Towards "Natural" Interactions in Search User Interfaces

Marti A. Hearst
UC Berkeley
102 South Hall
Berkeley, CA 94720
hearst@ischool.berkeley.edu

Keywords: *Search User Interfaces, Human Computer Interaction, Artificial Intelligence, Speech Recognition, Social Search, Video Search, Crowdsourcing*

Key Trends in Technology Use

What does the future hold for search user interfaces? The familiar web search interface of today works well for millions of people and billions of queries a year. Very few innovations in search interfaces gain wide enough acceptance to replace the standard type-keywords-in-entry-form, view-results-in-a-vertical-results-list interface. This is in part because search is a means toward another end, and reading text is a mentally demanding task; the fewer distractions while reading, the more usable the interface. Additionaly, search, like email, is used by nearly everyone who uses the Web, and so its features and functions must be understandable to an enormous and diverse range of people [14].

Therefore, future trends in search interfaces will most likely reflect trends in the use of information technology generally. And today there is a notable trend towards more "natural" user interfaces – pointing with fingers rather than mouses, speaking rather than typing, viewing videos rather than reading text, writing full sentences rather than artificial keywords. (The term "natural interface" has been promoted by researchers at Microsoft, among others.) It is not surprising that people are drawn to interfaces that allow them to think and move in a manner similar to that of their non-computing lives, but only recently has technology been able to support natural interaction.

There is also a trend towards social rather than solo usage of information technology. And importantly, these multi-person interactions are often recorded, stored, and indexed for later viewing. Again, many people would have preferred non-isolated computer usage from the start, but technology (and user interface design) did not support it well until recently.

Technology is also advancing in the integration of massive quantities of user behavior and large scale human-generated knowledge bases. Search today already benefits from the tracking of search behavior over hundreds of millions of queries to improve ranking, offer accurate spelling suggestions, auto-suggest query terms in real time while the user is typing, and suggest concepts related to a query. Integration with databases and more sophisticated processing place search at the cusp of being able to support smarter, data-driven, focused interfaces for advanced search needs.

These trends are, or will be, interweaving in various ways, which will have some interesting ramifications for search interfaces, and suggest promising directions for research.

Speech Input

Speech-based user interfaces generally, and speech for search input in particular, are likely to gain a much stronger presence in the coming few years.

At least three trends support this: First, the prevalence of phone-based mobile devices provide a natural way to capture speech, since phones are used in large part for spoken conversations.

Second, the technology for speech recognition, after years of only incremental progress, is improving by leaps and bounds, thanks to huge data repositories being generated by mobile phone usage. (To gather a large training set of spoken corrected data for its speech recognition system, Google hosted a free 411 information service for phones [28].)

Third, touch screen interfaces are becoming increasingly popular, especially when paired with mobile devices. Neither small devices nor touch screens lend themselves well to typing, and so make spoken input more attractive. (Although clever finger-swiping based input methods, such as ShapeWriter for entering text [39] and Gesture Search for menu navigation [19], provide compelling alternatives to typing.)

These trends suggest that usage of voice-activated queries and commands are likely to increase rapidly in the next few years as response time and accuracy continue to improve.

The next likely development following on voice-based input is a dialogue-like give and take. Although not a reality yet, recent advances are bringing closer the dream of an intelligent interactive agent. The Siri system provides an inter-

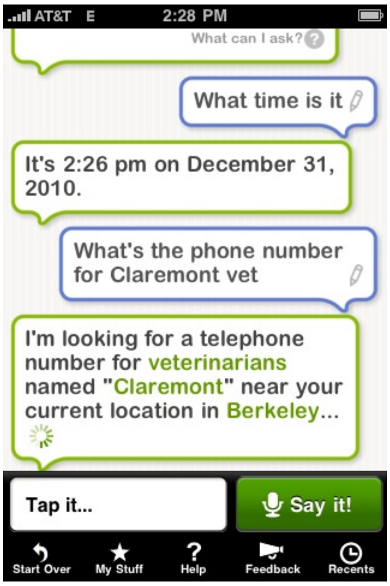

Figure 1: The Siri interface, which accepts speech as input and attempts to support a dialogue. In the first phase, a query for a phone number initially shows a message conveying what the system's understanding of the question is while it does a search.

