

Viewpoint

Opportunities and Challenges in Search Interaction

Seeking to address a wider range of user requests toward task completion.

INTERACTING WITH SEARCH systems, such as Web search engines, is the primary means of information access for most people. Search providers have invested billions of dollars developing search technologies, which power search engines and feature in many of today’s virtual assistants (including Google Assistant, Amazon Alexa, Microsoft Cortana, and others). For decades, search has offered a plentiful selection of research challenges for computer scientists and the advertising models that fund industry investments are highly lucrative. Given the phenomenal success, search is often considered a “solved problem.” There is some truth to this for fact-finding and navigational searches, but the interaction model and the underlying algorithms are still brittle in the face of complex tasks and other challenges, for example, presenting results in non-visual settings such as smart speakers.¹⁵ As a community, we need to invest in evolving search interaction to, among other things, address a broader range of requests, embrace new technologies, and support the often underserved “last mile” in search interaction: task completion.

Search Interaction

The retrieval and comprehension of information is important in many settings. Billions of search queries reach search engines daily and searching skills are now even taught in schools. Search interaction has been studied



by information science, information retrieval (IR), and human-computer interaction (HCI) researchers. Information scientists have examined the cognitive and behavioral mechanisms in the search process. IR researchers have developed new methods to collect and find information, including, recently, increased use of machine learning. HCI researchers have studied interactions with technology to develop interfaces to support activities such as information finding and sensemaking. Future opportunities are plentiful, including the three areas discussed in this Viewpoint:

- For more than a decade, search interaction has been immersed in a data revolution, using big (population) data¹ and small (personal) data³ to model search activity and improve search experiences. This has used traditional

data sources (queries, clicks), but richer data (browse, cursor, physiology, spatial context, and so forth) is emerging that enables search systems to more fully represent interests and intentions, unlocking sophisticated modeling methods such as deep learning.

- Support for search interaction has focused on helping searchers build queries and select results. Search systems must evolve to support more complex search activities, leveraging technological advances to meet people’s growing expectations about search capabilities.

- Virtual assistants offer an alternative means to engage with search systems. Assistants support rudimentary question answering but will soon more fully comprehend question semantics, understand intent through dialog, and support task completion through skill chaining and skill recommendation.

Data Revolution

There have been three documented “revolutions” in search-related research: cognitive (targeting intellectual processes); relevance (understanding different relevance types and criteria); and interactive (providing search support and capturing searcher preferences).¹² The interactive revolution continues to this day. We are also in the middle of a fourth revolution: the *data* revolution, driven by enhanced capabilities and interest in recording, analyzing, and learning from user data, both in the aggregate and individually. Application of data mining and machine learning models to query-click data has yielded improvements in ranking, query suggestion, and search advertising, as well as better understanding searchers, their activities, and their satisfaction and success.

Going forward, search *tasks* must be regarded as first-class elements in the search process. Session data is still only used to augment individual queries³ when the focus should be on supporting end-to-end task completion. Web browser trails capture behavioral traces in online environments that can help direct others.¹³ Search providers could mine such activity sequences to harness the procedural search knowledge of populations. These could be used in strategic search support, such as guided tours, that span full tasks, not just the starting points offered in today’s search results.⁴

Web search engines have made substantial progress in better understanding the intended meaning of a query, for example, recognizing mentions of entities and common query patterns. These are used not only to improve the precision and recall of query results, but also to attempt to provide a direct “best answer” for the most likely query meaning. Beyond query text, many new signals are available to search systems through new interaction modalities such as touch and gesture, as well as sensors tracking signals including physiology, eye gaze, and locomotion. Cursor movements can also be collected at scale and help interpret user activity in the absence of click-through data. Beyond interactions, search engines are also increasingly using semantic data to better understand document content. This data is sourced from background knowledge graphs and

The wealth of opportunity should not translate to dramatically increased complexity.

from the documents themselves in the form of embedded semantic data using the common schema.org ontology and markup standards (microdata, RDFa and JSON-LD). Cloud services mean data collected from users and elsewhere is no longer siloed in specific machines or applications. Longitudinal data about searchers helps build rich models of their interests and expertise. Search personalization is operationalized using short- and long-term data from individuals,³ which can be scaled to cohorts if data is sparse. Even non-search services (for example, productivity applications) can offer data to enrich contextual models and improve search effectiveness. For example, a search for “VAR” from a spreadsheet provides evidence the intent is variance, and not value at risk, for example. Other signals such as spatial context and time offer rich information about the search situation.

