A Study of Distinctiveness in Web Results of Two Search Engines

Rakesh Agrawal*
Data Insights Laboratories
ragrawal@acm.org

Behzad Golshan*
Boston University
behzad@bu.edu

Evangelos Papalexakis*
Carnegie Mellon University
epapalex@cs.cmu.edu

ABSTRACT

Google and Bing have emerged as the diarchy that arbitrates what documents are seen by Web searchers, particularly those desiring English language documents. We seek to study how distinctive are the top results presented to the users by the two search engines. A recent eye-tracking has shown that the web searchers decide whether to look at a document primarily based on the snippet and secondarily on the title of the document on the web search result page, and rarely based on the URL of the document. Given that the snippet and title generated by different search engines for the same document are often syntactically different, we first develop tools appropriate for conducting this study. Our empirical evaluation using these tools shows a surprising agreement in the results produced by the two engines for a wide variety of queries used in our study. Thus, this study raises the open question whether it is feasible to design a search engine that would produce results distinct from those produced by Google and Bing that the users will find helpful.

Categories and Subject Descriptors

H.2.8 [**Database Applications**]: Data mining; H.3.3 [**Information Search and Retrieval**]: Search process

Keywords

Web search; search engine; search result comparison; Google; Bing

1. INTRODUCTION

The World Wide Web is now widely recognized as the universal information source. The fairness doctrine enunciated several decades ago contends that citizens should have access to diverse perspectives [12]. The normative impetus behind this doctrine is the idea that exposure to different views is beneficial for citizens. Without question, content representing diverse perspectives exist on the Web almost on any topic However, this does not automatically guarantee that audiences encounter them [31].

---

*Work partially done at the erstwhile Microsoft Research in Silicon Valley. Author order is alphabetic.

Search engines have become the dominant tool used to access the web content [27]. In the physical world, one way people gain access to diverse perspectives is by subscribing to different newspapers, listening to different radio stations, tuning into different television channels, or manually selecting different publications and books. We seek to study whether users can garner different perspectives by obtaining results for the same query from different search engines. For this purpose, we study how distinctive are the web search results produced by Google and Bing - the two most popular search engines of English language documents (Yahoo's web search is currently powered by Bing).

In addition to the information about the documents that the search engine deems most relevant to the query (the so called "organic results"), a search engine result page (SERP) often contains a variety of other information. This may include inter alia sponsored listings, images, videos, maps, definitions, or suggested search refinements. We focus on comparing the top-10 organic results on the first SERP because they are the ones that get almost all of the clicks [11]. Users unsatisfied with the top results frequently retype a query, instead of looking at results at lower positions [14]. An organic result normally includes the title of the document, a snippet of the document, and URL of the full version.

A recent eye-tracking has shown that the web searchers decide whether to look at a document primarily based on the snippet and secondarily on the title of the document on the web search result page, and rarely based on the URL of the document [25]. Given that the snippet and title generated by different search engines for the same document are often different, we first develop tools appropriate for conducting this study. We then use these tools to study the extent of agreement in the results produced by the two engines for a wide variety of queries.

Contributions

In this work, our main contribution is quantifying how distinctive are the organic search results produced by Google and Bing. In order to achieve that, we also make the following technical contributions:[1]

---

[1] While designed to effectively analyze search engine results, our tools have broader applicability. For instance, consider a set of questions, possibly coming from a Massive Open Online Course (MOOC) exam; these questions can either be multiple choice or in free-text form. In the setting of a MOOC, there will be potentially hundreds, or thousands, of students responding to those questions. There are also multiple exams, as well as multiple MOOCs on the same subject, offered by different providers. Using these tools, we are able to quantify the similarity of students across different MOOCs, as well as similarity of MOOCs in terms of how students respond to exam questions (which could be an indicator of how well students learn from a particular MOOC).

Visualization and exploratory analysis: We introduce TENSOR-COMPARE, an exploratory tool for visualizing and analyzing pairwise differences between search engines.

Quantitative comparison of search engines results: We also introduce CROSSLEARNCOMPARE, a tool that uses machine learning and quantifies the similarity of results between two search engines by framing it as a prediction problem.

Paper Layout

The structure of the rest of the paper is as follows. We begin by discussing related work in Section 2. We then describe the new tools we designed for carrying out the comparative study in Section 3. Section 4 presents the empirical evaluation. We conclude with a discussion of the significance of the work and future directions in Section 5.

