
WORKSHOP REPORT

HCIR 2013: The Seventh International Symposium on Human-Computer Interaction and Information Retrieval

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

Luanne Freund
University of
British Columbia
luanne.freund@ubc.ca

Catherine L. Smith
Kent State University
csm141@kent.edu

Mark D. Smucker
University of Waterloo
mark.smucker@uwaterloo.ca

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

We summarize the 2013 Symposium on Human-Computer Interaction and Information Retrieval. The two-day symposium was held October 3-4 in Vancouver, BC. The event brought together researchers and practitioners from academia and industry for in-depth discussions in an informal setting. We accepted rigorously peer-reviewed full papers that will be archived and published in the ACM Digital Library, as well as short papers describing work in progress.

1 Introduction

Research in human-computer information retrieval (HCIR) focuses on information retrieval (IR) techniques that integrate human intelligence and algorithmic search to help people explore, understand, and use information more effectively (Marchionini, 2006). The fields of human-computer interaction (HCI) and IR have developed techniques to meet these objectives, and evaluate the performance either with human participants or user simulations of various kinds. However, the insights from the development and evaluation of these systems often do not traverse disciplinary boundaries. HCIR capitalizes on synergies between the fields of IR and HCI to foster the development of search systems and associated evaluation methodologies that, among other things, support fluid and meaningful interaction, promote critical thinking, and facilitate cognitive development (White et al., 2013).

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 to bridge the two communities. The seventh HCIR symposium (HCIR 2013) took place October 3-4, 2013 in Vancouver, BC, Canada, hosted by the University of British Columbia iSchool. This year's symposium featured both short-paper and full-paper presentations and posters, with significant opportunity for discussion among attendees. As we strive to maximize the quality of the event, we tried two new things this year: (1) hosting the symposium outside of the United States for the first time, and (2) experimenting with an expanded format that spanned two days, with two keynotes, author presentations, and plenty of time for interactions and discussion. We also held a special event to remember Gene Golovchinsky, our friend and HCIR 2013 program committee chair, who passed away suddenly in August of this year.

2 Tribute to Gene Golovchinsky's Research

On August 15, 2013, HCIR co-organizer and long-time HCIR participant and supporter, Gene Golovchinsky died of cancer. Gene's death was both sad and shocking. Many in the IR community had talked with Gene only a few weeks prior to his passing at the SIGIR 2013 conference in Dublin, Ireland.

The first session of HCIR 2013 was a tribute to Gene's research career. Since earning his Ph.D. in industrial engineering from the University of Toronto in 1996, Gene had worked as a Senior Research Scientist at FX Palo Alto Laboratory. Prior to his Ph.D., and after earning his undergraduate degree in electrical engineering from UCLA, he worked for Kaiser Electronics, UCLA, and IBM (Fisher, 2013).

After an overview of Gene's career, the session comprised three talks given by researchers who had worked closely with him at different points during his research career.

The first talk was given by Rick Kopak. Rick was a graduate student at the University of Toronto at the same time as Gene, and they both worked in the lab of Gene's doctoral advisor, Mark Chignell. Rick discussed Gene's early work on the integration of dynamic hypertext and IR. Rick then played a video recording of an interview that he had conducted with Mark Chignell. In the video, Mark recounted how Gene had been an engineer at Kaiser working on a heads-up display when Gene read Don Norman's work and decided that he wanted to pursue research in human factors. Rick and Mark discussed many aspects of Gene and his research. Of note, Mark commented that Gene's work on digital annotation is an area that still holds great promise and should be picked back up by today's researchers.

Cathy Marshall, a long-time colleague and friend of Gene, gave the second talk. She focused on Gene's work on XLibris, an e-book device that was designed to go beyond the capabilities of paper. Notable aspects of XLibris were its pen-based annotation capabilities which also allowed the user to navigate among documents based on the annotations. Cathy also spoke about some little-known customizations of XLibris developed for applications in government and their work together on that project.