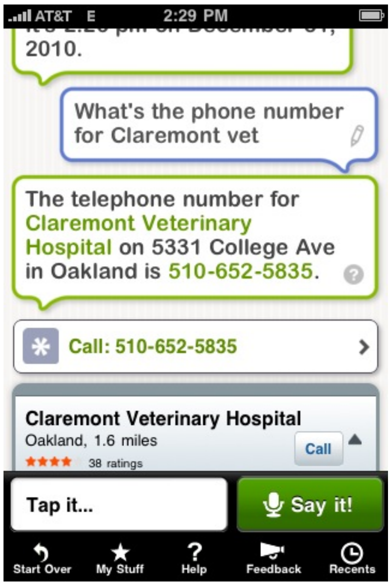

Figure 2: In the next part of the dialogue, the system has replaced the confirmation bubble with the answer. This is followed by a command to make the phone call. Note the use of graphics alongside text to enrich the output.

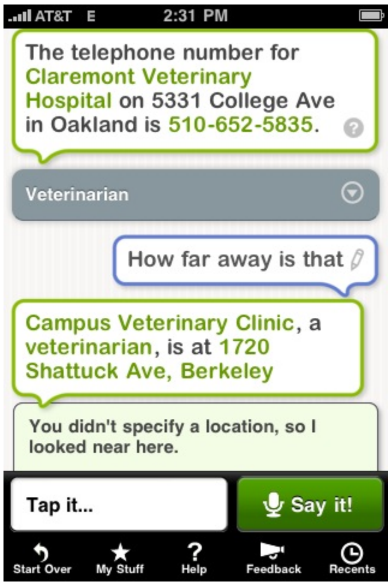

Figure 3: In the next step, the system partly understands the intention of the question (find the location, not find the distance) but does not correctly infer what the target location is.

face that combines local information, speech recognition, a well designed interface that allows easy editing of the voice recognition, and visual display of search results. Siri, which was recently bought by Apple, originated from a DARPA research project called CALO which involved hundreds of computer science researchers developing machine learning, reasoning, knowledge bases, and other technology to create an intelligent personal assistant [6, 35]. Although the ability to accurately follow up one request with the next is still limited, good interface design helps bridge the gap in the back end, since the user can see alternatives and make corrections. (See Figures 1 – 3.)

Note that Siri also attempts to makes use of the user's contextual information, such as their current location. An enormous amount of research (such as [8] and [20]) and commercial development focuses on the use of time, location, and other contextual cues for search and other applications, and this will continue to increase in importance, especially for mobile platforms.

There are of course drawbacks with the use of voice input, the most significant being that speaking makes noise, and can disturb people around the speaker. An exciting research advance would be the invention of a microphone that can successfully uptake the words that the speaker is saying but somehow prevent those around the speaker from hearing those words, like a science fiction "cone of silence." Such an invention would have wide-ranging utility for mobile phones.

Social Search: Collaborative Search

Although observational studies show that people often search collaboratively, only recently have tools been developed that

Figure 4: Two people collaborating in a video search task using the technology described in Pickens et al. Each user views a different unique interface along with a shared view; the results of the work of one person change the rank ordering of what is seen by the other.

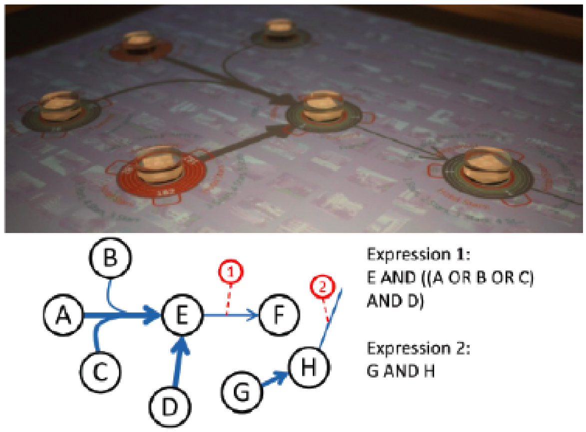

Figure 5: A table-based collaborative search formulation tool making us of a large table display, physical input devices and visualization. From Jetter et al.

explicitly support people searching together. This reflects a larger research renaissance in tools for supporting real-time shared activity, such as shared online whiteboards and document editing tools.

One exciting development in collaborative search is work by Pickens et al. [29, 7] that assumes that the ranking algorithm should allow people to work at their own pace, but be influenced in real-time by the different teammates' search activities. The different searchers should not step on one another's proverbial toes: if one searcher decides to issue a new query, others should not be interrupted in their trains of thought.