2. RELATED WORK

More than four decades ago, Lancaster and Fayen [19] in 1973 listed six criteria for assessing the performance of information retrieval systems: 1) Coverage, 2) Recall, 3) Precision, 4) Response time, 5) User effort, and 6) Form of output. Since the advent of search engines in early 90's, there are several reported studies that evaluated their performance on one or more these criteria. See [9] and references therein for examples of some early studies. See [23] for a recent compilation of various issues and studies related to the evaluation of web search engines. We will focus our discussion on prior works that studied the overlap of results between different search engines, the thrust of our paper.

An early overlap study is due to Ding and Marchionini, who measured the result overlap between the then popular three search engines: InfoSeek, Lycos, and OpenText. Five queries were used to conduct searches with these services. They observed a low level of result overlap among the services [10]. Around the same time, Selberg and Etzioni found, in the context of their metacrawler work, that none of Galaxy, Infoseek, Lycos, OpenText, Webcrawler and Yahoo was able to return more then 45% of the references followed by users. They also observed that each of the engines returned mostly unique results [28]. Also in 1996, Gauch, Wang and Gomez found that a metasearch engine that fused the results of Alta Vista, Excite, InfoSeek, Lycos, Open Text, and WebCrawler provided the highest number of relevant results [13].

Bharat and Broder estimated the size of the Web to be 200 million pages in November 1997 and the overlap between the websites indexed by HotBot, Alta Vista, Excite and InfoSeek to be only 1.4% [5]. Lawrence and Giles published their study of AltaVista, Excite, HotBot, Infoseek, Lycos, and Northern Light in 1998. They found that the individual engines covered from 3 to 34% of the indexable Web, based on their estimate of the size of the Web at 320 million pages. Combining the six engines in their study covered about 3.5 times as much of the Web as one engine [20].

Fast forwarding a bit, Gulli and Signorini estimated that by January 2005 the indexable Web had increased in size to about 11.5 billion pages and that Google's coverage rate was 76.2%, Yahoo's 69.3% and that of MSN Search (predecessor of Bing) 61.9% [15]. Spink et al. studied the overlap between the results of four search engines, namely MSN, Google, Yahoo and Ask Jeeves, using data from July 2005. Their findings showed that the percent of total first page results unique to only one of the engines was 84.9%, shared by two of the three was 11.4%, shared by three was 2.6%, and shared by all four was 1.1% [29]. In an update two years later, they noted that the first page results of the four engines continued to differ from one another and in fact they included fewer results in common in 2007 than in 2005 [30].

More recently, Pirkola investigated how effectively the websites of Finnish, French, and U.S. domains were being indexed by two US-based and three Europe-based search engines [26]. The results showed that Google and Live Search (predecessor of Bing) indexed US sites more effectively than Finnish and French sites, the Finnish www.fi indexed only Finnish sites and the French Voila only French sites, and the European engine Virgilio indexed European sites more effectively than US sites. In another interesting study, Wilkinson and Thelwall compared the results of seventeen random queries submitted to Bing for thirteen different English geographic search markets at monthly intervals [32]. They found there were almost no ubiquitous authoritative results: only one URL was always returned in the top-10 for all search markets and points in time and that results from at least three markets needed to be combined to give comprehensive results. There also have been studies pointing out that the search engine results are not stable even in short windows of time [3, 23].

We did not find much discussion in prior work of the techniques used for determining if two result pages contained links to the same web document. For example, [29, 30] simply state that this determination is done using string comparison of URLs. It is not clear what URL normalization [21, 22], if any, was done before string comparison. It is also not clear what, if anything, was done to address the problem of DUST - Different URLs with Similar Text [4]. Finally, there is no mention of short URLs, although the first notable URL shortening service, namely tinyURL, dates back to 2002 [1].

To summarize, all of prior work found little overlap between the first page results produced by different engines for very many queries. Some plausible reasons have also been put forward for this low overlap. They include that the search engines are constrained in the portions of the Web they index due to network bandwidth, disk storage, computational power, or a combination of these items. Search engines use different technologies to find pages and indexing them. And they deploy proprietary algorithms to determine the ranking of the results and their presentation to the users. Fingers have also been pointed at implicit personalization [16].

Why another study?

Given the rich prior literature we have outlined, it is natural to question the need for a new study on the overlap between the search engine results. We believe that much has changed in the recent times in the search engine market landscape and in search engine technologies to warrant such a study. With Bing powering Yahoo search, we now essentially have a diarchy in Google and Bing that arbitrates user access to the English language Web that a very large fraction of humanity accesses on daily basis to get information. But there is no recent comparative study of Google and Bing search results. It is imperative to periodically analyze what people are able to see and read. Such studies also lead to the creation of new analysis tools and the questioning of conventional wisdom, thus contributing to the advancement of science.

