

Reconsidering Relevance and Embracing Interaction

by Daniel Tunkelang

Search 2009

We live in an age of ubiquitous search. Children growing up today in the developed world would probably find a world without search engines as inconceivable as the rest of us would find a world without electricity and running water. Moreover, search engine users – that is to say, almost all of us – expect those engines to work like magic, extrapolating our information needs from a pair of words (on average) and responding with the content most relevant to those needs.

Of course, search engines do not always meet these expectations. Thousands of engineers work at – or with – search engine companies to tune the ways those engines model and compute relevance. Relevance, however, is not a one-size-fits-all measure. For example, web search engines typically measure the relevance of web pages using bibliometric considerations such as Google’s PageRank measure and the “anchor text” associated with links connecting into pages. In contrast, enterprises cannot typically rely on a rich network connecting the documents stored in their repositories; hence, their search engines look for relevance signals in document structure (for example, title and keyword fields) and consider factors specific to their domains or applications. Of course, almost every search engine’s filtering and ranking algorithms take into account the relationship between words in the search query and words that occur in documents.

The Problem with Relevance

But relevance as a concept is problematic. In 1964, William Goffman wrote, “[T]he relationship between the document and the [search] query,

though necessary, is not sufficient to determine relevance” [1]. In other words, the query does not provide the search engine with enough information to reliably determine how relevant a document is to the user’s information need. Library and information scientists like Tefko Saracevic and Nicholas Belkin have been saying as much for years, emphasizing an approach to information seeking that focuses on the user [2, 3]. Ironically, the success of commercial search engines has marginalized many of their concerns: search engines do well enough without considering the users – at least for web search – that many search engine researchers and developers have been content to ignore Goffman’s warning.

Lately, however, there have been signs that our honeymoon with purely ranked retrieval systems is coming to a close. Bing, Microsoft’s recent re-launch of its web search offering, touts itself as a “decision engine” rather than simply an engine to match search queries to documents. More substantively, Bing’s interface offers users a variety of ways to interact with the search engine. For example, a search for “Barack Obama” offers a set of query refinements that includes news, biography and interviews. Google, though not quite as aggressive on this front, has been slowly promoting similar features, initially relegating them to its advanced search and experimental interfaces but increasingly raising their visibility. Last but not least, tools like Yahoo!’s Search Assist provide real-time query suggestions to users as they enter search queries. On the web, we are seeing the initial signs of search engines engaging users in a more interactive query elaboration process.

Meanwhile, search applications outside of general web search have been embracing interaction for several years. In particular, faceted search has become a staple of online retail sites. Endeca established faceted catalog search as an industry standard, and mega-retailers Amazon and eBay have

Daniel Tunkelang is chief scientist at Endeca. He can be reached by email at DTunkelang<at>Endeca.com

since built out their own implementations. Today, site developers have a variety of commercial and open-source options available to them.

Faceted search combines conventional text search with the ability to refine query results using multiple facets of those results, enabling users to navigate content in multiple ways, rather than in a single, pre-determined, hierarchical order. For example, a shopping site might allow users to refine product results by such facets as product type, brand and price, while a directory might allow users to refine people results by age, gender and location.

Online retailers are some of the most visible users of faceted search, but they are by no means the only ones. Libraries are increasingly using it for their next-generation online public access catalogs (OPACs), almost a century after the Indian librarian S. R. Ranganathan promoted the advantages of a faceted classification scheme as an alternative to hierarchical schemes like the Dewey Decimal system [4]. More broadly, site and enterprise search applications have learned that the best way to help users satisfy their information needs is to provide users with a richer interface and user experience. The “10 blue links” interface is clearly not enough to satisfy a wide variety of information seeking applications.

Human-Computer Information Retrieval

Guided query refinement and faceted search fit into a class of approaches called human-computer information retrieval (HCIR), a term coined by Gary Marchionini [5]. The HCIR approach advocates for tools that bring human intelligence and attention actively into the search process. Rather than guessing what users need, these tools provide users with opportunities to clarify and elaborate their intent. If the engine isn’t sure what users want, it just asks them.

In economic terms, HCIR aims to offer users better return on investment. Instead of slavishly accepting the constraints of the current interaction metaphor (users enter two words as input and see a ranked list of ten results as output) and attempting to optimize the user experience within those constraints, a search engine can allow users to get more if they give more. But what should it ask users to give? And what will users get in return?

We answer those questions in reverse order. A search engine should provide users with an interface that offers transparency about the engine’s

internal state, control over the filtering and ranking of results, and guidance to inform an adaptive information seeking strategy. In return, users can and will help the search engine help them, using the available query elaboration options to clarify and adjust their expressed information needs throughout the information seeking process.

Let us now consider the three goals of transparency, control and guidance.

Transparency

The premise of transparency is simple: users should know why a search engine returns a particular response to their query. Note the emphasis on “why” rather than “how” since most users don’t know – or care – what algorithm a search engine uses to in order to assemble a response to a query. What users do care about is whether the engine “understood” their query – that is, whether the engine is at least attempting to address the information need that the user intended to express.

A recent study by Autobytel and Kelyon Research, entitled “The State of Search,” reported that most users expect search engines to read their minds [6]. Of course, such an expectation is unreasonable – even our closest friends cannot read our minds. What frustrates users most, however, is when a search engine not only fails to read the user’s mind, but also gives no indication of where the communication broke down, let alone how to fix it. In short, the search engine fails to provide transparency.

We are familiar with the quip that “to err is human, but to really foul things up requires a computer.” It shouldn’t surprise us that a search engine may really foul things up when it attempts to extrapolate an information need from a couple of words. If, however, the search engine provided the user with a clear picture of how it arrived at its response, then the user could adapt to the engine’s limited cognitive capabilities – in particular, the engine’s inability to read the user’s mind.