Pickens et al. address this issue by developing an algorithm that combines multiple rounds of queries from multiple searchers during a single search session. They use two criteria for weighting results where both are functions of the ranked list of documents returned for a given query. The first variable is "freshness," which is higher for documents not yet viewed. "Relevance" is higher for documents that closely match the query. The two factors are combined to complement one another, and are continuously updated based on new queries and searcher-specified relevance judgements.

In addition, Pickens et al. assign different roles to the members of a team. One role, the Prospector, is in charge of issuing new queries to explore new parts of the information space, while the other acts as Miner, whose job is to look at the retrieved results to determine which are relevant. Documents that have not yet been looked at are queued up for the Miner interface according to the freshness/relevance weighted scores. The Prospector is shown new query term suggestions based on how they differ from queries already issued, and on the relevance judgements being made by the Miner. Each role has its own distinct interface, but a third computer screen is used to show continually-updating information about the queries that have been issued, the documents that have been marked as relevant, and system-suggested query terms based on the actions of both users.

Another approach to supporting real-time search collaboration by Jetter et al. [16] makes use of a large work surface and input devices that combine physical manual manipulation with virtual markings. The interface was evaluated on a complex collaborative search task, that of a group of people selecting a product given that each member of the group has different preferences which act as constraints (e.g., when choosing a hotel, one person needs a heated pool, another wants it to receive at least 4 stars of recommendation, and a third person wants the price below a certain level). Jetter et al.'s solution uses a combination of faceted navigation [37] and filter-flow visualization [38] that shows how many constraints are met by a set of items given which constraints have been set. The visualization is displayed on a large workspace, where the controls are manipulated using physical selectors (see Figure 5). Collaboration is facilitated by allowing each user to work privately on a corner of the workspace and then let the results from their piece of the query flow into the rest of the group's query specification. A careful usability study found that this approach produced results as good as those using a standard web-based faceted navigation interface, but with more bonhomie among the collaborators.

Social Search: Asking Other People

Research suggests that much of online interaction on social sites is primarily for the social experience of the interaction, rather than for problem-centric information seeking [13]. Reflecting this, a study by Morris et al. [24] found that the kinds of things asked of others via social networks are not necessarily the kinds of information that can be found in a static web page. Survey respondents were asked to supply questions that they had posed to their social networks on Twitter and Facebook. The researchers classified the 249 examples manually, finding that only 17% of the queries were for factual information that one would typically seek from web pages (e.g., how to put an Excel file into LaTex). The most common categories of questions were requests for recommendations (29%), opinions (22%), rhetorical questions (14%), requests for others to join social events (9%), for favors (4%), or for social connections, such as job openings (3%) and offers of various kinds (1%).

A study of the Aardvark expert social question answering system found similar results: 65% of a random sample of 1000 queries had a subjective aspect to them [15]. The new social question answering site Quora also asks questions that are often subjective and opinion-based; for instance, the question "What does Dustin Moskovitz think of the new Facebook movie?" was answered by the subject of the question himself.

At this point it is unclear what the best user interfaces are for representing this more social kind of search. Freyne et al. [12] conducted a small study in which different kinds of social cues were shown via icons next to search results listings. Subjective results showed a positive preference towards cues that showed which articles had been frequently read by or annotated by others. Yahoo experimented with the MyWeb system in which search results were augmented with an avatar of the person in the user's social network who had recommended the page, along with that recommendation. Google recently introduced a social search tool called "+1" with a similar interface. It is likely that we will see significant user interface innovations in this space.

When using a social network to try to gain information, especially in a work situation, research is ongoing about how best to distribute questions among experts, either within an organization or across the Internet generally [18, 21]. Recent work by Richardson and White [34] deployed and studied an instant-messaging-based question answering service that matched the asker's questions against pre-defined profiles of over 2000 potential answerer's expertise, based on their availability. Three experts were contacted at a time, in descending order of how well their profile matched the content of the question. If an offer to answer was not received within a fixed set of time, the request was sent out to a wider circle of experts. If an answerer accepted a request, the other outstanding requests were cancelled. The tool then mediated the conversation between questioner and answerer, and asked the questioner to rate their satisfaction with the answer.