3. ANALYTICAL TOOLS

We designed two tools to be able to analyze and compare search engine results. One, which we call TensorCompare, uses tensor analysis to derive low-dimensional compact representation of search results and study their behavior over time. The other, which we call CrossLearnCompare, uses cross-engine learning to quantify their similarity. We discuss them next.

3.1 TensorCompare

Postulate that we have the search results of executing a fixed set of queries at certain fixed time intervals on the same set of search

engines. These results can be represented in a four mode tensor $\underline{\mathbf{X}}$, where (query, result, time, search engine) are the four modes [18]. A result might be in the form of a set of URLs or a set of keywords representing the corresponding pages. The tensor might be binary valued or real valued (indicating, for instance, frequencies).

This tensor can be analyzed using the so-called canonical or PARAFAC decomposition [17], which decomposes the tensor into a sum of rank-one tensors: $\underline{\mathbf{X}} \approx \sum_{r=1}^{R} \lambda_r\, \mathbf{a}_r \circ \mathbf{b}_r \circ \mathbf{c}_r \circ \mathbf{d}_r$, where the (i, j, k, l)-th element of $\mathbf{a} \circ \mathbf{b} \circ \mathbf{c} \circ \mathbf{d}$ is simply $\mathbf{a}(i)\mathbf{b}(j)\mathbf{c}(k)\mathbf{d}(l)$.

The vectors $\mathbf{a}_r, \mathbf{b}_r, \mathbf{c}_r, \mathbf{d}_r$ are usually normalized, with their scaling absorbed in λ_r. For compactness, the decomposition is represented as matrices $\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}$. The decomposition of $\underline{\mathbf{X}}$ to $\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}$ gives us a low rank embedding of queries, results, timings, and search engines respectively, corresponding to the aforementioned clusters. This implies that we are able to track the temporal behavior of a cluster of semantically similar search engines for a set of queries.

The factor matrix $\mathbf{D}$ projects each one of the search engines to the R-dimensional space. Alternatively, one can view this embedding as soft clustering of the search engines, with matrix $\mathbf{D}$ being the cluster indicator matrix: the (i, j) entry of $\mathbf{D}$ shows the participation of search engine i in cluster j.

This leads to a powerful visualization tool that captures similarities and differences between the search engines in an intuitive way. Say we take search engines A and B and the corresponding rows of matrix $\mathbf{D}$. If we plot these two row vectors against each other, the resulting plot will contain as many points as clusters (R in our particular notation). The positions of these points are the key to understanding the similarity between search engines.

Figure 1 serves as a guide. The (x, y) coordinate of a point on the plot corresponds to the degree of participation of search engines A and B respectively in that cluster. If all points lie on the 45 degree line, this means that both A and B participate equally in all clusters. In other words, they tend to cluster in the exact same way for semantically similar results and for specific periods of time. Therefore, Fig. 1(a) paints the picture of two search engines that are very (if not perfectly) similar with respect to their responses. In the case where we have only two search engines, perfect alignment of their results in a cluster would be the point $(0.5, 0.5)$. If we are comparing more than two search engines, then we may have points on the lower parts of the diagonal. In the figure, we show multiple points along the diagonal for the sake of generality.

Figure 1(b), on the other hand, shows the opposite behavior. Whenever a point lies on either axis, this means that only one of the search engines participate in that cluster. If we see a plot similar to this figure, we can infer that A and B are very dissimilar with respect to their responses. In the case of two search engines, the only valid points on either axis are $(0, 1)$ and $(1, 0)$, indicating an exclusive set of results. However, for generality, we show multiple points on each axis.

Note, of course, the cases shown in Fig. 1 are the two extremes, and we expect to observe behaviors bounded by those extremes. For instance, in the case of two search engines, all points should lie on the line $\mathbf{D}(1, j)x + \mathbf{D}(2, j)y = 1$, where $\mathbf{D}(1, j)$ is the membership of engine A in cluster j, and $\mathbf{D}(2, j)$ is the membership of engine B in cluster j. This line is the dashed line of Fig. 1(a).

TENSORCOMPARE also allows us to track the behavior of clusters over time. In particular, given the i-th group of semantically similar (query, result, search engine) cluster, as given by the decomposition, the i-th column of matrix $\mathbf{C}$ holds the temporal profile


Figure 1: Visualization guide for TENSORCOMPARE.

of that cluster. Suppose we have T days worth of measurements. If the search engines of that cluster produce similar results for the given set of queries for all T, the temporal profile will be approximately constant and each value will be approximately equal to $\frac{1}{T}$. Otherwise, there will be variation in the profile, correlated with the variation of the particular results. In the extreme case where a result appeared only on a single day, the time profile will have the value approximately equal to one corresponding to that day, and approximately zero for the rest of the days.