The third talk was given by Jeremy Pickens via recorded video. Jeremy's talk focused on his collaboration with Gene while at FXPAL and their work together on collaborative search. Jeremy discussed how working with Gene often involved Gene taking one's ideas and then improving them and giving them back. For example, while working with Gene on collaborative search, Jeremy envisioned that it could resemble a "jam session" of musicians, but initially imagined the roles as symmetric; each searcher would have access to the same tools and engage in the same processes. Gene took that idea and pushed for a notion that collaborative search should involve people with asymmetrical roles, much in the same way that musicians in a jam session do not all play the same instrument. This latter approach proved much more fruitful. Jeremy also discussed Gene's work on Querium, where almost any aspect of the system can be used as a search query. In his conclusion, Jeremy explained that working with Gene "made you a better researcher," and that it was always a shared, collaborative process.

The session concluded with comments and remembrances from the audience.

3 Keynotes

3.1 Ed Chi, Google

The keynote talk on the first day was given by Dr. Ed Chi, Staff Research Scientist at Google, and was entitled *Blurring the Boundary between Search and Recommendation*. Dr. Chi started the talk by reviewing each of the previous HCIR keynotes, providing historical context and a summary of major interest areas for HCIR research. He then established the central premise of his talk: that the boundaries between search systems and recommendation systems are blurring. Intertwined with this blurring are the tradeoffs between "push" and "pull" models of information access. Dr. Chi described how search systems are becoming more personalized and driven by question-answering capabilities, and how recommendation systems include not only search capabilities, but also features such as faceted search. He

argued that as search and recommendation systems move forward together, there are new user interface and user experience challenges that need to be addressed. Dr. Chi presented several research systems and studies to illustrate his ideas about how to meet these challenges. First, he described his research on applying information theory to social tagging systems such as Delicious, showing that such an approach could provide important insights into the creation of tags and their use in retrieval. As an example of a system that blurs the boundary between search and social recommendation systems, he presented Mr. Taggy (mrtaggy.com), a system that allows people to search using Delicious social tags. Next, Dr. Chi addressed the emerging problem of how to search and filter information on social media streams, describing his work on zerozero88, a social news recommender. In concluding this section of his talk, Dr. Chi noted that social recommendation systems encode good quality information sources, and that these sources can be used as a potential source of serendipitous information in search systems, again blurring the boundaries between search and recommendation. In the last part of his talk, Dr. Chi discussed the need for future search systems to understand search context, including social and personal factors. To illustrate these issues, he presented results from a study to examine the influence of including social annotations in search-result lists. Using gaze tracking techniques, he and colleagues found that the size and position of the social recommendation image influenced whether or not users noticed it. However, in many cases, users did not notice the social annotations, especially if the URLs and titles provided sufficient information. In concluding the talk, Dr. Chi discussed the importance of user interfaces and context, and noted that previous HCIR keynote talks had correctly predicted recent research trends in HCIR.

3.2 Cristina Conati, University of British Columbia

The keynote on the second day was entitled *Who are my users and how I can help them? The quest of user-adaptive interaction*, and was delivered by Dr. Cristina Conati, Associate Professor of Computer Science at the University of British Columbia. In the talk, Dr. Conati mapped out her work on user modeling and adaptive systems, which spans the fields of artificial intelligence, cognitive science, and human computer interaction, and which can be considered cognate to HCIR. The first part of the talk focused on open-ended learning systems and described a series of studies in which students' interactions with a learning application were logged and used to derive models of behavior for high and low learners. The models were used to identify particular behaviors associated with learning and to predict learning capabilities. Combining system logging with gaze data enabled accuracy in this task of approximately 80% after only 20% of actions with this specific system. The second part of the talk focused on user-adaptive visualizations. Dr. Conati presented studies examining the relationships between cognitive traits, tasks, and interaction behaviors when using familiar types of visualizations, such as bar charts and radial graphs. The work shows that tasks and some cognitive traits can be inferred from behavior, which creates an opportunity to offer dynamic help to users through guidance or adaptation of the visualization to suit their needs. Dr. Conati suggested that future work should focus on the specific nature and timing of help interventions. This work mirrors much of the predictive modeling work being done in HCIR and provides some valuable new perspectives on methods and approaches.

