human-computer information retrieval (HCIR) is the study of information retrieval (IR) techniques that bring human intelligence into the search process. 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 workshop provides a venue for the presentation of advances in research 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. The fourth HCIR workshop (HCIR 2010) was held on Sunday, August 22, 2010 at Rutgers University in New Brunswick, NJ. This year the workshop attracted 70 attendees, making it the largest HCIR workshop to date, and was co-located with the Information Interaction in Context Symposium (IiX 2010).

2 The Workshop

Activities at the workshop included a keynote, six paper presentations, a highly-interactive poster and system demonstration session, and a first-time HCIR Challenge in which participating groups submitted novel search interfaces using an annotated corpus of articles from The New York Times.

2.1 Keynote

The keynote presentation for HCIR 2010 was given by Dan Russell, a leading researcher in search quality and user experience, and a research scientist at Google. In his talk, “Why is search easy and hard? Understanding serendipity and expertise in search,” Russell began by congratulating the IR and HCI communities for creating systems that now make many searches “easy,” but quickly outlined the challenges in supporting searches that are currently “hard.” These include exploratory searches, searches involving non-textual and visual elements, and searches for information that is not yet well-represented in existing retrieval systems, such as botanical images or ancient languages.

Russell described how improvements in search technology can change the nature of the questions that people can ask. As an example, he described using a new audio search tool to identify a song being played live by an Irish folk music group. Further expanding on the idea of how technology can change the questions that can be asked of search engines, Russell discussed the emerging strategy of crafting queries in terms of how users think others have asked the same or similar questions, leveraging the power of search suggestions based on other people’s queries. He gave as an example a successful search using the query *[oh oh oh oh]* and the real-time search completions offered by the search engine to find a song played at Stanford basketball games with those lyrics. In another example of the importance of how questions are asked, Russell presented results from a study that he conducted in which the use of different synonyms in the search task presentation had a dramatic effect on the time to find results and overall user satisfaction.

An overarching theme in Russell’s talk was that search technologies and interfaces are in constant flux, and these fluctuations alter the types of questions that we can ask search systems and how we ask our questions. He claimed that users are continually faced with challenges in how to present their queries and how to interpret and locate desired information from search results. HCIR researchers need to identify the strategies successful searchers use and work to not only support, but also to teach these strategies when users interact with a search interface.

Russell also argued for the need for better education and teaching about how to conduct searches. He described a comparison of search interface improvements that his group performed. They found that teaching someone to use the keystroke combination Control-F to search for a desired word or phrase on a Web page made a larger improvement than 16 alternative interface enhancements they had tested, with a reduction of 12% in their time-to-result measure. Furthermore, they found that over 90% of all US Internet users and 50% of all US teachers do not know how to use Control-F.

Russell described strategies for educating users about search features using both in-person instruction and offering interface elements that provide hints and tips at points of need. Specifically, Russell argued that teaching meta-skills for search was an effective strategy and outlined five important meta-skills to teach: (i) how to learn (looking for new terms, judging salience), (ii) how to manage attention, (iii) how to do research, (iv) how to search first, and (v) how to assess credibility.

2.2 Papers, Posters, and System Demonstrations

We accepted six papers for the presentation session; sadly, one of the presenters could not attend. The five presentations covered a variety of topics relating to tools, models, and evaluation for HCIR.

The first paper session focused on tools. Klaus Berberich described *NEAT*, a system to support the exploration of news archives. It extracts temporal expressions from within the articles as well as using publication dates to support temporal search. Search results are shown in a timeline based on both publication and internal temporal references. The timeline also shows major events, collected via crowdsourcing, to put the results into context. John Stasko presented *Jigsaw*, a system that integrates automated document analysis and interactive visualization to support interactive search via a set of tightly-coupled visualizations. Jigsaw offers grid, list and document views to allow the searcher to interactively sort, cluster, and explore document attributes and inter-relationships.

The second paper session focused on models and evaluation. Mark Smucker reported on a user study that investigated the relationship between retrieval precision and perceived search difficulty. By varying retrieval precision and measuring users' reported difficulty, Smucker and Jethani determined that precision affects both the perceived difficulty and the enjoyability of the search experience. Pertti Vakkari argued that exploratory search can be modeled as conceptual exploration, and that evaluation of exploratory search systems should address both their impact on the search process as well as the outcomes of the search. Using a concept-based approach, he described several measures of success in explicating information needs, ways to measure conceptual changes, and ways to assess term selection and query formulation. Max Wilson argued that "casual-leisure searching" breaks our current models of exploratory search. He presented several example searches based on a diary study and analysis of Twitter postings (tweets), including looking for entertainment on television, and "needless" browsing on eBay. He argued that these are exploratory but—in contrast to what recent research has focused on—they do not have an information-oriented goal. That is, they are not motivated by an information need. This sparked considerable discussion on the definition of a search.