Theoretical Foundation

We next provide a Lemma that connects the plots provided by TENSORCOMPARE to the degree of semantic overlap of two search engines. Suppose that for a given cluster j, we denote the membership of search engine A as $x = D(A, j)$ and the membership of search engine B as $y = D(B, j)$. For ease of exposition, consider the case of two search engines and assume that we have a three mode tensor: (query, result, search engine).

LEMMA 1. *Assume a binary (query, result, search engine) tensor that has exactly one rank one component. Let search engine A correspond to the x coordinate, and search engine B correspond to the y coordinate of a TENSORCOMPARE plot. For the particular component, if search engine B has p_1 fraction of queries in common with A, and p_2 portion of the result in common with A, then*

$$y \leq p_1 p_2 x.$$

PROOF. See Appendix. □

In the case of a four-mode tensor, with p_3 percent overlap in the time mode, the bound is $y \leq p_1 p_2 p_3 x$. The above Lemma provides an upper bound, however, we experimentally validated that this bound is in practice tight.

3.2 CrossLearnCompare

An intuitive measure of the similarity of the results of two search engines is the predictability of the results of a search engine given the results of the other. Say we view each query as a class label. We can then go ahead and learn a classifier that maps the search result of search engine A to its class label, i.e. the query that produced the result. Imagine now that we have results that were produced by search engine B. If A and B return completely different results, then we would expect that classifying correctly a result of B using the classifier learned using A's results would be difficult, and our classifier would probably err. On the other hand, if A and B returned almost identical results, classifying correctly the search results of B would be easy. In cases in between, where A and B bear some level of similarity, we would expect our classifier to perform in a way that it is correlated with the degree of similarity between A and B.

Note we can have different accuracy when predicting search engine A using a model trained on B, and vice versa. This, for instance, can be the case when the results of A are a superset of the

Algorithm 1: CROSSLEARNCOMPARE

Input: $\mathcal{R}_A, \mathcal{R}_B$ are instances of results of engines A and B. Each instance is in the form (query, result representation in chosen feature space)

Output: Similarity measures $c_{A,B}$ and $c_{B,A}$ between search engines A, B.

1: Train a model $\mathcal{M}_A$ based on the instances $\mathcal{R}_A$, using the query as a class label.
2: Train a model $\mathcal{M}_B$ based on the instances $\mathcal{R}_B$, using the query as a class label.
3: For all instances in $\mathcal{R}_B$, use $\mathbf{M}_A$ to predict the query. Set $c_{A,B}$ as a measure of the classifier's accuracy (e.g. Area Under the Curve).
4: For all instances in $\mathcal{R}_A$, use $\mathbf{M}_B$ to predict the query. Set $c_{B,A}$ likewise.

---

results of B. Algorithm 1 shows an outline of CROSSLEARNCOMPARE.

4. EMPIRICAL EVALUATION

We now present the results of the empirical study we performed, applying the tools just described on the search results from Google and Bing for a wide variety of queries.

4.1 Data Set

We conducted the evaluation for two sets of queries. The TRENDS set (Table 1) contains the most popular search terms from different categories from Google Trends during April 2014. We will refer to them as *head queries*. The MANUAL set (Table 2) consists of hand-picked queries by the authors that we will refer to as *trunk queries*. These queries consist of topics that the authors were familiar with and were following at the time. Familiarity with the queries is helpful in understanding whether two sets of results are different and useful. Queries in both the sets primarily have the informational intent [6]. The total number of queries was limited by the budget available for the study.

Albert Einstein	American Idol	Antibiotics	Ariana Grande
Avicii	Barack Obama	Beyonce	Cristiano Ronaldo
Derek Jeter	Donald Sterling	Floyd Mayweather	Ford Mustang
Frozen	Game of Thrones	Harvard University	Honda
Jay-Z	LeBron James	Lego	Los Angeles Clippers
Martini	Maya Angelou	Miami Heat	Miami Heat
Miley Cyrus	New York City	New York Yankees	Oprah Winfrey
San Antonio Spurs	Skrillex	SpongeBob SquarePants	Tottenham Hotspur F.C.
US Senate			