4 Papers

The symposium included a selection of full papers and short papers on a variety of topics.

4.1 Full Papers

Five full papers were accepted for presentation at the symposium as long talks. Accepted full papers will also appear in the ACM Digital Library. The following papers were accepted:

- *Are Some Tweets More Interesting Than Others? #HardQuestion*, by Omar Alonso, Catherine Marshall and Marc Najork. In this paper, the authors sought to devise a method to generate high-quality labels for Twitter postings, and apply these labels to automatically predict the interestingness of
-

tweets. The authors performed a series of studies using crowdsourcing involving a diverse set of workers with different perspectives. They explored different labeling strategies, including varying the judges, the labels they applied, the datasets, and other aspects of the labeling task. The authors then used crowdsourcing to assemble a set of tweets that were rated as interesting or not; they scored these tweets using textual and contextual features; and they used these scores as inputs to a classifier. They reported being able to achieve moderate agreement between the best classifier and the human assessments, reflecting the challenges associated with this judgment task.

- *Adaptation of a Search User Interface towards User Needs: A Prototype Study with Children and Adults*, by Tatiana Gossen, Marcus Nitsche, Jana Vos and Andreas Nuernberger. In this paper, the authors describe search interfaces that can adapt to changes in the search capabilities of their users. Search interfaces are typically optimized for general searchers and assume that search goals and cognitive skills are similar for everyone. However, especially for young and elderly users, the design requirements change quickly due to changes in users' abilities, so flexible modifications of the interface may be needed. The authors propose an evolving search user interface that adapts dynamically based on the derived capabilities of those interacting with it. They present a prototype implementing these ideas that considers the special requirements of children and is customizable to their abilities. The authors' system offers adaptation in menu type and structure, search results visualization, surrogate structure, font, audio, theme and other interface properties. They performed a user study with children and adults and describe its implications for search-system design.
 - *Sewing the Seams of Sensemaking: A Practical Interface for Tagging and Organizing Saved Search Results*, by Marti Hearst and Duane Degler. In this paper, the authors present the steps taken to develop a user-centered tool for use by analysts to save and organize their search results. In contrast to existing literature that suggests users would like to create "piles" of documents, the authors found that such designs are not efficient and thus disliked by searchers needing to perform large amounts of document triage, where time may be important. A key finding of the paper is that group creation and naming must be fast and integrated into the triage process. The solution developed by the authors allows a single keyboard letter press to save a document to a category or create a new category. For example, if "b" is pressed for a document, the document would be categorized under "b" or a new "b" category would be created. Likewise, instead of allowing full flexibility in placement of possibly overlapping piles, the authors found that grouping documents in an orderly fashion and fixed sized lowered the work needed by the users during triage. Finally, once a group of documents exists, it should be able to be used as a query to find other similar documents.
 - *Why Do Users Perceive Search Tasks As Difficult? Exploring Difficulty in Different Task Types*, by Jingjing Liu and Chang Suk Kim. In this paper, the authors report on a laboratory study to examine the reasons why searchers perceive particular search tasks to be difficult, and to examine how these reasons vary across different types of tasks. They asked 32 participants to complete four search tasks that varied across dimensions shown previously to contribute to task difficulty (product, task goal, complexity, and naming). Their results show correlation between perceived task difficulty and factors such as pre-task knowledge, previous experience, and pre-task interest, although the correlations were weak to moderate. On pre- and post-task questionnaires, participants were asked to provide reasons for their perceived task difficulty ratings. The authors reviewed these responses and created a two-level coding scheme. Frequent reasons provided by participants included: complexity, specific task requirements, low topic knowledge, uncertainty about the information need, too many choices, and insufficient information.
 - *Slow Search: Information Retrieval without Time Constraints*, by Jaime Teevan, Kevyn Collins-Thompson, Ryen W. White, Susan T. Dumais, and Yubin Kim. This paper describes slow search, a class of search where traditional speed requirements are relaxed in favor of a high quality search experience. Significant time and effort has been devoted to reducing search engine response time, which has been shown to affect searcher engagement with search results. In this paper, the authors
-

