Trends in Search Interaction

Ricardo Baeza-Yates[1], Paolo Boldi[2], Alessandro Bozzon[3], Marco Brambilla[3], Stefano Ceri[3], and Gabriella Pasi[4]

[1] Yahoo! Research
Avinguda Diagonal 177, 08018 Barcelona, Spain
rbaeza@acm.org
[2] Università degli Studi di Milano, Dipartimento di Scienze dell'Informazione
Via Comelico 39/41, I-20135 Milano, Italy
boldi@dsi.unimi.it
[3] Politecnico di Milano, Dipartimento di Elettronica e Informazione
Piazza Leonardo da Vinci 32, I-20133 Milano, Italy
{firstname.lastname}@polimi.it
[4] Università degli Studi di Milano Bicocca, DISCO
Viale Sarca 336, I-20126 Milano, Italy
pasi@disco.unimib.it

Abstract. This paper reports the main findings of a panel about trends in search engine interaction, focused upon the use of search engines for performing complex processes[1]. The discussion focuses on the different evolutionary path followed by search engines with respect to other Web and information management solutions, making end users acquainted with the simplistic and never changing keyword-based query paradigm. The analysis delves into the pros and cons of personalization, contextualization, and exploration of Web information, with special attention to the presentation and user interaction aspects. In the end, we also wonder if the keyword-based query paradigm will ever change.

1 Introduction

The technology of search engines and the amount and quality of services they offer have radically evolved in the last twenty years, in parallel with their ever-increasing market value: not only they largely remain the main entry point to the Web for the vast majority of users, but they are constantly trying to develop new features that are more loosely related to search and should serve to enhance the Web-user experience; such functionalities (integrated Webmail, Web-editing of documents, personal calendars, image and video hosting facilities, personalized information delivery, to cite only a few) are actually changing our way of thinking about computations and Web services in general. As often observed, they are in a sense pointing back to the idea of centralized computation, making browsers similar to old-days' terminals.

Notwithstanding this evolution, the basic service provided by search engines has remained the same, and it is surprising to observe how the basic textual keyword-based

[1] Panel Session, Workshop on "Search as a Process", Politecnico di Milano, May 26 2010.

S. Ceri and M. Brambilla (Eds.): Search Computing II, LNCS 6585, pp. 26–32, 2011.

search interface has not substantially evolved or changed over time: alternative proposals (using portals, employing other more sophisticated tools for information hunting, etc.) did not (yet) take off, except possibly in some specialized niches.

Such a long-lasting success is most uncommon in computer science and should suggest that the very concept of keyword-based search is extremely robust although it not always matches users' expectations and satisfaction.

Of course, behind the scenes, search engines changed a lot. Companies such as Google, Yahoo! and Microsoft have tremendously improved the search engine technology, in particular in scalable distributed systems and sophisticated query processing and ranking techniques. While those innovations did not affect the way in which users approach search through simple keyword-based interfaces, they are changing the way in which results are presented, by trying to guess and satisfy better the user's needs. Let us make an attempt to understand more thoroughly in which directions they are moving to reach this goal.

2 Helping the User

A general motto that is repeated constantly by engineers working in search is that one "should help the user find what (s)he's looking for". But not all users are the same, and not all searches are the same, so helping the user really does not mean the same thing for all searches and for all users. The "one size fits all" paradigm commonly applied by search engines does not adhere to the complex and dynamic characteristics of typical search tasks. In a complex search, the user knows what (s)he wants, and this is the result of a complex task (s)he has in mind, while in exploratory search the user does not have a precise idea of what (s)he wants [20], although the result may be fairly simple. Hence, these two dimensions are somehow orthogonal. In the former case, the system should understand that such a search mission is ongoing, and react by offering diverse suggestions and help, possibly keeping track of what the user has collected so far and trying to make sense of it. In the latter case, a more liberal (and prudent) attitude should be adopted, avoiding to interfere too much with the user's search, and act in a way that may sometimes be felt as unwanted or even intrusive.

Many recent works [3][7][9][10] focus on how one may understand whether a complex search mission is underway, for example, by storing and using the amount of knowledge collected from what other users did in the past (an obviously precious source of information that is collected in search engine query logs). It is important to observe that a complex search mission is probably a long-lasting one. Tasks are "long lived" and occur across sessions – people start to plan their trip to Greece one day and they continue the next day - therefore it would be useful for a search engine to recall/reconstruct previous interactions from the same user and to build personalized histories. A personalized and collaborative search approach can be of help in this context. A further difficulty is that users often "intertwine missions", i.e., they switch from a task to another (e.g., alternate work with leisure) - therefore systems should be able to capture interest shifts (especially temporary ones): this ability has proven to be difficult to achieve, although in many cases intertwined missions actually boil down to simple diversions.