Richardson and White examined log data for this system to make predictions to form an interruption cost model, including how many people to send a question to in order to minimize disruption while maximizing likelihood of receiving an informed answer, whether a question will be answered, and how well the asker will be satisfied by the answer received.

As expert solicitation systems become more sophisticated about targeting people with the right expertise and who are in the right state of mind to address a request, they are likely to become a fixture in knowledge-centric workplaces as well as in various volunteer causes, such as the Peer2Patent project for community input of patent prior art [26].

Social Search: Crowdsourcing

The term collaboration as used above refers to a set of people working together closely, usually synchronously, to achieve a goal. By contrast, "crowdsourcing" refers to large groups of people who are not necessarily working together knowingly, but rather who each contribute in small ways, leading to a larger greater whole, as seen, for example, in Wikipedia.

Crowdsourcing in the realm of information seeking is seen in web sites in which communities curate and rate information and share this information with others. These include question answering websites as mentioned above, as well as product reviewing sites, bookmark sharing websites like Delicious and news ranking and aggregation websites like Digg. The more explicitly networked social tools like Twitter and Facebook also act as real-time socially targeted information sources.

There have been a number attempts to use explicit user input to improve search results ranking, but most have not survived. For instance, Google's SearchWiki, which allowed users to comment on web sites and see other people's comments in the search results page, is no longer supported. The Blekko web search engine is a new attempt to use sophisticated algorithms combined with community curation to improve results rankings. (Its founder also started the Open Directory Project, a crowdsourced yellow pages for the web.) With Blekko, users can create "vertical" (subject specific) search by labeling web pages with a category label preceded by a slash. Users can also mark pages as spam. These two operations together impose crowdsourced quality control over retrieved web pages. The search engine also has a social aspect in that it allows the user to see if their friends have marked particular pages with a "/like" slashtag. It remains to be seen if explicit crowdsourcing will scale for search results ranking.

Crowdsourcing usually refers to people explicitly contributing to an effort, but it can be argued that web search engines have been using a form of *implicit* crowdsourcing for years, by modifying ranking algorithms based on huge quantities of user clickthrough data [17] or predicting which vertical subject area (e.g., music, news, travel) to augment a query with [9]. Richer user behavior data such as mouse movements, page dwell time, and in one line of work, tracking searchers' click path many steps from the search results page, even across domains to their destination page, has been used to yield useful suggestions of pages that are not related to the original page by a close keyword match [36].

Natural Language-like Queries

Although keyword querying is still standard on the Web, savvy users have been typing in more detailed queries for years, and Web search engines have greatly improved their ability to handle long queries. Research shows that people prefer natural expression of queries over keywords [30, 5] and web search engine query length continues to increase each year. According to Experian Hitwise [22], when comparing queries over a four-week period in August-September 2010 to the same four-week period in 2009, searches of 5 - 8 words were up 10% while 1 - 4 word searches were down 2%. Spoken queries are also likely to be full sentences once speech recognition becomes faster and more accurate.

Longer queries are now being helped by the online existence of more colloquial language. When most of the content is technical or scientific – as was more characteristic of the early Web – there is less likely to be an easy to find match between a lay user's words and the terms used in the informative documents. The popularity of question answering sites (such as Answers.com, Quora, and Yahoo Answers), which store user-generated content, bridges colloquial and formal language directly in relevant documents.

For example, if a searcher needs a device to connect both a Wii and a DVD player to a TV, but does not know what that device is called, keyword queries may very well fail. But a query on "how do I connect wii and dvd to my tv" turns up a nearly perfect match on a question answering site, with the solution being an item called either a video selector or a 2-way A/V switcher. The point of this example is that although the searcher does not have the vocabulary to look up what is needed, they do have the same vocabulary as other people who were in the same cognitive situation. The combination of text worded colloquially and search engines that do a good job with sentence-length queries helps reduce the vocabulary problem. There has been considerable work on how to effectively search question answering sites [3, 4]; ranking algorithms that increasingly make use of these mappings will continue to improve results for difficult queries.

Another technical development that may help users who issue long queries is the increasing number of systems that support quasi-natural language interfaces, or "sloppy commands." The new syntax is tolerant of variations, relatively robust, and "exhibit slight touches of natural language flexibility" [25]. These interfaces are seen in Web search engines that support a variety of wordings for certain kinds of questions that will retrieve answers from a database or a pre-defined program ("Istanbul time", What is the time in Istanbul?", "what time is it? Istanbul".) Blekko allows query modification by a simple slash notation to refine results to pre-defined categories, e.g., "istanbul /tech" for search results about technology and "istanbul /people" for results that have been labeled relevant to people.

