

DOI:10.1145/2018396.2018414

Users will speak rather than type, watch video rather than read, and use technology socially rather than alone.

BY MARTI A. HEARST

‘Natural’ Search User Interfaces

WHAT DOES THE future hold for search interfaces for users? Today’s familiar Web search interface works well for tens of millions of people and billions of queries a year, but 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 partly 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. Additionally, search, like email, is used by nearly everyone using the Web, so its features and functions must be understandable to an enormous and diverse population.¹³

Future trends in search interfaces will most likely reflect trends in the use of IT generally. Today, there is a notable trend toward more “natural” user interfaces: pointing with fingers rather than mice, speaking rather than typing, viewing videos rather than reading text, and writing full sentences rather than artificial

keywords. (The term “natural interface” is promoted by researchers at Microsoft, among others.) Not surprisingly, people are drawn to interfaces that allow them to think and move in a manner like what they use in their non-computing lives, but only recently has technology been able to support it.

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

Technology is advancing toward integration of massive quantities of user behavior and large-scale human-generated knowledge bases. Search today 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 as the user types, 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.

These trends are, or will be, interweaving in various ways, with interesting ramifications for search interfaces

» key insights

- “Natural” modes of interaction are starting to be commonplace in hardware and software tools, influencing search interfaces in interesting ways.
- These changes lend urgency to the research problems of analyzing video content, interpreting spoken and written natural language, and supporting collaboration among people seeking information together.
- Content analysis over huge collections of user behavior data, combined with interactive user-interface design could lead to breakthroughs in such long-standing problems as human-computer dialogues for question answering.

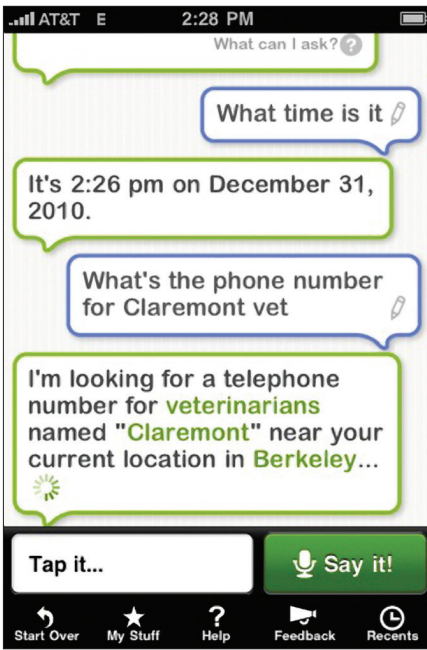


Figure 1. The Siri interface accepts speech as input, attempting to support a dialogue; in the first action, a query for a phone number shows a message conveying the system's understanding of the question as it performs a search.

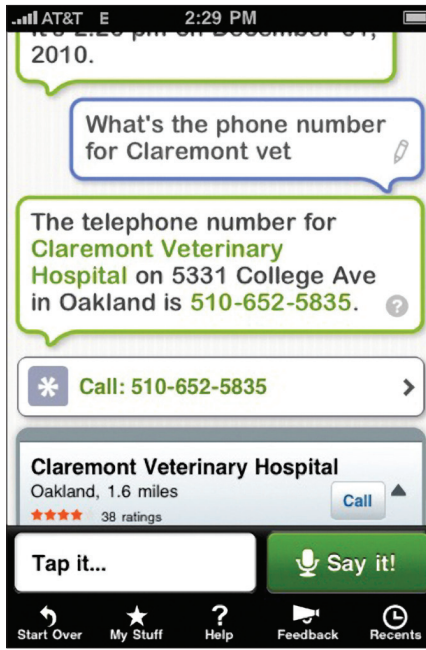


Figure 2. In the second action, the system replaces the confirmation bubble with the answer, offering a command to make the phone call; note the use of graphics alongside text to enrich the output.