point to recent advances in the understanding of how people find information that suggest that there are scenarios where a search engine could take significantly longer to return relevant content to searchers. The authors perform large-scale log analysis and user surveys to examine how individuals value their time during search. As expected, they confirm that speed is important, but also show that there are many situations where search-result quality is also important. The paper highlights new opportunities for search systems to provide high-quality content that takes time to identify.

4.2 Short Papers

Sixteen short papers were accepted for presentation at the symposium as short talks or posters, and poster presenters participated in a “poster booster” session where they highlighted their research contributions to all attendees in short presentations. The short papers are summarized briefly below.

Four papers were given the opportunity to be presented with talks. Edie Rasmussen and Michael Huggett presented the Meta-Dex system and reported on the results of a user study of the Meta-Dex user interface, which allows people to search through the aggregated indexes of collections of digital books. The study identified common tasks and solicited perceptions of this novel type of knowledge structure. Tatiana Gossen, Michael Kotzyba, Sebastian Stober and Andreas Nuernberger presented research on children’s speech based interactions with Knowledge Journey, a search engine designed for children, conducted through a Wizard of Oz experiment. Results point to some interesting features of children’s speech-based interactions. A novel search system for a clinical trials database that employs an animated speaking agent was presented by Dina Utami, Barbara Barry, Timothy Bickmore and Michael Paasche-Orlow. The results of an evaluation show comparable effectiveness and higher satisfaction with the animated agent interface over a keyword-based search system. Clarisa Pérez-Jasso, Julio Rodríguez and Rolando Salazar’s paper reported work on providing query formulation support to users of a collaborative search system, based on a semantic knowledge sharing approach. The approach uses latent semantic analysis to suggest additional query terms to group members, taking into consideration different types of search tasks.

The following papers were presented as posters in a lively and interactive poster session. Heather O’Brien, Rafa Absar, and Helen Halbert discussed the application of the User Engagement (UE) framework in the study of mobile information interaction, arguing that UE provides a unifying structure for work in mobile user experience and a basis for theory building in the study of information interaction. Rafa and Heather also proposed a task-based framework for analysis of mobile information interaction and presented preliminary results from a study of mobile search tasks. Aiman Al-Harbi and Mark Smucker presented work in progress on uncertainty in binary relevance judgments for IR evaluation, reporting that assessors express degrees of certainty when thinking aloud while making judgments. The authors proposed that a measure of judgment certainty may be useful for understanding inter-assessor differences in relevance judgments. Erik Choi proposed the study of motivations for, and expectations of, online Q&A, with the goal of developing a conceptual framework for the contexts in which Q&A sites are used. Michael Zarro presented a preliminary model that describes how patients seek information for the day-to-day self-management of chronic illness. The model associates stages of disease self-management with changing information needs and resources. Wei Jeng, Jiepu Jiang, and Daqing He described work in progress on searchers’ perceptions of query input errors and reformulation strategies in voice search, reporting that searchers perceive three major causes of difficulty in voice search: recognition errors, topic familiarity, and topic complexity. Xiaojun Yuan, Nicholas Belkin, and Ning Sa reported results from an experiment comparing search behavior in two interactive modes: a standard keyboard interface and a spoken-language and touch interface; they found less interaction with the spoken interface, including fewer query iterations, fewer viewed documents, and fewer clicks. Chathra Hendaheewa and Chirag Shah reported preliminary results from an analysis approach that focuses on the search process as a sequence of segments or stages; the goal of the method is to detect and identify search stages in order to recommend optimal actions to improve search results. Marcus Nitsche and