Miller et al. [23] have developed tools for "sloppy commands" meaning the user has a lot of flexibility in how the express the command, and so memorization is not required to make use of them. The "linguistic command line" of Enso (later Ubiquity) [33, 10] experimented with leniency in operating system command lines. The Quicksilver application for Apple operating systems supports a hybrid command/GUI interface, using continuous feedback to whittle down the available choices to include what the user has typed so far that still matches available commands.

The Wolfram Alpha search engine provides a wide range of pre-defined query types that mix structured forms with some flexibility in word order and a knowledge-base and computational back end that knows how to handle certain combinations of these. For instance, a query on "2 slices of pizza with pepperoni" is decomposed into the base information need (information about pizza) refined by units (slices), the quantity (2) and modifications of the baseline concept (with pepperoni). The result is a table that shows calorie and nutrition information. However, the interpretive range is limited; a query on "recipe for pizza with pepperoni" returns the same measurement information as "pizza with pepperoni."

This hybrid of improved language analysis and command languages, making use of structured knowledge bases paired with clever uses of interaction, may well lead to more intelligent interfaces and an expansion of dialogue-like interaction (as discussed in the Siri system above). The IBM Watson project which recently beat top human champions in the game of *Jeopardy!* also employs massive language analysis, knowledge base analysis, and speech recognition, and will likely set the stage for future highly advanced natural language question answering systems [11].

The Growing Importance of Video Content

There is increasing evidence of a preference among ordinary information consumers for video and audio content over textual content. Movies have already replaced books as cultural touchstones in the U.S. A recent report by Pew Research included a quote from a media executive who said that emails containing podcasts were opened 20% more often than standard marketing email [32]. Also according to Pew, 52% of U.S. adults have watched online videos, with 7 in 10 Internet users saying they do [31]. According to Hitwise, the YouTube video sharing site was the fifth most visited website in 2010 [2] and Comscore reported in March 2010 that YouTube had a higher search volume than Yahoo or Bing [1].

Video communication is taking on some of the trappings of textual communication. For instance, YouTube supports the notion of a video "reply." And when video questions were accepted for the 2008 U.S. presidential primary debates, most of the citizen-submitted videos that were selected by the moderators consisted simply of people pointing the camera at themselves and speaking their question out loud, with a backdrop consisting of a wall in a room of their home. There were few visual flourishes; the video did not add much beyond what a questioner in a live audience would have conveyed. Video is becoming a mundane way to communicate.

Mobile devices now make it easier to capture video, thus increasing the likelihood of this becoming an ever more important form of communication. According to Pew, almost a fifth of American adults have tried video calling either on their phones or computers, and 23% of U.S. Internet users have used a video chat service such as Skype. Further, 14% of Internet users have created and uploaded videos [31].

Without doubt the technology to support full video usage is significantly lacking behind that of text, but we can surmise that some handy inventions are not far off. Better tools for quick edits are likely to be seen soon, as they have been for image processing (a currently popular mobile app called Instagram allows people to snap a photo with their phone, quickly and easily apply filters to produce an "artsy" look, and then immediately share the image with a social network. Instagram achieved 1 million users within two months of its introduction.)

Still lacking are truly useful tools for quickly and cogently skimming the content of video, summarizing it in a meaningful way, and more to the point, searching within and across it, although research is active in this area [27]. YouTube has introduced tools that automatically provide textual closed captioning over spoken language, which can also be used for search, as has a startup company called SpeakerText. Faceted navigation [37] has proven to be the method of choice for information access within image collections; perhaps the same will be true of video collections. That said, serious breakthroughs are needed for both image and video content search before their performance can rival that of text search.

The time constraints imposed by YouTube have resulted in a culture of short videos, characterized by focus topics, which makes title search more effective than it would be if most online videos were longer in duration. For instance, the excellent educational video courses of the Khan Academy each are less than 10 minutes long and their subject matter is easily browsable by title (e.g., "Circles: Diameters, Radius, and Circumference," or "Distributive Property of Matrix Products"). But just as search over collections of books is still not particularly sophisticated, search over movie-length videos may well prove problematic and require alternative approaches.