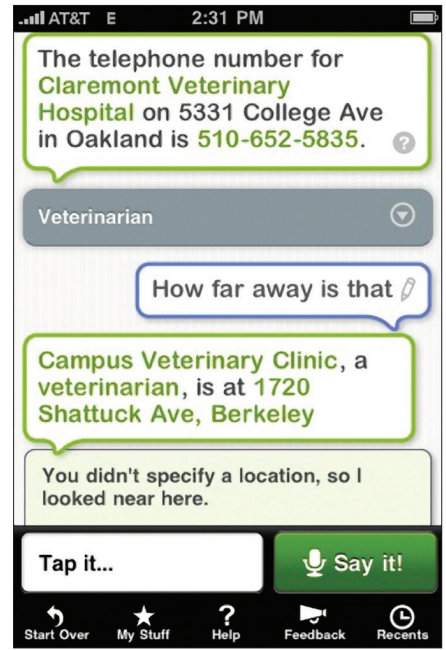


Figure 3. In the concluding action, the system partly understands the intention of the question (about a location) but incorrectly interprets a request for a distance as a request for a location.

and suggesting 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 years. At least three technological trends support the move toward spoken queries: First, 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 through the use of mobile phones. (To assemble a large training set of spoken corrected data for its speech-recognition system, Google hosted, from 2007 to 2010, a free 411 information service for phones.²⁸) And third, touch-screen interfaces are increasingly popular, especially when paired with mobile devices. Neither small devices nor touch screens lend themselves well to typing, making spoken input more attractive, though clever finger-swipe-based input methods (such as ShapeWriter for entering text³⁹ and Gesture Search for menu naviga-

tion¹⁹) provide compelling alternatives to typing.

These trends suggest 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. Though not yet a reality, recent advances are bringing closer the dream of an intelligent interactive agent. For example, the Siri system provides an interface combining local information, speech recognition, easy editing of voice recognition, and visual display of search results. Siri, which was acquired by Apple in April 2010, originated from a Defense Advanced Research Projects Agency research project called CALO (<http://www.ai.sri.com/project/CALO>), in which dozens of computer-science researchers developed machine learning, reasoning, knowledge bases, and other technology to create an intelligent personal assistant.^{4,35}

Though the user's ability to accurately follow up one request with another is limited in Siri, good interface design helps bridge the gap in the back end, since the user sees alternatives and is able to make corrections

(see figures 1–3). Note that Siri also attempts to use searchers' contextual information, including current location. Enormous research interest^{5,20} and commercial development focuses on using time, location, and other contextual cues for search and related applications, and will continue to increase in importance, especially for mobile platforms.

Voice input also has drawbacks, the most significant being that speaking makes noise and can disturb people around the speaker. An exciting research advance would be a microphone that uptakes the words the speaker says but somehow prevents those around the speaker from hearing the words, like a science fiction "cone of silence." Such an invention would have wide-ranging utility for mobile phones.

Social Search: Collaboration

Though observational studies have found that people often search collaboratively, tools have only recently been developed to explicitly support people searching together. Such support reflects a broader research renaissance in tools for real-time shared activity (such as shared online whiteboards and document-editing tools).

One exciting development in collaborative search, from Pickens et al.,^{11,29} assumes the ranking algorithm should allow users to work at their own pace but be influenced in real time by their teammates' search activities. The searchers should not step on one another's proverbial toes; if one person issues a new query, others' thoughts should not be interrupted.

Pickens et al.^{11,29} addressed this issue by developing an algorithm that combines multiple rounds of queries from multiple searchers during a single search session (see Figure 4), using two criteria for weighting results—both 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, while “relevance” is higher for documents closely matching the query. These two factors are combined and continuously updated based on new queries and searcher-specified relevance judgments.

In addition, Pickens et al.^{11,29} assigned different roles to the members of a team. For example, the “Prospector” is in charge of creating new queries to explore new parts of the information space, and the “Miner” looks at the retrieved results to determine which are relevant. Documents not yet looked at are queued up for the Miner interface according to freshness/relevance weighted scores. The Prospector is shown new query-term suggestions based on how they differ from queries already issued, as well as on the relevance judgments made by the Miner. Each role has its own interface; a third view is used to show continually updating information about the queries that have been issued, the documents that have been marked as relevant, and the system-suggested query terms based on the actions of the users.