A subtle point to be considered here is that complex searches often present multiple facets and may grow in different directions: therefore, reconstructing the state of a search process may include detecting branches in that process and then associating each branch with different results, as diversity and richness of results is not less important than appropriateness and completeness [16].

3 Wisdom of Crowds

An aspect that we already touched upon is the central role of collective past behavior in understanding what a user is doing and in making an effort to be of help. Albeit such activity is worth in general, using collective behavior to help understanding what is going on is particularly important when the user is engaged in a complex search activity. Gathering data about the decisions taken by a large group of individuals in order to assist the user is one of the fundamental ideas behind the so-called wisdom of crowds [18]: if the group is large enough, and if its members are sufficiently independent from one another, their collective behavior contains at the same time sufficient homogeneity and enough diversity to allow one to extract clear trends to cover most of the possible user's intents and needs. In the Web, wisdom of crowds is an alternative name of Web data mining, in particular Web usage mining.

Using collective information in Information Retrieval is not something new. The basic ranking schemes used in the 1960's use the collective wisdom of people writings (e.g. TF-IDF). Later, in the Web, links were used (e.g. in PageRank), that is the collective wisdom of webmasters (that is not true today as any person can add links thanks to the Web 2.0). Today, we can also use the collective wisdom of users, reflected in weblogs as clicked pages or in query logs as queries and clicked results. Hence, the wisdom of crowds is crucial to rank many kinds of objects (e.g. documents) or to find subgroups of experts in social networks.

This form of large-group knowledge is also at the very heart of Web 2.0 [15], and appears in many different forms (collaborative tagging, blog harvesting, query-log mining etc.). Albeit important, though, one should always have a clear understanding of the limits and possible pitfalls behind this idea; not only the information collected may be biased due to the presence of spam [6], but also (and more simply) because often, rephrasing Gresham's law, "bad information drives out good". Even under optimistic assumptions, mediocrity tends to prevail, thus producing suggestions of low or limited quality; and, in some cases, outdated or blatantly wrong data may outweigh good ones.

Some measure of trust or reliability of users may alleviate this problem; yet, even in the presence of a perfectly dependable source of information, a further point cannot be ignored: employing the wisdom of crowds to bias the results of a query is only worth if the user shares the crowd's values, else it may be more harmful than beneficial. This is the case of the collaborative search approach [17], where like-minded people actions are considered to leverage the search undertaken by an individual of the considered community. Collaborative search is an example of context-aware search, where the considered context is the user's social context. An important aspect of the problem is to understand user's trusted fellows, because adding such information to the result may be very relevant in convincing that result's additions and

corrections are acceptable. Making a user satisfied requires some kind of sentiment analysis [14] to determine what the user really wants; otherwise, departing from the "neutral" result of a query (the one natively produced by the search engine algorithm) could cause angriness or disaffection.

We remark that our observations should not be overstated; in most cases, using defensive algorithms and leveraging on large unbiased datasets is enough to obtain useful signals. Our main point was that helping users is quite difficult and may be risky, and that in some cases helps may have gone too far without really improving their precision and recall relative to the user's query. Finally, introducing a bias in the results based on a "machine driven" and too user independent interpretation of the user's preferences yields an ethical question: should the search engine present the answer a user wants, or should it present the "bare truth"?

4 Contextualization and Personalization

Contextualization and personalization are an essential ingredient of modern search, as suggested also in the previous section. The notion of context in the search task may be referred to several components (search context, document context etc.), among which a central role is played by user's context.

As previously outlined, making a search outcome tailored to a specific user (or group of users) is a difficult although challenging task: how to collect and represent users' preferences is one of the most important aspects, which has been extensively investigated in the last years [11][13]. Personalization has been recently addressed not only as a pool of techniques aimed to improve search by taking advantage of a single user's preferences, but also by exploiting the user social context, as previously outlined with reference to collaborative search [17][19].

One of the main issues in personalization is the construction of users' profiles: their definition is based either on the implicit monitoring of the user's behavior (Web logs, past queries, Web pages copied on the user's desktop, etc.) or on explicit user indications or on both. There are only a few things in the user's profile that don't change, e.g., her/his birth date, being a fan of a given soccer team (or of a kind of music), but the rest may change, e.g., the preferences for restaurants may change depending, for example, on the user's current location. For this reason profiles should be defined as dynamic and adaptive pieces of knowledge, and possibly represented to reflect the multi-dimensional and faceted nature of the multiple user's interests. In fact, as each of us shows multiple personalities (a.k.a. *personas*) when interacting with a search engine, the user profile should be able to represent multiple user preferences, according to the context in which the search task is being executed.