Twenty papers were selected for poster presentation, including eight research papers and twelve position papers. Presenters introduced their work to attendees in short, one-minute "boasters" at the beginning of the workshop. At the poster session later in the day, some presenters elected to accompany their poster with a demonstration of their system. The research papers covered a broad range of IR genres, including library catalogs, health information, Web exploration, and Web, book, and image search. Research approaches included user studies, experiments, a log analysis of natural search, a literature review, and a simulation study. Research topics included faceted search, use of tag data, interface evaluation, eye-tracking, and sources of behavioral evidence. The position papers also covered a broad range of topics and issues. Two papers argued for the value of domain-specific research (patent search, for its task-specific characteristics; museum, archive, and special collections, for its highly structured content). Two papers elaborated on systems built for the HCIR Challenge (see Section 2.3 for more details on the challenge). Other papers addressed topics such as the nature of context, the role of domain knowledge in behavior, evidence of affect, integrated interfaces, user-centered design, network models for information exchange, and the support of task environments.

2.3 The HCIR Challenge

New to this year's workshop was an HCIR Challenge session. The challenge encouraged researchers and practitioners to build and demonstrate information access systems emphasizing user interaction and exploration, to address information needs ill-served by query-response systems. Thanks to the generosity of the Linguistic Data Consortium, participants were provided with free access to The New York Times (NYT) annotated corpus for the purpose of developing their challenge entries. The corpus comprised over 1.8 million articles published by The NYT between January 1, 1987 and July 19, 2007, annotated with rich metadata.

Participants were instructed to design their systems—and, if possible, perform pilot evaluations—to support users performing a variety of exploratory information-seeking tasks that emphasized the archival nature of the corpus. Tasks required users to research topics over the time span of the corpus, understand the competing perspectives on controversial topics, and answer questions that could not be resolved from individual documents. For example, one task required the user to enumerate the main arguments that have been made for and against rent control in New York.

Participants submitted four-page reports describing their work. The six challenge submissions were included in a workshop session, where participants presented their systems and optional evaluations. Afterward, workshop attendees voted for their favorite, which received the “people’s choice” trophy.

The participants demonstrated a diverse set of approaches. Corrado Boscarino presented a system that he, Arjen de Vries, and Wouter Alink developed based on a field theory of journalism, to help users identify the trustworthiness of news sources and authors. Christian Kohlschütter demonstrated *NewsClub*, a system already in production, that he applied to The NYT corpus to identify relevant terms for sets of documents and visualize their network of associations. Raman Chandrasekar presented *News Sync*, a system developed in partnership with Ann Paradiso, Jim St. George, Jeroen van den Eijkhof, and V. G. Vinod Vydiswaran. As the name suggests, the system uses summarization to help users catch up on news from a particular time period, location, or about a specific topic. Vladimir Zelevinsky demonstrated a faceted search system that lets users create custom facets based on WordNet. Wei Zheng presented a system, developed with Hui Fang, that combines topic modeling and sentiment analysis to identify and associate competing perspectives on an issue.

Finally, the “people’s choice” award went to *Time Explorer*, a system built by Michael Matthews, Pancho Tolchinsky, Roi Blanco, Jordi Atserias, Peter Mika, and Hugo Zaragoza of Yahoo! Labs. Time Explorer provides users with a way to explore the evolution of a topic over time. In addition to using publication dates; the system parses absolute and relative dates and time spans—even in the future—from article text. The system also includes visualizations that allow users to discover unexpected relationships between entities at particular points in time.

The challenge attracted extraordinary efforts from participants and intense interest from attendees, and is expected to be a permanent addition to the workshop.

3 Concluding Remarks

Synergies between IR, HCI, and other communities can help create search experiences that empower people to find, use, and share information effectively. HCIR 2010 provided a valuable opportunity for researchers to discuss ideas in an interdisciplinary setting. The feedback received from our attendees was extremely positive and we are grateful to everyone who contributed to the event’s success.

The future is bright for HCIR. The next workshop will be held at Google headquarters in Mountain View, CA in October, 2011. HCIR will continue to be a venue for early-stage research on all aspects of HCIR (from design to evaluation and deployment), but with perhaps more emphasis on systems, demonstrations, and practical implications for 2011. We will again run the HCIR Challenge, perhaps more formally and with a new dataset. We are also guest editing a special topic issue of the journal *Information Processing and Management on HCIR*, with a call for papers timed to coincide with the 2011 workshop. The aim of the special issue is to provide authors of workshop papers—and others in the community—with an opportunity to present principled advances in the area of HCIR.

We look forward to seeing you at the next workshop and receiving your journal submissions!

4 References

Marchionini, G. (2006). Toward human-computer information retrieval. *Bulletin of the American Society for Information Science*, June/July.