Another approach to supporting real-time search collaboration, described by Jetter et al.,¹⁶ used a large work surface and input devices combining 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, where each member of the group has different preferences that act

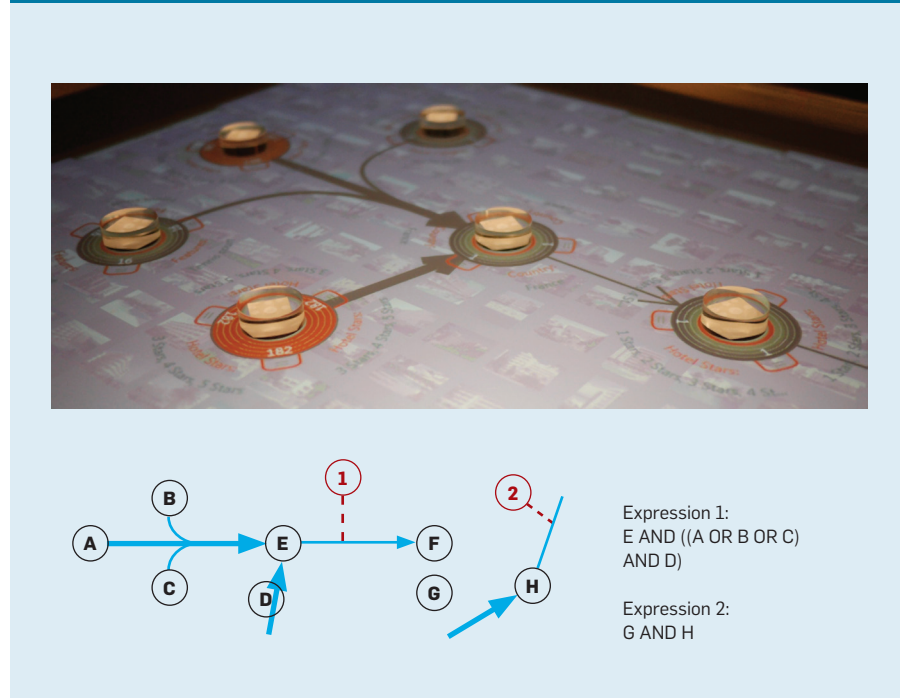
as constraints (such as when choosing a hotel, one needs a heated pool, another wants one that received at least four stars of recommendation, and a third wants the price below a certain amount). Jetter et al.'s solution used a combination of faceted navigation³⁷ and filter-flow visualization,³⁸ showing how many constraints are met by a set of items, given certain constraints. The visualization

was displayed on a shared horizontal workspace, where the controls were manipulated through physical selectors (see Figure 5). Collaboration was facilitated by allowing each user to work privately on a corner of the workspace, then let the results from each piece of the query flow into the rest of the group's query specification. A careful usability study by Jetter et al. found this approach produced results



Figure 4. Collaborating on a video-search task using technology developed by Pickens et al.²⁹; each user views a different unique interface, as well as a shared view. The results of one person's work change the rank ordering of what is seen by the other person.

Figure 5. A collaborative search-formulation tool making use of a large table display, physical input devices, and visualization; from Jetter et al.¹⁶



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 online interaction on social sites is for the social experience of the interaction, rather than for problem-centric information seeking.¹² Reflecting this, a study by Morris et al.²⁴ found the questions asked of others via social networks do not necessarily involve the kinds of information found on static Web pages. Morris et al. asked survey respondents to supply questions they had posed to their social networks on Twitter and Facebook, manually classifying the 249 examples and finding only 17% were for factual information one would typically seek from a Web page (such as how to, say, put an Excel file into LaTeX). The most common categories were requests for recommendations (29%), opinions (22%), rhetorical questions (14%), requests for others to join social events (9%), favors (4%), and social connections, including job openings (3%) and offers of various kinds (1%).

A study of the Aardvark expert social-question-answering system (<http://www.vark.com>) found similar results, with 65% of a random sample of 1,000 queries reflecting a subjective attitude.¹⁵ The questions asked on the social-question-answering site Quora also tend to be subjective and opinion-based; for instance, “What does Dustin Moskovitz think of the new Facebook movie?” was answered by the subject of the question himself.

Unclear is what the best user interfaces are for representing this more social kind of search. Freyne et al.¹⁰ conducted a small study in which different kinds of social cues were shown via icons alongside search-results listings. Subjective results showed a positive preference toward cues showing which articles were read frequently or annotated by others. Yahoo experimented (2005–2009) 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 the recommendation. In March 2011, Google introduced a social-

search tool called “+1” with a similar interface. Significant experimentation on incorporating social information into search results listings is likely over the next few years.

