Enhancing Web Search by Promoting Multiple Search Engine Use

Aditya Srivastav ;Tanya Sharma & Karishma Sinha

Abstract

Any given Web search engine may provide higher quality results than others for certain queries. Therefore, it is in users' best interest to utilize multiple search engines. In this paper, we propose and evaluate a framework that maximizes users' search effectiveness by directing them to the engine that yields the best results for the current query. In contrast to prior work on meta-search, we do not advocate for replacement of multiple engines with an aggregate one, but rather facilitate simultaneous use of individual engines. We describe a machine learning approach to supporting switching between search engines and demonstrate its viability at tolerable interruption levels. Our findings have implications for fluid competition between search engines.

1. INTRODUCTION

Web search engines such as Google, Yahoo!, and Live Search provide users with keyword access to Web content. According to statistics aggregated by surveys, users occasionally use multiple search engines, they are typically loyal to a single one even when it may not satisfy their needs, despite the fact that the cost of switching engines is relatively low . While most users appear to be content with their experience on their engine of choice, it is conceivable that many users dislike the inconvenience of adapting to a new engine, may be unaware how to change the default settings in their Web browser to point to a particular engine, or may even be unaware of other Web search engines that exist and may provide better service. Performance differences between

Web search engines may be attributable to ranking algorithms and index size, among other factors. It is well understood in the Information Retrieval (IR) community that different search systems perform well for some queries and poorly for others which suggests that excessive loyalty to a single engine may actually hinder searchers.

To address this problem, this paper describes a machine learning approach that allows users to leverage multiple search engines by unobtrusively recommending the most effective engine for a given query. The approach relies on a classifier to suggest the topperforming engine for a given search query, based on features derived from the query and from the properties of search result pages, such as titles, snippets, and URLs of the top-ranked documents. We seek to promote *supportedsearch engine switching* operations where users are encouraged to temporarily switch to a different search engine for a query on which it can provide better results than their default search engine. Unsupported switching, whereby users navigate to other engines on their own accord, is a phenomenon that may occur for a number of reasons: users may be dissatisfied with search results or the interface, they may be lured to the engine by advertising campaigns or word of mouth, or they may switch by accident. Results of a log-based study that we present in the paper show that only around 10% of search sessions currently involve more than one search engine. We conjecture that by proactively encouraging users to try alternative engines for appropriate queries (hence increasing the fraction of sessions that contain switching)

we can promote more effective user searching for a significant fraction of queries. Empirical results presented in this paper support this claim.

2. RELATED WORK

Prior work in search engine switching has sought to characterize the behavior with the goal of developing metrics for competitive analysis of engines in terms of estimated user preference and user engagement , or switching prediction. Other work has focused on building conceptual and economic models of search engine choice. Telang et al. proposed a qualitative model of search engine choice that is a function of the search engine brand, the loyalty of a user to a particular search engine at a given time, user exposure to banner advertisements, and the likelihood of a within-session switch from the engine to another engine. Mukhopadyay et al. develop an economic model of search engine competition assuming that the switching cost between engines is very low. These studies have focused on understanding and characterizing *existing* switching behaviors in Web search. Although we provide summary statistics on the nature of switching from our observations, our objective is not to characterize switching behavior. Instead, we demonstrate that the utilization of multiple search engines can be advantageous to users and propose a framework that proactively promotes switching.

Commercial meta-search engines such as Clustyand Dogpile attempt to provide access to multiple engines. Given the ranked lists of documents returned by multiple search engines in response to a given query, the objective of meta-search engines is to combine these lists in a way which optimizes the performance of the combination. The IR community has studied meta-search in great detail, with the emphasis on how to merge results from multiple engines, rather than on encouraging people to switch engines as we do in this work. Proactive switching support is an attractive alternative to meta-search for the following reasons: (i) strong brand loyalty may discourage users from migrating to a meta-search engine, (ii) meta-search engines merge search results and obliterate the benefits of interface features of the individual engines, and (iii) meta-searching may be discouraged by search engines as it can negatively impact brand awareness and advertising revenue. We propose an approach whereby users can use their favorite engine but have an alternate engine suggested to them when it is expected to perform better for their current query. In some respects, this is similar to distributed IR , although we are interested in directing users to the best engine rather than the best collection of documents, and do not merge the search results, as is common practice in that sub-discipline.

Supporting engine switching in real-time requires computationally efficient estimation of relative search result quality across several engines. Measuring quality of search results via metrics such as precision and recall has been central in driving research in IR algorithm design, particularly in the Text REtrieval Conference (TREC) community. Hawking et al. employed a methodology similar to TREC to compare the performance of multiple Web search engines. Others, such as Rorvig and CronenTownsend et al. , have looked at techniques for predicting the quality of results using the dispersion of the top documents or computing the entropy between the language model for the results and the collection as a whole. Leskovec et al. used properties of search result sets

projected onto the Web graph to estimate result quality. Despite their effectiveness at computing result quality, some of techniques depend on relevance judgments, meaning that they cannot scale to unseen queries, and some are computationally expensive, meaning that real-time computation is unfeasible. One key distinction of our work from these approaches is that we directly model *relative* quality of multiple search result sets instead of the quality of any individual result set.