Summary

The future for user interfaces generally is in those that support more "natural" human interactions, be they gesturing with the fingers, speaking rather than typing, watching video rather than reading, or using information technology socially rather than alone. This article has argued that these trends will also affect user interfaces for search, and has highlighted some of the recent work that reflects these trends. Utilizing advanced processing techniques over huge sets of behavioral data, future search interfaces will better support means of finding other people to answer questions or provide opinions, more natural dialogue-like interactions, and finding information expressed as non-textual content, using non-textual means of input. More natural modes of interaction have long been goals of interface design but recent developments have brought them closer to realization.

1. REFERENCES

[1] Comscore releases march 2010 u.s. search engine rankings, March 2010. http://www.comscore.com/Press_Events/Press_Releases/2010/4/com-Score_Releases_March_2010_U.S._Search_Engine_Rankings.

[2] Facebook was the top search term in 2010 for second straight year, Dec 29 2010. http://www.hitwise.com/us/press-center/press-releases/facebook-was-the-top-search-term-in-2010-for-sec/.

[3] L.A. Adamic, J. Zhang, E. Bakshy, and M.S. Ackerman. Knowledge sharing and yahoo answers: everyone knows something. In *Proceeding of the 17th international conference on World Wide Web*, pages 665–674. ACM, 2008.

[4] J. Bian, Y. Liu, E. Agichtein, and H. Zha. Finding the right facts in the crowd: factoid question answering over social media. In *Proceeding of the 17th international conference on World Wide Web*, pages 467–476. ACM, 2008.

[5] D. Bilal. Children's use of the Yahooligans! Web Search Engine: I. Cognitive, Physical, and Affective Behaviors on Fact-Based. *Journal of the American Society of Information Science*, 51:646–65, 2000.

[6] V.K. Chaudhri, A. Cheyer, R. Guili, B. Jarrold, K.L. Myers, and J. Niekrasz. A case study in engineering a knowledge base for an intelligent personal assistant. In *the Proc. of the 2006 Semantic Desktop Workshop, Athens, GA*. Citeseer, 2006.

[7] Gene Golovchinsky Chirag Shah, Jeremy Pickens. Role-based results redistribution for collaborative information retrieval. *Information Processing and Management*, 46(6):773–781, 2010.

[8] K. Church, J. Neumann, M. Cherubini, and N. Oliver. The Map Trap?: an evaluation of map versus text-based interfaces for location-based mobile search services. In *Proceedings of the 19th international conference on World wide web*, pages 261–270. ACM, 2010.

[9] F. Diaz and J. Arguello. Adaptation of offline vertical selection predictions in the presence of user feedback. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, pages 323–330. ACM, 2009.

[10] M.Y. Erlewine. Ubiquity: Designing a Multilingual Natural Language Interface. In *SIGIR Workshop on Information Access in a Multilingual World*, July 2009.

[11] D. Ferrucci, E. Brown, J. Chu-Carroll, J. Fan, D. Gondek, A.A. Kalyanpur, A. Lally, J.W. Murdock, E. Nyberg, J. Prager, et al. Building Watson: An Overview of the DeepQA Project. *AI Magazine*, 31(3):59, 2010.

[12] J. Freyne, R. Farzan, P. Brusilovsky, B. Smyth, and M. Coyle. Collecting community wisdom: integrating social search & social navigation. In *Proceedings of the 12th international conference on Intelligent user interfaces*, pages 52–61. ACM, 2007.

[13] F.M. Harper, D. Moy, and J.A. Konstan. Facts or friends?: distinguishing informational and conversational questions in social Q&A sites. In *Proceedings of the 27th international conference on Human factors in computing systems*, pages 759–768. ACM, 2009.

[14] M. Hearst. *Search user interfaces*. Cambridge Univ Pr, 2009.

[15] D. Horowitz and S.D. Kamvar. The anatomy of a large-scale social search engine. In *Proceedings of the 19th international conference on World wide web*, pages 431–440. ACM, 2010.

[16] H-C. Jetter, J. Gerken, M. Zöllner, H. Reiterer, and N. Milic-Frayling. Materializing the Query with Facet-Streams – A Hybrid Surface for Collaborative Search on Tabletops. In *Proceedings of the 29th international conference on Human factors in computing systems*. ACM, 2011. To appear.

[17] T. Joachims, L. Granka, B. Pan, H. Hembrooke, and G. Gay. Accurately interpreting clickthrough data as implicit feedback. In *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 154–161. ACM, 2005.