When using a social network to try to answer questions, especially in a work situation, research is ongoing about how best to distribute the related information needs among experts, either within an organization or across the Internet generally.^{18,21} Recent work by Richardson and White³⁴ deployed and studied an instant-messaging-based question-answering service that matched the asker’s questions against predefined profiles of more than 2,000 potential answerers’ expertise, based on their availability. The system contacted three experts at a time, in descending order of how well their profiles matched the content of the question. If an offer to answer was not received within a fixed time limit, the request was sent 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, asking questioners to rate their satisfaction with the answer.

Richardson and White³⁴ examined log data for this system to form an interruption cost model, including how many people should be sent a question in order to minimize disruption while maximizing the likelihood of receiving an informed answer, whether a question will be answered, and how well the asker will be satisfied with the answer received.

Expert solicitation systems that are sophisticated about targeting people with the right expertise and state of mind to address a request are likely to become a fixture in knowledge-centric workplaces, as well as in volunteer causes (such as the Peer2Patent project for community input of patent prior art²⁶).

Social Search: Crowdsourcing

The word “collaboration” as it is used here refers to a set of people working together closely, usually synchronously, to achieve a goal. “Crowdsourcing” refers to large groups of people not necessarily working together knowingly but each contributing in small ways,

leading to a greater whole, as seen, in, for example, Wikipedia editing.

Crowdsourcing in information seeking is seen in Web sites in which communities curate and rate information and share it with others, including question-answering sites, and in product-reviewing sites, bookmark-sharing sites like Delicious, and news-ranking and aggregation sites like Digg. The more-explicitly networked social tools (such as Twitter and Facebook) also function as real-time socially targeted information sources.


Multiple efforts have sought to use explicit user input to improve search-results ranking, though few survive; for instance, Google’s SearchWiki, which allowed users to explicitly reorder search results and share this re-ranking information with others, was shut down in 2010. The Blekko Web search engine, launched October 2010, is an 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,” or subject-specific, search by labeling Web pages with a category label preceded by a slash; they can also mark pages as spam. These two operations together impose crowdsourced quality control over retrieved Web pages. Blekko also provides a social feature allowing users 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 Web search engines have used a form of implicit crowdsourcing for years, by modifying ranking algorithms based on huge quantities of user clickthrough data¹⁷ or predicting which vertical subject area (such as music, news, and travel) to use to augment a query.⁷ Richer user behavior data (such as mouse movements, page dwell time, and searchers’ click paths many steps from the search results page—even across domains—to their destination page) has helped produce useful suggestions of pages not related to the original page through close keyword matches.³⁶


Natural Language-like Queries

Though keyword querying remains standard practice on the Web, savvy users have been typing more detailed queries for years, and Web search engines have greatly improved their ability to handle long queries. Research has shown that people prefer natural expression of queries over keywords,^{3,30} and Web search engine query length continues to increase. According to Experian Hitwise,²² a global online competitive intelligence service, when comparing queries over a four-week period (August–September 2010) to the same four-week period in 2009, found that searches of from five to eight words were up 10%, while searches of from one to four words were down 2%. The growth of query length suggests a desire to express one's information needs more thoroughly and may pave the way toward full-sentence queries. Spoken queries are also likely to be full sentences when speech recognition is faster and more accurate.

Longer queries are also being helped by the online use of colloquial language. When most content is technical or scientific (as was characteristic of the early Web), there is less likely an easy-to-find match between a lay user's words and the words used in the informative documents. Popular question-answering sites (such as Answers.com, Quora, and Yahoo Answers) that store user-generated content bridge 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, a keyword query could fail. But the query "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 a product called either "video selector" or "two-way A/V switcher." The point is that, though the searcher lacks the vocabulary to look up what is needed, the searcher has the same vocabulary as other people in the same cognitive situation. The combination of text worded colloquially and search engines that do a good job with sentence-length queries helps resolve the vocabulary problem. Considerable work has focused on how



Though observational studies have found that people often search collaboratively, tools have only recently been developed to explicitly support people searching together.



to search question-answering sites^{1,2}; ranking algorithms that make use of these mappings will continue to improve results for difficult queries.