Using activity data at massive scale offers incredible potential to understand the human condition. Although logs lack ground truth about search intent, success, experience, and attention, they can still help characterize search behavior, build machine-learned models, and make meaningful discoveries, for example, forecasting influenza in populations.⁶ Access to this data is restricted to search providers or only available for purchase from analytics companies for a significant fee, hindering scientific progress. To help address this, search providers have released limited search log data and other resources (for example, Microsoft recently released MARCO,^a a machine reading comprehension dataset), and some researchers broadly share user study data (an encouraging

a <http://www.msmarco.org/>

trend). Open data movements such as data.gov promote data availability, but not for search data, at least not yet.

In working with search interaction data (or any user data) to make intelligent inferences, privacy and data reliability are paramount. Privacy concerns stemming from the construction of user profiles and detailed surveillance of people’s activities must be addressed. Systems should obtain user consent and offer clear explanations about what is being recorded and how it is being used. Search providers must act responsibly and correct any biases in search results,⁷ in data collected from users and in user sampling. Humans are affected by many factors impacting recorded activities¹⁴ (for example, cognitive biases, behavioral biases, common misconceptions, and misinformation and rumor). These factors can skew behavioral signals such as click-through rates used in ranking algorithms, creating “filter bubbles.”¹¹ This must be considered during data collection and experimental analyses.⁵

Many of these lessons apply in domains beyond Web search. Much of search is domain specific, including legal, medical, and intellectual property. Even within Web search, there are different verticals (including images, video, news) each with its own presentation format and interaction method (for example, “infinite scrolling” in image search). Boundaries between vertical and generic search are blurring as content from verticals bleeds into general result pages, affecting search interactions.¹⁰

Evolving Capabilities and Expectations

Advances in data availability coupled with new interaction paradigms (such as touch, gaze, large displays, gesture, spoken dialog), mobile computing capabilities (including tablets, smartphones, smartwatches), and the democratization of artificial intelligence, have created new opportunities for information access and use. Searchers can now interact with search systems in more lightweight and natural ways,⁷ including while engaging in non-search tasks. Information visualization tools such as Microsoft SandDance^b

b <https://www.sanddance.ms/>

help people explore and understand data, building on prior HCI research on visualization.² Machine learning advances yield significant gains in conversational intelligence and question answering. Improvements in near- and far-field speech recognition coupled with new dialog research make conversational search feasible. Even within current interaction paradigms, deeply understanding query and document semantics can help provide more intelligent responses; for example, medical symptom answers on Google and multi-perspective answers on Bing.

Mobile devices such as smartphones and tablets are powerful and versatile. The integration of hardware such as accelerometers, gyroscopes, and proximity sensors provides rich contextual signals about user activities that are useful for search and recommendation. Evidence from self-reports and log analysis suggests people now demand search support in more situations—to resolve a diverse set of questions (or arguments!)—and question complexity continues to rise. Complex tasks spanning devices are also more frequent. Search systems can utilize downtime between task activities to perform “slow searches,” for example, finding sets of relevant resources or using crowdworkers to compose answers.

Wearable and augmented reality applications support the presentation of relevant information just in time, in anticipation of its use. Hardware such as hearables (for example, Google Pixel Buds) or head-mounted displays (such as Google Glass, Microsoft HoloLens) provide continuous information access in any setting. For some tasks (for example, monitoring activities), relevant information can be offered proactively, capitalizing on signals such as user preferences and location. Proactive notifications need to be carefully gated and privacy must be respected, including the privacy of any collocated individuals.

The wealth of opportunity should not translate to dramatically increased complexity. The prevalence of the Google interface design has meant searchers expect simplicity, and rightly so: search activities are already sufficiently complex. Any new capabilities must be intuitive, simple, and add clear value.