[18] H. Kautz, B. Selman, and M. Shah. Referral Web:

combining social networks and collaborative filtering. *Communications of the ACM*, 40(3):63–65, 1997.

[19] Y. Li. Gesture search: a tool for fast mobile data access. In *Proceedings of the 23nd annual ACM symposium on User interface software and technology*, pages 87–96. ACM, 2010.

[20] C. Liu, P.L.P. Rau, and F. Gao. Mobile information search for location-based information. *Computers in Industry*, 61(4):364–371, 2010.

[21] Y. Liu and E. Agichtein. You've got answers: towards personalized models for predicting success in community question answering. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics on Human Language Technologies: Short Papers*, pages 97–100. Association for Computational Linguistics, 2008.

[22] Matt McGee. The long tail is alive and well, September 16 2010. http://www.smallbusinesssem.com/long-tail-alive-well/3659/ and http://twitter.com/Hitwise_US/status/24041444164.

[23] R.C. Miller, V.H. Chou, M. Bernstein, G. Little, M. Van Kleek, and D. Karger. Inky: a sloppy command line for the web with rich visual feedback. In *Proceedings of the 21st annual ACM symposium on User interface software and technology*, pages 131–140. ACM, 2008.

[24] M.R. Morris, J. Teevan, and K. Panovich. What do people ask their social networks, and why?: a survey study of status message q&a behavior. In *Proceedings of the 28th international conference on Human factors in computing systems*, pages 1739–1748. ACM, 2010.

[25] D. Norman. The next UI breakthrough: command lines. *interactions*, 14(3):44–45, 2007.

[26] B.S. Noveck. Peer to Patent: Collective Intelligence, Open Review, and Patent Reform. *Harvard Journal of Law & Technology*, 20(1):123–162, 2006.

[27] P. Over, G. Awad, J. Fiscus, B. Antonishek, and M. Michel. TRECVID 2010 – An introuction to the goals, tasks, data, evaluation mechanisms and metrics. 2010.

[28] Juan Carlos Peres. Google wants your phonemes. *InfoWorld*, October 23 2007. http://www.infoworld.com/t/data-management/google-wants-your-phonemes-539.

[29] J. Pickens, G. Golovchinsky, C. Shah, P. Qvarfordt, and M. Back. Algorithmic mediation for collaborative exploratory search. In *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 315–322. ACM, 2008.

[30] A. Pollock and A. Hockley. What's Wrong with Internet Searching. *D-Lib Magazine*, 1997. www.dlib.org.

[31] K. Purcell. The State of Online Video, June 3 2010. http://www.pewinternet.org/ /media//Files/Reports/2010/PIP-The-State-of-Online-Video.pdf.

[32] L. Ranie. Digital 'Natives' Invade the Workplace, September 28 2006. http://pewresearch.org/pubs/70/digital-natives-invade-the-workplace.

[33] A. Raskin. The linguistic command line. *interactions*, 15(1):19–22, 2008.

[34] M. Richardson and R. White. Supporting Synchronous Social Q&A Throughout the Question Lifecycle. In *Proceedings of WWW 2011*, 2011. To appear.

[35] Wade Roush. The story of siri, from birth at sri to acquisition by apple – virtual personal assistants go mobile. *Xconomy*, June 2010. http://www.xconomy.com/san-francisco/2010/06/14/the-story-of-siri-from-birth-at-sri-to-acquisition-by-apple-virtual-personal-assistants-go-mobile/.

[36] R.W. White, M. Bilenko, and S. Cucerzan. Studying the use of popular destinations to enhance web search interaction. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 159–166. ACM, 2007.

[37] K.P. Yee, K. Swearingen, K. Li, and M. Hearst. Faceted metadata for image search and browsing. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 401–408. ACM, 2003.

[38] D. Young and B. Shneiderman. A graphical filter/flow representation of Boolean queries: a prototype implementation and evaluation. *Journal of the American Society for Information Science*, 44(6):327–339, 1993.

[39] S. Zhai, P.O. Kristensson, P. Gong, M. Greiner, S.A. Peng, L.M. Liu, and A. Dunnigan. Shapewriter on the iPhone: from the laboratory to the real world. In *Proceedings of the 27th international conference extended abstracts on Human factors in computing systems*, pages 2667–2670. ACM, 2009.