Another technical development that may help users who express themselves through long queries is systems that support quasi-natural language interfaces. The new syntax is tolerant of variations, relatively robust, and "exhibit[s] slight touches of natural language flexibility."²⁵ These interfaces are seen in Web search engines supporting various wordings for certain kinds of questions that retrieve answers from a database, as in "Istanbul time," "What is the time in Istanbul?," and "What time is it? Istanbul." Blekko allows query modification through a simple slash notation to refine results to predefined categories (such as "istanbul /tech" for search results about technology and "istanbul /people" for results labeled relevant to people).

Miller et al.²³ developed tools for "sloppy commands," meaning users have a lot of flexibility as to how they express the command, so memorization is not required to make use of them. The "linguistic command line" of Enso (later Ubiquity)^{8,33} experimented with leniency in operating system command lines. The Quicksilver application lookup tool 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 range of predefined query types that mix structured forms with some flexibility in word order, along with a knowledge base and computational back-end able to handle certain combinations of these inputs. For instance, the query "2 slices of pizza with pepperoni" is decomposed into the base information need (information about pizza) refined by units (slices), the quantity (two), and modifications of the baseline concept (with pepperoni). The result is a table listing calorie and nutrition information. However, the system's interpretive range is limited; the query "recipe for pizza with pepperoni" returns the same measurement information as "pizza with pep-

peroni” instead of a recipe.

This hybrid of improved language analysis, command languages making use of structured knowledge bases, and interaction may well lead to more intelligent interfaces and expanded dialogue-like interaction, as discussed earlier regarding the Siri system. The IBM Watson project, which famously beat the top two human champions in the television game show “Jeopardy!” in February 2011, also employs massive language analysis, knowledge-base analysis, and speech recognition, likely setting the stage for future highly advanced natural-language question-answering systems.⁹


Importance of Video Content

Increasing evidence reflects a preference among ordinary information consumers for video and audio content over textual content. Movies have generally replaced books as cultural touchstones in the U.S. A report by Pew Research included a quote from a media executive saying email messages containing podcasts were opened 20% more often than standard marketing email messages.³²


Also according to Pew, 52% of U.S. adults have watched online videos, with seven in 10 Internet users saying they have.³¹ According to Hitwise, the YouTube video-sharing site was the fifth most visited Web site in the U.S. in 2010,¹⁴ and comScore reported in March 2010 that YouTube users generated a greater search volume than Yahoo or Bing.⁶

Video communication is taking 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 citizen-submitted videos selected by the moderators consisted of people pointing the camera at themselves and speaking their question aloud, with a backdrop consisting of a wall in a room in their homes. There were few visual flourishes, and the video did not add much beyond what a questioner in a live audience would have conveyed. Video is fast becoming a conventional way to communicate.

Mobile devices make it easier to cap-



Still lacking are truly useful tools for cogently skimming video content, summarizing it in a meaningful way, and, more to the point, searching within and across it, though research is active in this area.



ture video, increasing the likelihood of video becoming an even more important form of communication. According to Pew, almost 20% of American adults had, as of 2010, tried video calling on phones or computers, and 23% of U.S. Internet users had used a video chat service (such as Skype). Further, 14% of U.S. Internet users had created and uploaded videos.³¹


No doubt, the technology to support full video use lags significantly behind that of text, but we can surmise that some handy inventions are not far off. Better tools for quick edits are also likely soon, as they have been for image processing; a popular mobile iPhone app called Instagram allows users to snap a photo with their phones, quickly apply filters to produce an “artsy” look, then immediately share the image with a social network. Instagram claimed it attracted one million users within two months of its introduction, October 2010, and seven million by August 2011.

Still lacking are truly useful tools for cogently skimming video content, summarizing it in a meaningful way, and, more to the point, searching within and across it, though research is active in this area.²⁷ YouTube provides tools that automatically provide textual closed captioning over spoken language and can also be used for search; so has a startup company called SpeakerText. Faceted navigation³⁷ has become the method of choice for browsing image collections; perhaps the same will be possible with video collections. However, serious breakthroughs are still needed for both image and video content analysis before such search performance rivals that of text search.

Time constraints imposed by YouTube have resulted in a culture of short videos characterized by focused topics, making 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 (<http://www.khanacademy.org>) are each shorter than 10 minutes, with subject matter easily browsable by title (as in “Circles: Diameter, Radius, and Circumference” and “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.