Virtual Assistants

Integration with virtual assistants such as Amazon Alexa, Google Assistant, or Microsoft Cortana allows search systems to extend their capabilities to better understand needs and support higher-order search activities such as learning, decision making, and action.⁹ Search engines can provide an entry point to virtual assistants when search requests demand additional engagement (for example, are non- navigational). Search technology already powers some virtual assistants, and knowledge bases created for information finding have utility herein. End-to-end task completion (that is, from search interactions to action in the physical world) has traditionally been underserved by search engines. This can be achieved via first- and third-party skills in virtual assistants. Skills best suited to the current context can be recommended by assistants and even chained together to support multistage tasks.

Virtual assistants are particularly amenable to supporting search interaction: they are personal and contextual, they support dialog, and they are ubiquitous (across applications and devices). Deep understanding of searchers and their contexts is necessary to adapt system responses to the situation. Natural interactions, including multi-turn dialogs, enable search systems to clarify searcher needs. Conversational search is already attracting significant interest.⁸ Ubiquity has advantages beyond availability, that is, richer data enables sophisticated inferences such as automatically detecting task completion or estimating task duration, as well as supporting rapid task resumption.

Despite its promise, search-assistant integration is not without challenges that require rethinking several aspects of search interaction. For example, although virtual assistants can foster dialog, natural language conversations can be inefficient ways to obtain answers or complete tasks. Virtual assistants often manifest in headless devices such as smart speakers and personal audio, making it difficult to communicate result lists or discover assistant capabilities.¹⁶ Also, the traditional search-advertising model depends on visual attention and does not scale well to audio-only settings.

Looking Ahead

We are just beginning a journey to a more enlightened society facilitated by interactions with search systems. Looking ahead, the data revolution in search interaction will gather pace, searchers will engage with search systems in new ways, and virtual assistants will serve as comprehensive search companions. Building on these and other pillars, search systems will empower people and support the activities they value. This important effort will only succeed given the expertise, collaboration, and commitment of communities within computer science and beyond. **□**

References

1. Agichtein, E., Brill, E., and Dumais, S. Improving web search ranking by incorporating user behavior information. In *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2006, 19–26.
2. Ahlberg, C., Williamson, C., and Shneiderman, B. Dynamic queries for information exploration: An implementation and evaluation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (1992), 619–626.
3. Bennett, P.N. et al. Modeling the impact of short- and long-term behavior on search personalization. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2012, 185–194.
4. Bush, V. As we may think. *The Atlantic Monthly* 176, 1 (Jan. 1945), 101–108.
5. Eckles, D., Karrer, B., and Ugander, J. Design and analysis of experiments in networks: Reducing bias from interference. *Journal of Causal Inference* 5, 1 (Jan. 2017).
6. Ginsberg, J. et al. Detecting influenza epidemics using search engine query data. *Nature* 457, 7232, (2009), 1012–1014.
7. Kay, M., Matuszek, C., and Munson, S.A. Unequal representation and gender stereotypes in image search results for occupations. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems ACM*, 2015, 3819–3828.
8. Hearst, M.A. (2011). “Natural” search user interfaces. *Commun. ACM* 54, 11 (Nov. 2011), 60–67.
9. Marchionini, G. Exploratory search: From finding to understanding. *Commun. ACM* 49, 4 (Apr. 2006), 41–46.
10. Metrikov, P. et al. Whole page optimization: How page elements interact with the position auction. In *Proceedings of the 15th ACM Conference on Economics and Computation*. ACM, 583–600, 2014.
11. Pariser, E. *The Filter Bubble: How the New Personalized Web is Changing What We Read and How We Think*. Penguin, New York, NY, 2011.
12. Robertson, S.E. and Hancock-Beaulieu, M.M. On the evaluation of IR systems. *Information Processing and Management* 28, 4 (Apr. 1992), 457–466.
13. 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*. ACM, 159–166, 2007.
14. White, R.W. Beliefs and biases in web search. In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 3–12, 2013.
15. White, R.W. *Interactions with Search Systems*. Cambridge University Press, New York, NY, 2016.
16. White, R.W. Skill discovery in virtual assistants. *Commun. ACM* 61, 11 (Nov. 2018), 108–115.

Ryen W. White (ryenw@microsoft.com) is a Partner Researcher and a Research Manager at Microsoft Research AI, Redmond, WA, USA.

Copyright by author.