Designing a Tool for Exploratory Information Seeking

Gene Golovchinsky
FX Palo Alto Laboratory, Inc.
3174 Porter Drive
Palo Alto, CA 94304
gene@fxpal.com

Tony Dunnigan
FX Palo Alto Laboratory, Inc.
3174 Porter Drive
Palo Alto, CA 94304
tonyd@fxpal.com

Abdigani Diriye
University College London
London, UK
a.diriye@ucl.ac.uk

CHI'12, May 5–10, 2012, Austin, Texas, USA.
ACM 978-1-4503-1016-1/12/05.

Abstract
We describe an on-going design process in building Querium, a search system for multi-session exploratory search tasks. Querium extends a conventional search interface with a sophisticated search history interface that helps people make sense of their search activity over time. Information seeking is a cognitively demanding process that can benefit from many kinds of information. Our design process has focused on creating interfaces that facilitate on-going sense-making while keeping the interaction efficient, fluid, and enjoyable.

Author Keywords
Multi-session search, search interface, interaction design, persuasive computing.

ACM Classification Keywords
H.5.2 User Interfaces

General Terms
Human Factors

Introduction
Exploratory search describes information seeking activity whose goal is to find a thorough description of some topic. It may involve many queries, many retrieved documents, and may evolve over time. This

Design principles

Exploratory interaction style
People should be able to experiment and to backtrack without much interaction overhead.

Expressing intent
The system should allow people to express their information need ways appropriate to their current (but evolving) knowledge and interest.

Don't trust the ranking algorithm
The design should encourage deeper exploration of search results.

Data and metadata
Often the value of a document may be assessed only by actually skimming, browsing, or reading it, rather than only looking at metadata.

kind of search is poorly supported by most systems because they tend to treat each query as an isolated event. Yet in exploratory search, the queries are not independent; they form a dialog with the system through which the searcher tries to understand and evolve the information need. Thus, to support effective exploratory search, designs need to take into account the history of interaction with respect to a specific information need [7].

We set out to create a web-based platform to study people's fundamental information-seeking behaviors in a natural setting. Our system supports session-orientation and intentional collaboration [5] through up-front design rather than as retrofits. The basic set of exploratory search operations can be distilled from a large list (see [2][6]) to a more manageable set: running *ad hoc* queries, examining search results, reviewing the session history, performing relevance feedback, and following links.

Evolution of design

To set the context for our current activity, we briefly describe the evolution of our design (Figure 2 - Figure 5), highlighting the differences among the interfaces.

SACK
The SACK [6] interface was organized around a set of views in tabs, including the current result set, the set of related documents (obtained dynamically via a relevance feedback query), documents saved by the user, a query history, and a document history.

The system was effective at demonstrating the potential of integrating session history into the search process, but it suffered from several usability problems. The purpose of the tabs was confusing to users; the display was poorly formatted and did not show relevant snippets, and the relevance judgment interface was too complex for people familiar only with web search.

Querium 0.5
Based on feedback on this early design, we refactored the UI to leverage people's familiarity with existing interfaces while preserving the functionality. The tabular layout was replaced by a more conventional composition. The document view was retained, but most of the tabs were removed; only the search result and fusion tabs (formerly called document history) were retained. The query history was moved to a drop-down associated with the query box ( ); saved documents were obtained by filtering the document list.

One advantage of the filter design was the ability to manipulate the search *process* metadata in addition to the retrieved *document* metadata. Search process metadata includes information about which documents were already retrieved, seen, saved, etc. The ability to combine and filter results from multiple queries in the fusion view let people reflect on what they had accomplished, and on what remained to be done.

Problems remained with the UI, however: It was not clear how to save documents, and the spatial association between the "find by useful" button and "useful" documents was also tenuous. Also, system performance in computing the fusion results was poor due to complex queries against the metadata database used to find documents retrieved by multiple queries.

Both SACK and Querium 0.5 supported collaboration in that they the assessments of collaborating searchers

UI components

Results list

Query matches include document metadata and snippets, and may contain widgets for saving documents. Queries in a session often re-retrieve the same documents; we display a document's retrieval history in a histogram where bar height indicates degree of match to each query.

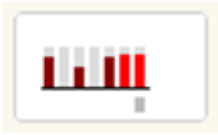

Figure 1. Histogram showing a document's revtrieval history.

Document view

Selected documents appear in a frame alongside the search results. The view may also let people comment on the displayed document.