Conclusion

The future of user interfaces will involve support for natural human interaction, gesturing with fingers, speaking rather than typing, watching video rather than reading, and using IT socially rather than alone. This article has explored why these trends will also affect user interfaces for search, highlighting recent work reflecting these trends. Using advanced processing techniques over huge sets of behavioral data, future search interfaces will better support finding other people to answer questions or provide opinions, more natural dialogue-like interaction, and information expressed as non-textual content through non-textual input. More-natural modes of interaction have long been goals of interface design, but recent developments have brought them closer to reality. 

References

1. Adamic, L.A., Zhang, J., Bakshy, E., and Ackerman, M.S. Knowledge sharing and Yahoo answers: Everyone knows something. In *Proceedings of the 17th International Conference on the World Wide Web* (Beijing). ACM Press, New York, 2008, 665–674.
2. Bian, J., Liu, Y., Agichtein, E., and Zha, H. Finding the right facts in the crowd: Factoid question answering over social media. In *Proceedings of the 17th International Conference on the World Wide Web* (Beijing). ACM Press, New York, 2008, 467–476.
3. Bilal, D. Children's use of the Yahoo!igans! Web search engine: I. Cognitive, physical, and affective behaviors on fact-based search tasks. *Journal of the American Society of Information Science* 51, 7 (2000), 646–65.
4. Chaudhri, V.K., Cheyer, A., Guili, R., Jarrold, B., Myers, K.L., and Niekraz, J. A case study in engineering a knowledge base for an intelligent personal assistant. In *Proceedings of the 2006 Semantic Desktop Workshop* (Athens, GA, 2006).
5. Church, K., Neumann, J., Cherubini, M., and Oliver, N. 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 the World Wide Web* (Raleigh, NC, Apr. 26–30). ACM Press, New York, 2010, 261–270.
6. comScore. comScore releases March 2010 U.S. search engine rankings (Mar. 2010); http://www.comscore.com/Press_Events/Press_Releases/2010/4/comScore_Releases_March_2010_U.S._Search_Engine_Rankings
7. Diaz, F. and Arguello, J. 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* (Boston, July 19–23). ACM Press, New York, 2009, 323–330.
8. Ertwine, M.Y. Ubiquity: Designing a multilingual natural language interface. In *Proceedings of the SIGIR Workshop on Information Access in a Multilingual World* (Boston, July 19–23). ACM Press, New York, 2009.
9. Ferrucci, D., Brown, E., Chu-Carroll, J., Fan, J., Gondek, D., Kalyanpur, A.A., Lally, A., Murdock, J.W., Nyberg, E., Prager, J., et al. Building Watson: An overview of the DeepQA Project. *AI Magazine* 31, 3 (2010).