Andreas Nuernberger presented the design for a search interface and filtering system that uses tags generated by a search engine. Julian Brinkley and Nasseh Tabrizi reported results from a pilot study of visually impaired users. The authors found differences in the online behavior of sighted and visually-impaired users, and that particular difficulties were found in using social media. Luanne Freund, Rick Kopak, Kristof Kessler, and Tracey Vantighem discussed the iterative, user-centered design process for the development of an e-government search engine called FRED. They described the development of the faceted interface and system functionality. Eric Meyers, Natalie Schembri, and Emily Orr discussed the design of retrieval tasks for the study of youth search practices, calling for the use of tasks that represent the day-to-day information needs and behaviors of young people.

5 Symposium-Wide Discussion

The symposium concluded with a discussion of the major themes and key ideas that arose over the two days of presentations and on future directions for HCIR. Participants emphasized the unique value of the symposium as a forum for the exchange of ideas across the HCI and IR research areas. More specifically, the importance of spoken dialog with search systems was highlighted as a particular area of interests since this was a topic covered in a number of the papers. The need for specialized interfaces for children, older adults and users with low literacy or technology skills was raised together with the observation that researchers need to move away from working with university students as study participants to incorporate more diversity in system designs and evaluation. Participants noted the need for new evaluation methods and metrics for user outcomes such as learning, synthesis, confusion, and certainty and wanted to see more of this type of work presented at the HCIR event.

Ideas for future HCIR meetings included more focus on recommendation and personalization systems, a greater emphasis on search user interface design, and the incorporation of system demos. Participants noted that the session honoring Gene Golovchinsky's research contributions added a valuable historical perspective, which allowed participants to see the trajectory of a body of work and to revisit valuable ideas relevant to current research. On that basis, participants proposed that tutorials or presentations reviewing key areas of research be included in future HCIR meetings.

Workshop participants also discussed possibilities for the structure of future HCIR events. Many expressed support for HCIR to continue to stress its unique qualities by restoring the HCIR Challenge to spotlight new innovations in a demonstration format, continuing to bring together researchers from the various subareas of IR that are typically segregated to different conferences (SIGCHI / SIGIR / ASIST), a focus on early-stage work, and its crossover with industry. Ideas for improvement were to incorporate tutorials, demonstrations, invitations to authors of notable SIGCHI and SIGIR papers to present their research at HCIR, and consider a format similar to CIDR (Conference on Innovative Data Systems Research, cidrdb.org) or HCIC (Human Computer Interaction Consortium workshop) allowing significant discussion of presented work. Participants also encouraged HCIR organizers to find sponsors who can cover all event costs, as had been done for many of the previous offerings of HCIR.

6 Concluding Remarks

HCIR 2013 was very successful. We experimented with an expanded format spanning two full days (including both short- and full-paper presentations and two keynotes), and a non-US location for the first time. We also provided extensive opportunity for discussion among attendees, which participants of previous events have informed us that they have found especially useful about the HCIR series. The feedback we received was positive and we are grateful to everyone who contributed to the event's success. Our special thanks goes to the University of British Columbia iSchool for hosting the event, and to our other sponsors: FXPAL, Microsoft Research, the GRAND network (grand-nce.ca), and ACM SIGIR, who provided student travel awards. We are extremely grateful to our keynote speakers, and the presenters of papers and posters, all of whom did an excellent job of presenting their work in a clear and accessible manner. Plans are already afoot for HCIR 2014. Look for a call for papers coming soon!

7 References

Fisher, L. (2013). Gene Golovchinsky, 1965 – 2013. *ACM News*, ACM. URL:
<http://cacm.acm.org/news/167019-gene-golovchinsky-1965-2013/fulltext>

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

White, R.W., Capra, R., Golovchinsky, G., Kules, B., Tunkelang, D., and Smith, C. (2013). Introduction to special issue on human-computer information retrieval. *Information Processing and Management*, 49(5): 1053–1057.