Results assessment

People can indicate whether documents are useful or not with respect to their current task. Not useful results are suppressed in subsequent displays, although they are available with the "show all" filter setting.

when performing relevance feedback. This kind of algorithmic mediation was shown to be useful by Pickens *et al.* [10] for increasing the diversity of search results. Neither system, however, had much support for communication over the search task. SACK had none; Querium 0.5 introduced the notion of comments that could be attached to documents for others to see.

Querium 1.0

We revisited the design to introduce communication into the collaboration process, and to simplify some of the visual clutter of the results list. We also switched the collection CiteSeer [3]. This collection contains well over 1 million academic papers collected from the web, covering a range of disciplines. It also includes automatically-extracted metadata such as authors' names, titles, dates of publication, etc. CiteSeer is an important web resource for academic research, and this makes it an attractive target for understanding information seeking in a naturalistic setting. We believe, however, that these lessons can be applied more broadly.

As Figure 4 shows, the filters became a faceted search interface in the sidebar, the query history was moved to the lower part of the sidebar, and the results layout was streamlined and made even more conventional. Relevance judgments were replaced with "thumbs up" (useful) and "thumbs down" (not useful) buttons, and a facility to share (🔗) documents and queries with collaborators explicitly was introduced. A resizing control was introduced to change the relative size of the search result vs. document view to accommodate fixed-layout PDF documents.

We also introduced a summary view that offered a chronological account of people's activity in a search task. The summary view showed collaborators who ran which query, saved which document, etc., and when the actions were performed. We conducted a small deployment of this system to users outside our lab, and discovered some usability and performance issues with the system. These problems included some interface bugs, slow response times for fusion queries, some data collision issues on the server during synchronous collaboration, and out-of-date data. (The CiteSeer snapshot we had available at the time only went through 2007; we've since acquiring a much more recent snapshot.) But people's reactions to the system caused us to reflect on what we could improve not just by fixing glitches, but in reconsidering some more fundamental aspects of the interaction. These considerations led us to design Querium 2.0, the system we are currently building.

Querium 2.0: the ongoing redesign

We identified a number of technical solutions to improve the performance of the server, and reconsidered some important aspects of the user interface in an attempt to streamline the user experience. A new client architecture that implemented a different notion of the client's responsibilities with respect to interaction was introduced. The Querium 2.0 design (Figure 5) retains some of the features of earlier systems, but introduces the following changes in response to the feedback we got from our participants.

Client architecture

Whereas the older clients were merely responsible for receiving data from the server and displaying, the new architecture moves all of the session data to the client. We also adopted a more streamlined data schema for the client, rather than replicating the server schema.

Relevance feedback (RF)
RF is the process of expanding and re-evaluating an existing query based on information extracted from "useful" documents. One design issue is how to specify which documents should be used to expand the query.

Document history
The document history (also called "fusion" and "fusion query") combines ranked lists of all queries in a session to produce a single list of documents retrieved in a session.

Query history
To facilitate sense-making, particularly in a collaborative setting [9], this view shows the queries that have been run, along with some metadata (who ran it, when, how many documents were retrieved, etc.). Most implementations also add color-coded representations of the assessments of retrieved documents.

These simplifications made it possible to perform client-side computation much faster than equivalent computation would be on the server.

Views and sub-views
The summary view is split into two sub-views, one for queries and one for documents. Rather than representing an explicit chronology of actions, both sub-views can be sorted by the time the queries or documents were last visited. The document summary sub-view reverted to the document history view of the SACK and Querium 0.5, with the important difference that fusion is no longer calculated on the server, but rather on the client. Whereas for the server this was an expensive database operation, client has all the relevant information in memory, making it nearly instantaneous.

Queries
Querium 1.0 had four different kinds of queries (keyword, like-based relevance feedback, *ad hoc* relevance feedback, and fusion), but we found that this complexity was hard for people to understand. Even though relevance feedback links appeared in response to user actions, this interaction seemed to confuse people. We therefore changed the query interface to a single, integrated design. Keyword queries are entered as before, and relevance feedback documents are shown below the query box.

Engagement, persuasion, and emotional design
Exploratory search interfaces have to support a more complex set of interactions compared with the familiar web-based search engine interfaces. One challenge is to design useful interactions for exploration; another is to get people to see their and to get them to adopt new ways of thinking about information seeking interfaces.

We have applied an emotional design [8] approach to the arrangement, ornamentation and animation of the UI. Our goal is to provide users with an orchestrated understanding of how this system works so that they might interact with it more effectively. We've tried to incorporate elements that are amusing, interesting or otherwise aesthetically pleasing to create a more complete understanding of the application.