10. Freyne, J., Farzan, R., Brusilovsky, P., Smyth, B., and Coyle, M. Collecting community wisdom: Integrating social search & social navigation. In *Proceedings of the 12th International Conference on Intelligent User Interfaces* (Honolulu, Jan. 28–31). ACM Press, New York, 2007, 52–61.
11. Golovchinsky, G., Shah, C., and Pickens, J. Role-based results redistribution for collaborative information retrieval. *Information Processing and Management* 46, 6 (2010), 773–781.
12. Harper, F.M., Moy, D., and Konstan, J.A. 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* (Boston, Apr. 4–9). ACM Press, New York, 2009, 759–768.
13. Hearst, M. *Search User Interfaces*. Cambridge University Press, 2009.
14. Hitwise. 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-second-straight-year>
15. Horowitz, D. and Kamvar, S.D. The anatomy of a large-scale social search engine. In *Proceedings of the 19th International Conference on the World Wide Web* (Raleigh, NC, Apr. 26–30). ACM Press, New York, 431–440.
16. Jetter, H.-C., Gerken, J., Zöllner, M., Reiterer, H., and Millic-Frayling, N. 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* (Vancouver, Canada, May 7–12). ACM Press, New York, 2011.
17. Joachims, T., Granka, L., Pan, B., Hembrooke, H., and Gay, G. Accurately interpreting clickthrough data as implicit feedback. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Salvador, Brazil, Aug. 15–19). ACM Press, New York, 2005, 154–161.
18. Kautz, H., Selman, B., and Shah, M. Referral Web: Combining social networks and collaborative filtering. *Commun. ACM* 40, 3 (Mar. 1997), 63–65.
19. Li, Y. Gesture search: A tool for fast mobile data access. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology* (New York, Oct. 3–6). ACM Press, New York, 87–96.
20. Liu, C., Rau, P.L.P., and Gao, F. Mobile information search for location-based information. *Computers in Industry* 61, 4 (May 2010), 364–371.
21. Liu, Y. and Agichtein, E. 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* (Columbus, OH, June 15–20). Association for Computational Linguistics, Stroudsburg, PA, 2008, 97–100.
22. McGee, M. The long tail is alive and well. *Small Business Search Marketing* (Sept. 16, 2010); <http://www.smallbusinesssem.com/long-tail-alive-well/3659/> and http://twitter.com/Hitwise_US/status/24041444164
23. Miller, R.C., Chou, V.H., Bernstein, M., Little, G., Van Kleek, M., and Karger, D. 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* (Monterey, CA, Oct. 19–22). ACM Press, New York, 2008, 131–140.
24. Morris, M.R., Teevan, J., and Panovich, K. 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* (Atlanta, Apr. 10–15). ACM Press, New York, 2010, 1739–1748.
25. Norman, D. The next UI breakthrough: Command lines. *Interactions* 14, 3 (2007), 44–45.
26. Noveck, B.S. Peer to patent: Collective intelligence, open review, and patent reform. *Harvard Journal of Law & Technology* 20, 1 (2006), 123–162.
27. Over, P., Awad, G., Fiscus, J., Antonishek, B., and Michel, M. TRECVID 2010: An introduction to the goals, tasks, data, evaluation mechanisms, and metrics. *Proceedings of the Eighth TRECVID Workshop*. National Institute of Standards and Technology, Gaithersburg, MD, 2010.
28. Peres, J.C. Google wants your phonemes. *InfoWorld* (Oct. 23, 2007); <http://www.infoworld.com/t/data->

- management/google-wants-your-phonemes-539
29. Pickens, J., Golovchinsky, G., Shah, C., Qvarfordt, P., and Back, M. Algorithmic mediation for collaborative exploratory search. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Singapore, July 20–24). ACM Press, New York, 2008, 315–322.
30. Pollock, A. and Hockley, A. What's wrong with Internet searching. *D-Lib Magazine* (Mar. 1997); <http://www.dlib.org/dlib/march97/bt/03pollock.html>
31. Purcell, K. *The State of Online Video*. Pew Internet & American Life Project, Washington, D.C., June 3, 2010; <http://www.pewinternet.org/~media/Files/Reports/2010/PIP-The-State-of-Online-Video.pdf>
32. Ranie, L. *Digital 'Natives' Invade the Workplace*. Pew Internet & American Life Project, Washington, D.C., Sept. 28, 2006; <http://pewresearch.org/pubs/70/digital-natives-invade-the-workplace>
33. Raskin, A. The linguistic command line. *Interactions* 15, 1 (2008), 19–22.
34. Richardson, M. and White, R. Supporting synchronous social Q&A throughout the question life cycle. In *Proceedings of the 20th International World Wide Web Conference* (Hyderabad, India, Mar. 28–Apr. 1, 2011).
35. Roush, W. 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. White, R.W., Bilenko, M., and Cucerzan, S. 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* (Amsterdam, The Netherlands, July 23–27). ACM Press, New York, 2007, 159–166.
37. Yee, K.P., Swearingen, K., Li, K., and Hearst, M. Faceted metadata for image search and browsing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Fort Lauderdale, FL, Apr. 5–10). ACM Press, New York, 2003, 401–408.
38. Young, D. and Shneiderman, B. A graphical filter/flow representation of Boolean queries: A prototype implementation and evaluation. *Journal of the American Society for Information Science* 44, 6 (July 1993), 327–339.
39. Zhai, S., Kristensson, P.O., Gong, P., Greiner, M., Peng, S.A., Liu, L.M., and Dunnigan, A. 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* (Boston, Apr. 4–9). ACM Press, New York, 2667–2670.

Marti A. Hearst (hearst@ischool.berkeley.edu) is a professor in the School of Information at the University of California, Berkeley, with an affiliate position in the Computer Science Division. She is the author of *Search User Interfaces*, Cambridge University Press, 2009.