Control

Transparency is necessary for users to establish effective communication with search engines, but it is not sufficient. Indeed, transparency might even cause more harm than good if the users have visibility into the system’s

misunderstandings but no ability to resolve them. To overcome this problem, search engines also need to offer users control. The ultimate arbiter of the information seeking process must be the user, not the system.

Interestingly, early search engines offered users more control than most modern ones. In particular, commercial search engines in the 1970s and 1980s, some of which are still used for applications like electronic discovery, offered users Boolean search, enabling users to construct arbitrarily complex queries that the system would follow precisely. Indeed, Boolean search interfaces not only provide control to users, but also offer them a high degree of transparency. Unfortunately, their lack of guidance (which we will discuss in a moment) proved fatal: most users – even professional information seekers – are not adept at query construction, especially when they are querying unfamiliar content repositories [7].

But many people learned the wrong lesson from the weaknesses of Boolean search interfaces. As a result, modern search engines generally offer a ranked retrieval model that not only is opaque to users, but also offers them minimal control over the information-seeking process beyond the ability to enter words into the search box. In particular, the ranking of results, an essential function of the search engine when the number of matching results far exceeds the number shown to the user, is typically outside the user's control.

There is no reason to prevent users from having full control over the filtering and ranking of results. Offering users control does not preclude (or excuse) a search engine from making a best effort to satisfy users who do not exercise that control. Web search providers have argued that the threat of spam and unscrupulous search engine optimization requires protecting users with secret filtering and ranking algorithms – an argument reminiscent of security through obscurity. To the contrary, giving users more control would actually offer them more robust protection, since the diversity of user behavior would dilute the spammers' efforts. All search engines should treat their users as the primary actors in the information-seeking process, rather than mostly passive recipients of their paternalism.

Guidance

While transparency helps users understand the system internal state, and

control allows users to override the system, these two ingredients, as shown by the early experience with Boolean search systems, are necessary but not sufficient. The third essential ingredient to ensure a successful user experience is guidance: The system must not only respond to users' queries, but also help users formulate those queries.

In the 1990s, Peter Pirolli and Stuart Card developed a theory of information foraging, establishing an analogy between the way users seek information and the way animals forage for food [8]. The core concept in this theory is "information scent," the indicators or clues that lead users to choose the next step in the information-seeking process with the expectation that it will help them progress toward their goal. Earlier work by Marcia Bates described a "berrypicking" model of information retrieval, in which a user does not satisfy an information need in a single query, but rather through a series of steps in which the user either finds information directly relevant to the information need or learns how to adapt the search strategy [9].

Like the information architects who organize the content on websites, search engine designers should aspire to provide users with scent at every step of their information-seeking process. Techniques like query suggestions, faceted search and results clustering all offer users the opportunity to make progress on their next step, rather than always having to restart the information-seeking process from scratch. Indeed, faceted search is a popular technique for offering users such guidance. While users are ultimately responsible for expressing their information needs, it is the search engine's job to act like a reference librarian and help the users in this process.

From Vision to Reality

Hopefully you are now convinced that search engines should offer users transparency, control and guidance. Granted, it is one thing to advocate these goals and another to actually achieve them. A key obstacle is inertia. Both users and search engine developers are comfortably familiar with a standard interface and wary of making radical changes to it. Nonetheless, the successes in highly visible domains like online shopping have demonstrated that users will take advantage of better interfaces, given the opportunity. Meanwhile, researchers and practitioners are continuing to invest in better

TUNKELANG, continued

methods to support the query elaboration process. The status quo is a significant challenge, but not an insurmountable one.

As users, we must demand to be treated as masters of our information-seeking destiny, embracing the responsibility concomitant with that power.

As search engine developers, we must place users first, giving them control in order to get their active participation in the information-seeking process. Together, we can bring forth a brave new world of human-computer information retrieval.

Resources Mentioned in the Article

- [1] Goffman, W. (1964). On relevance as a measure. *Information Storage and Retrieval*, 2(3), 201-203.
- [2] Saracevic, T. (2006). Relevance: A review of the literature and a framework for thinking on the notion in information science. Part II. *Advances in Librarianship*, 30, 3-71.
- [3] Belkin, N. (2008). Some(what) grand challenges for information retrieval. *ACM SIGIR Forum*, 42(1), 47-54. (Keynote address presented at the 30th Annual European Conference on Information Retrieval (ECIR '08)).
- [4] Ranganathan, S. R. (1951). *Philosophy of library classification*. Copenhagen: Ejnar Munksgaard.
- [5] Marchionini, G. (June/July 2006). Toward human-computer information retrieval. *Bulletin of the American Society for Information Science and Technology*, 32(5). Retrieved August 20, 2009, from www.asist.org/Bulletin/Jun-06/marchionini.html.
- [6] Autobyte. (October 10, 2007). Suffering from information overload? You aren't the only one! 72 percent of consumers report having "search engine fatigue," reports Autobyte [Press release about *The state of search* by Autobyte and Kelton Research]. Retrieved August 20, 2009, from www.autobyte.com/content/home/help/pressroom/pressreleases/index.cfm/action/template/article_id_int/2801
- [7] Turtle, H. (1994). Natural language vs. Boolean query evaluation: A comparison of retrieval performance. *Proceedings of 17th Annual ACM SIGIR International Conference on Research and Development in Information Retrieval*, 212-220.
- [8] Pirolli, P., & Card, S. 1999. Information foraging. *Psychological Review*, 106(4), 643-675.
- [9] Bates, M. (October 1989). The design of browsing and berrypicking techniques for the online search interface. *Online Review*, 13, 407-424.