Our framework relies on a classifier to estimate the differences in search result quality between the engines using features computed based on the query and the result pages. Yom-Tov et al. have proposed estimating query difficulty using a machine learning approach based on query-only features, validating it for a distributed IR setting with several collections of newswire documents, rather than Web search as we do in this work. Caption features have already been shown to be important to users in determining which search results to select , and query-caption features have been used in the development of ranking algorithms to improve search . As our empirical results demonstrate, utilizing multiple diverse feature sources is beneficial over query-only features, and is a key performance differentiator for accurate prediction of the most appropriate search engine for a given query in real-time.

3. THE CASE FOR MULTI-ENGINE USE

At the outset of our studies, we pursued general statistical clues that could provide insight into the extent to which users switched engines and the potential benefit to them of switching engines. To do so, we used the interaction logs of a large sample of consenting Web users. We begin by describing the statistical properties of search sessions extracted from the logs.

3.1 Search Sessions

We used the interaction logs of over five million consenting Web users over a five-month period from May 2007 to September 2007. These logs were anonymized, and all personally identifiable information, including IP addresses, was removed. The logs gave us access to user interactions with all search engines. From these logs, we extracted *searchsessions* that began with a query to benefit of supporting switching. Google, Yahoo!, or Live Search and terminated after 30 minutes of browsing inactivity. A similar threshold has been used to demarcate search sessions in previous work on search engine switching [16] and in related studies of user search behavior [20,26]. These sessions are used to analyze switching behavior and give insight into the potential.

3.2 Overview of Switching Behavior

Our analysis showed that 36.4% of searchers used more than one search engine in the duration of the logs.[7] The findings also showed that 6.8% of all sessions and 12.0% of sessions containing more than one query involved a switch between two or more search engines. Although the aim of the paper is not to characterize the nature of search engine switching, a visual examination of search engine usage patterns in the logs revealed three salient classes of switching behavior: *within-session*, *between-session*, and *long-term*. We now describe these classes and provide summary statistics:

- Within-session switching: Users switch between Web search engines within a single search session and may use multiple engines concurrently. Such switches may be associated with a desire

for topic coverage, dissatisfaction with any particular engine, and perhaps even automated applications that issue queries to multiple engines. Approximately 33.4% of the users in our sample exhibited this class of behavior.

- Between-session switching : Users switch engines for individual search sessions or groups of sessions. Switches of this nature may occur because a user feels that a particular engine is better suited for the current task due to an interface component or vertical supported. Approximately 13.2% of the users in our sample exhibited this type of switching behavior.

- Long-term switching: Users switch from one search engine to another and never return to the original engine. This appears to represent

  a change in their search engine preference. Approximately 7.6% of the users in our sample switched search engines and never

  returned to their original engine in the duration of the study.

Of these three classes, our component aims to support within session switches, where it might be in a user's interest to change search engines for the current query. While the above statistics demonstrate that search engine switching is a strategy employed by some users, the majority of users remain loyal to a single engine. Prior to describing our method for supporting search engine switching, the next section analyzes the potential benefit to users brought by utilizing multiple search engines.

3.3 Potential Benefit of Switching

To motivate our approach, we first quantify the potential benefit of multiple search engine use. That is, if a user is searching on a given engine, what is the likelihood that they would obtain better quality results if they were to issue the same query on a different engine. This is important, since encouraging users to switch when it is not in their interests to do so could lead to user dissatisfaction and ultimately distrust for our classifier

4. CONCLUSIONS

In this paper, we advocated for the use of multiple

Search engines to empower users to search more effectively. We described a log based study of Web search behavior with a particular emphasis on multiple search engine use, which demonstrated that search engine switching can substantially improve retrieval effectiveness. We proposed a machine learning-based approach for supporting switching that estimates in real time whether more accurate results exist on alternate search engines. Estimation is based on features of the query, the result set, and the titles, snippets, and URLs of the top-ranked search results. An empirical analysis of classification performance demonstrates that it is accurate at predicting when users would benefit from switching between engines at low recall levels. The promotion of multiple search engine use through application components such as that described has the potential to improve the retrieval experience for users of all search engines.

5. REFERENCES

[1] Buckley, C. & Walz, J. (1999). Overview of the TREC-8 query track. In Voorhees, E.M. & Harman, D.K. (Eds.) *Proc. TREC-8*, 65-76.

[2] Budzik, J. & Hammond, K.J. (2000). Users interactions with everyday applications as context for just-in-time information access. In *Proc. IUI*, 44-51.

[3] Pew Internet & American Life Project. (2005). *Search Engine Users*. Accessed October 16, 2007. Available at: http://www.pewinternet.org/pdfs/PIP_Searchengine_users.pdf

[4] Pitkow, J.E. & Pirolli, P. (1999). Mining longest repeating subsequences to predict world wide web surfing. In *Proc. USENIX Symposium*, 139-150.

[5] Rasolofo, Y., Abbaci, F. & Savoy, J. (2001). Approaches for collection selection and results merging for distributed information retrieval. In *Proc. CIKM*, 191-198.

[6] Rorvig, M. (2000). A new method of measurement for question difficulty. In *Proc. ASIS*, 372–378.

[7] Si, L. & Callan, J. (2003). A semi-supervised learning method to merge search engine results. *TOIS*, *21*(4). 457-491.

[8] Telang, R., Mukhopadhyay, T. & Wilcox, R. (1999). An empirical analysis of the antecedents of internet search engine choice. In *Proc. Workshop on Information Systems and Economics* (WISE, Charlotte NC).