We take a two-pronged approach, tackling engagement and persuasion separately [4]. To improve engagement, we restructured view to minimize scrolling, added animations to reinforce people's actions, and added a colorful background image. The overall effect is to make the interface feel more like an application than a web page.

We have also started investigating persuasive aspects of the interface. The goal is to gradually reinforce searchers' behaviors that lead to productive outcomes. There is a potentially rich and under-investigated space of which behaviors to reward and of the nature of rewards. We are interested in fostering exploration and assessment: we would like to get people to run diverse queries because they are useful for understanding topic breadth, and we would also like people to make many judgments of usefulness to drive relevance feedback. Our strategy is to use ambient displays of information to minimize disruptions of the primary task [1].

To encourage query diversity, our design animates the background image to reflect the contribution of the current query to the set of documents retrieved in the session. The more new documents a query retrieves, the clearer and more pleasant the background image becomes. Conversely, when a query retrieves mostly

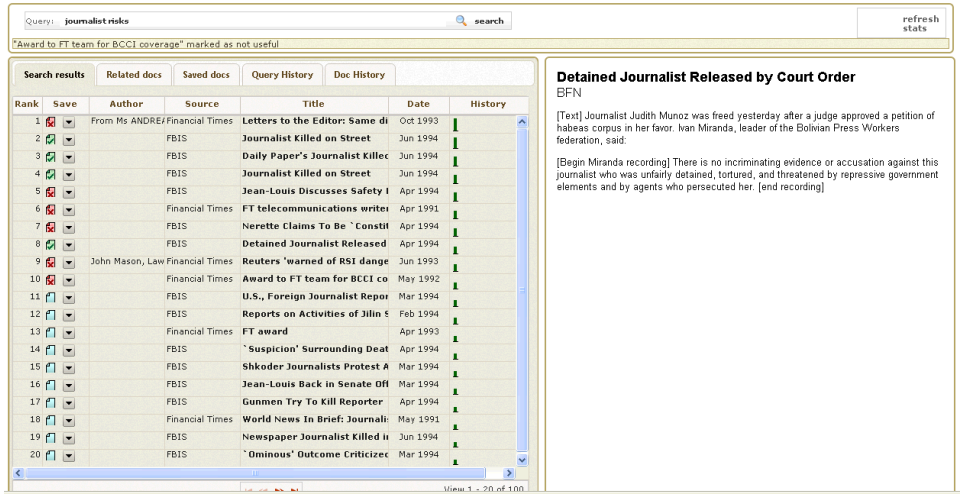

Figure 2. SACK Query Results view

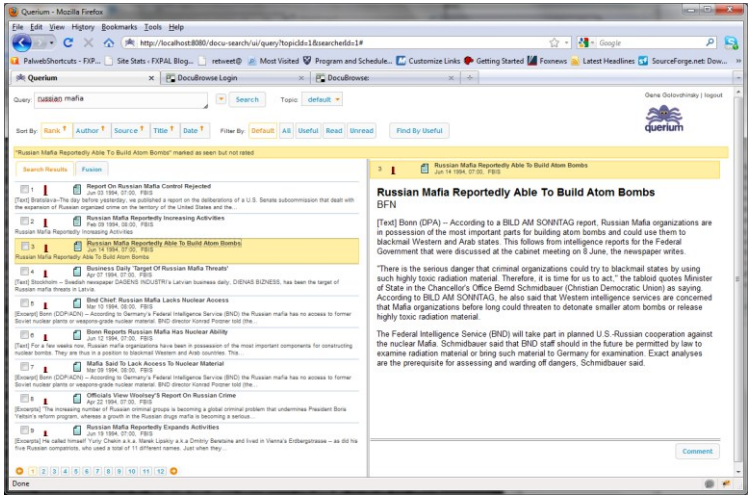

Figure 3. Querium 0.5 results view

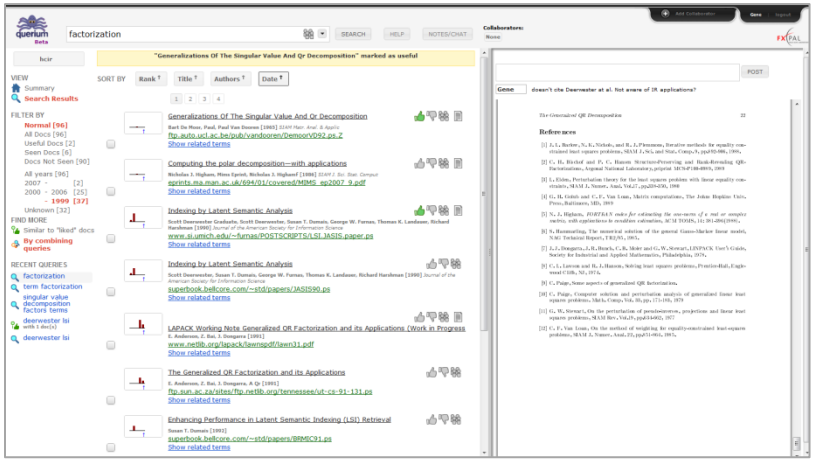

Figure 4. Querium 1.0 results view

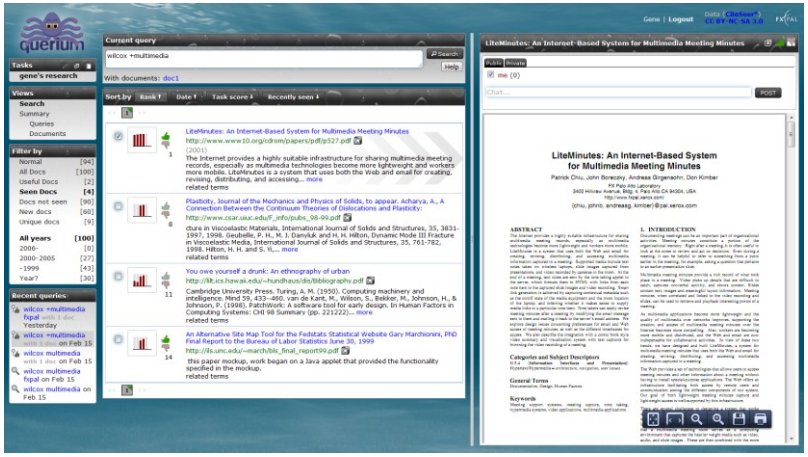

Figure 5. Querium 2.0 results view

the same documents as have already been found, the background gets darker and murkier. This change happens over a few seconds after the new results are shown, and represents our initial exploration of persuasive computing. The actual number of new documents is indicated in the filter.

Marking documents as "useful" or "not useful" is an extension of the familiar bookmarking behavior that is useful for filtering of results, for communicating with collaborators, and as a source for relevance feedback. To encourage people to mark documents, we are introducing subtle dynamic highlighting of the associated controls if they have not been used recently.