User profiles do represent long term topical interests: it is therefore hard to associate a specific topical preference with random sessions, much in the same way it may be difficult to guess one's preferences based on limited human interactions with short sentences. This is the reason why collecting and representing user's actions should be based on several techniques. Moreover, to better represent users' preferences, the user's cognitive (i.e. topical) context should be enriched with the knowledge of her/his geographic and social context. There is in fact a great convergence on the importance of contextualization in space and time: for instance, if one is searching on

his/her mobile, most likely (s)he is not attaching a complex task (e.g., next summer's trip to Greece), but rather (s)he is trying to localize a close-by service. Then the search answer (i.e. the user-tailored search process) should take the geographic context into account.

The consideration of multiple preference dimensions gives rise to the interesting problem of aggregation: how to aggregate the various contextual components to define the final ranking of the retrieved information items? Also the aggregation process could be driven by user preferences [4].

Personalization and contextualization in fact typically concern not only the results to be presented but also (and more importantly) their presentation and their order. However, the general approach to ranking sees everyone agreeing on the fact that total order is misleading, as users are only concerned by partial orders. Regarding multi-dimensional search, for many users could be easier to work on one dimension at a time; it could be then appropriate to build search protocols that suggest at each stage the results which optimize one dimension; of course, in such scenarios the ordering of dimensions chosen by the protocol becomes essential, but it could also be user-driven.

Another important issue is to go beyond the ranked list presentation of search results; a first step was mixing other results in the standard page result (e.g. universal search in Google), but the problem of combining different types of results in a single screen is mostly unsolved. Some recent and interesting approaches are aimed to help users to visually identify bad and good results through two or three-dimensional presentations of search results, by also taking into account the user's preferences [1].

An important topic is the impact of personalizing the result. To do personalization well we need enough data from the user as well as explicit consent, due to privacy issues (e.g. login authentication that implicitly approves the terms and conditions of the interaction); the overall percentage of users satisfying these two restrictions will be low. This problem is usually faced when developing client-side applications [12]. Despite of these privacy issues, Google personalized search offers an example of centralized solution to the "one size fits all" approach.

On the other hand, contextualizing results (e.g. ranking and displaying the results) according to the intent and context of the query without personalizing at an individual level is another direction that may help groups of people doing the same task regardless of their identity. As log analysis shows, as users perform similar tasks in the Web, this kind of contextualization could have a larger impact. In addition, as we do not need to know the single user, we remove the privacy issues related to personalization [2].

5 Raising the Complexity of the User Interface?

The final question concerns the interaction paradigm: will it always be bound to keyword-based search, or should we expect that the interaction paradigm complexity would rise, at least for complex search, so as to be more expressive? Query interfaces should anyway be kept very simple, featuring interactions where users implicitly choose dimensions as they become available, and where each dimension corresponds to a real-world object, and objects are connected by a graph. This interaction seems to be coherent with the emerging model of "objects-based search" presented by Yahoo

as an evolution of page-based search [2][5]. Interaction is obviously more complicated, although the user does not need to know that (s)he is traversing a graph, (s)he is just presented the valid options at each step. This form of interaction is already partially implemented in the form of query suggestion or as the "find-similar" feature of existing search engines.

Still, this seems a rather radical step for search engines at their current stage of development. There is consensus that, while niche users will enjoy more expressive languages - optimistically one could think of a 10% of sophisticated users – the standard user interface of search engines will not change. It is not in the search company's interest to train people in doing "better search": with millions and millions of new users per year, the average user skill is actually reducing, and search companies are most interested in pleasing & capturing the tail of newcomers.

Moreover, recent studies show that the same user may at time act as "dummy" and at time act as a very clever user [8]. Most people are satisfied of results that they get in the "dummy" mode, so they have little incentives in making more efforts, as the interface simplicity pays off. But a complex multi-domain search could still be engaged with a simple keyword-based interface and simple ways of interaction, so there is hope to see search computing queries at work below current interfaces, and not necessarily with a more expressive interaction paradigm.

On a different direction, a rich trend of research is trying to understand natural-language queries (maybe combined with speech recognition tools, and possibly with a verbal output through a voice synthesizer). This is a promising line of study, especially if you think of its application to mobile-based search, but at the moment it seems to be still immature for large-scale implementation.