The exact space of actions and schedule of rewards needs to be discovered empirically, but the idea that playful behavior for a serious purpose can improve people's attitudes toward a system and can be used to foster desirable behaviors seems interesting and promising to us.

Conclusion

Exploratory search is a rich and challenging environment for interface design. We have investigated several aspects of the space, including how best to structure cognitively demanding aspects of the interaction, the effectiveness of various means of expressing information need, how to help people make sense of their ongoing search activities, etc. We are about to deploy a new version of the system that implements a simplified design while retaining the required functionality and introduces new persuasive features. The deployment should provide us with ample fodder for the next iteration of this design.

References

[1] Adamczyk, P.D., and Bailey, B.P. (2004) If not now, when?: the effects of interruption at different moments within task execution. In *Proc. CHI '04*. ACM, 271-278.

[2] Bates, M. (1989) The Design of Browsing and Berrypicking Techniques for the Online Search Interface. *Online Review*, vol. 13 no. 5, 407-24.

[3] CiteSeer. Available online at http://citeseerx.ist.psu.edu/

[4] Fogg, B.J., Cuellar, G., Danielson, D. (2009) Motivating, Influencing, and Persuading Users. In *Human-computer Interaction*, Sears, A. and Jacko, J., Eds. CRC Press, 133-146.

[5] Golovchinsky, G., Qvarfordt, P., and Pickens, J. (2009) Collaborative Information Seeking. *Computer*, 42(3), IEEE Press, 47-51.

[6] Golovchinsky, G., and Pickens, J. (2010) Interactive Information Seeking via Selective Application of Contextual Knowledge. In *Proc. IIiX 2010*. ACM Press, 145-154.

[7] Morris, M.R. and Horvitz, E. (2007) S3: Storable, Shareable Search. In *Proc. Interact 2007*, pp. 120-123.

[8] Norman, D. (2003) *Emotional Design: Why We Love (or Hate) Everyday Things*, Basic Books.

[9] Paul, S. and Morris, M.R. CoSense: (2009) Enhancing Sensemaking for Collaborative Web Search. In *Proc. CHI 2009*, ACM, 1771-1780.

[10] Pickens, J., Golovchinsky, G., Shah, C., Qvarfordt, P., and Back, M. (2008) Algorithmic Mediation for Collaborative Exploratory Search. In Proc. *SIGIR 2008*. ACM, 315-322.