Users are so satisfied of the simplicity of search engine interfaces that there must be a huge incentive to convince them to go for a more complex kind of interaction; such a situation may happen in the future, at least in some context, but at the moment it appears to be out of reach. This sets the main challenge of search engine evolution: going much beyond the current answering capability without changing the simple user interface.

References

[1] Ahn, J.-W., Brusilovsky, P.: Adaptive Visualization of Search Results: Bringing User Models to Visual Analytics. Information Visualization 8(3), 167–179 (2009)

[2] Baeza-Yates, R., Raghavan, P.: Next Generation Web Search. In: Ceri, S., Brambilla, M. (eds.) Search Computing. LNCS, vol. 5950, pp. 11–23. Springer, Heidelberg (2010)

[3] Boldi, P., Bonchi, F., Castillo, C., Donato, D., Gionis, A., Vigna, S.: The query-flow graph: model and applications. In: Proceeding of the 17th ACM Conference on Information and Knowledge Management, Napa Valley, California, USA, October 26-30 (2008)

[4] da Costa Pereira, C., Dragoni, M., Pasi, G.: Multidimensional Relevance: A New Aggregation Criterion. In: Boughanem, M., Berrut, C., Mothe, J., Soule-Dupuy, C. (eds.) ECIR 2009. LNCS, vol. 5478, pp. 264–275. Springer, Heidelberg (2009)

[5] Dalvi, N.N., Kumar, R., Pang, B., Ramakrishnan, R., Tomkins, A., Bohannon, P., Keerthi, S., Merugu, S.: A Web of concepts. In: ACM Principles of Database Systems (PODS), pp. 1–12 (2009)

[6] Davison, B., Najork, M., Converse, T.: SIGIR Worksheet Report: Adversarial Information Retrieval on the Web, AIRWeb (2006)
[7] Donato, D., Bonchi, F., Chi, T., Maarek, Y.: Do you want to take notes?: identifying research missions in Yahoo! search pad. In: Proceedings of the 19th International Conference on World Wide Web, Raleigh, North Carolina, USA, April 26-30 (2010)
[8] Goel, S., Broder, A.Z., Gabrilovich, E., Pang, B.: Anatomy of the long tail: ordinary people with extraordinary tastes. In: Proceedings of WSDM 2010, pp. 201–210 (2010)
[9] He, D., Göker, A.: Detecting session boundaries from Web user logs. In: Proceedings of the BCS-IRSG 22nd Annual Colloquium on Information Retrieval Research, Cambridge, UK, pp. 57–66 (2000)
[10] Jones, R., Klinkner, K.L.: Beyond the session timeout: automatic hierarchical segmentation of search topics in query logs. In: Proceeding of the 17th ACM Conference on Information and Knowledge Management, Napa Valley, California, USA, October 26-30 (2008)
[11] Kelly, D., Teevan, J.: Implicit feedback for inferring user preference: A bibliography. SIGIR Forum 37(2), 18–28 (2003)
[12] Kobsa, A.: Privacy-Enhanced Web Personalization. In: Brusilovsky, P., Kobsa, A., Nejdl, W. (eds.) Adaptive Web 2007. LNCS, vol. 4321, pp. 628–670. Springer, Heidelberg (2007)
[13] Micarelli, A., Gasparetti, F., Sciarrone, F., Gauch, S.: Personalized Search on the World Wide Web. In: Brusilovsky, P., Kobsa, A., Nejdl, W. (eds.) Adaptive Web 2007. LNCS, vol. 4321, pp. 195–230. Springer, Heidelberg (2007)
[14] Pang, B., Lee, L.: Opinion Mining and Sentiment Analysis. Found. Trends Inf. Retr. 2, 1–2 (2008)
[15] O'Reilly, T.: What Is Web 2.0: Design Patterns and Business Models for the Next Generation of Software (2010), http://oreilly.com/web2/archive/what-is-web-20.html
[16] Rafiei, D., Bharat, K., Shukla, A.: Diversifying Web search results. In: Proceedings of the 19th International Conference on World Wide Web, Raleigh, North Carolina, USA, April 26-30 (2010)
[17] Smyth, B.: A Community-Based Approach to Personalizing Web Search (2007)
[18] Surowiecki, J.: The wisdom of crowds. Knopf Doubleday Publishing Group (2005)
[19] Teevan, J., Morris, M., Bush, S.: Discovering and using groups to improve personalized search. In: Proceedings of the ACM International Conference on Web Search and Data Mining, pp. 15–24 (2009)
[20] White, R.W., Kules, B., Drucker, S.M., Schraefel, M.C.: Supporting Exploratory Search, Introduction. Communications of the ACM 49(4), 36–39 (2006)