Table 1: TRENDS queries

Afghanistan	Alternative energy	Athens	Beatles	Beer
Coup	Debt	Disaster	E-cigarettes	Education
Gay marriage	Globalization	Gun control	IMF	iPhone
Iran	Lumia	Malaria	Merkel	Modi
Paris	Polio	Poverty	Rome	Russia
San Francisco	Self-driving car	Syria	Tesla	Ukraine
Veteran affairs	World bank	World cup	Xi Jinping	Yosemite

Table 2: MANUAL queries

We probed the search engines with the same set of queries at the same time of the day for a period 21 days for the TRENDS set, and 17 days for the MANUAL set, during June-July 2014. For Google, we used their *custom search API* (code.google.com/apis/console), and for Bing their *search API* (datamarket.azure.com/dataset/bing/search). For both, we recorded the top-k results. The value of k is set to 10 by default, except in the experiments studying the sensitivity of results to the value of k. Every time, we ran the same code from the same machine having the same IP address to minimize noise in the results. Because we were getting the results programmatically through the API, no cookies were used and there was no browser information used by Google or Bing in producing the results [16].

4.2 Representation of Search Results

While our methodology is independent of the specific representation of search results, we employ the snippets of the search results provided by the search engines for this purpose. The snippet of a search result embodies the search engine's semantic understanding of the corresponding document with respect to the given query. The users also heavily weigh the snippet in deciding whether to click on a search result [25]. The alternative of using URL representation must first address the well-known problems arising from short URLs [1], un-nomalized URLs [21, 22], and different URLs with similar text [4]. Unfortunately, there is no agreed upon way to address them and the specific algorithms deployed can have large impact on the conclusions. Furthermore, the users rarely decide whether to look at a document based on the URL they see on the search result page [25].

More in detail, for a given result of a particular query, on a given date, we take the bag-of-words representation of the snippet, after eliminating stopwords. Subsequently, a set of results from a particular search engine, for a given query, is simply the union of the respective bag-of-words representations. For TENSORCOMPARE, we keep all words and their frequencies; 0/1 features did not change the trends. For CROSSLEARNCOMPARE, we keep the top-n words and have binary features. Finally, the distribution of the snippet lengths for Google and Bing was almost identical for all the queries. This ensures a fair comparison between the two engines.

To assess whether snippets are appropriate for comparing the search results, we conducted the following experiment. We inspect the top result given by Google and Bing for a single day, for each of the queries in both TRENDS and MANUAL datasets. If for a query, the top result points to the same content, we assign the URL similarity score of 1 to this query, and the score of 0 otherwise. We then compute the cosine similarity between the bag-of-word representations of the snippets produced by the two search engines for the same query. Figure 2 shows the outcome of this experiment. Each point in this figure corresponds to one query and plots the URL and snippet similarity scores for this query. For clarity, the X and Y axes show ranges beyond [0,1].


(a) TRENDS query set   (b) MANUAL query set

Figure 2: Comparing URL similarity with snippet similarity

We see that for most of the queries for which the snippet similarity was low, the results pointed to different documents; on the other hand, when the similarity of snippets is high, the documents are identical. In both TRENDS and MANUAL, there exist some outliers with pointers to identical documents yet dissimilar snippets (e.g. the query Tesla in TRENDS and US Senate in MANUAL). Yet, overall, Fig. 2 indicates that snippets are good instruments for content comparison.

Note that we do not consider their ordering in our representation of the search results. Instead, we study the sensitivity of our conclusions to the number of top results, including top-1, top-3, and top-5 (in addition to top-10).

4.3 Results of TensorCompare

Having chosen a bag-of-words representation for the results, the input tensor to TENSORCOMPARE has modes (query, term, date, search engine). Our data collection results in a $32 \times 36631 \times 21 \times 2$ tensor for the TRENDS dataset and a $35 \times 39725 \times 17 \times 2$ tensor for the MANUAL set. For fitting the PARAFAC decomposition, we use the algorithm from [8] that is appropriate for sparse, count data. More specifically, we use Tensor Toolbox from Matlab [2], which contains an efficient implementation of this algorithm. The number of components we chose was $R = 20$; however, qualitatively similar behavior was observed for various values for R. The results of TENSORCOMPARE analysis are shown in Figs. 3 and 4. Figure 3 shows the similarity of search results, while Fig. 4 shows the temporal profile of each one of the points in Fig. 3.

(a) TRENDS query set   (b) MANUAL query set

Figure 3: Visualization of TENSORCOMPARE for Google and Bing. Values on the x-axis correspond to the membership of Google to a cluster, and values on the y-axis correspond to the membership of Bing. Thus, an (x, y) point on this plot represents one of the clusters of TENSORCOMPARE. The closer the points are to the 45-degree line, the more similar are the two search engines.

(a) TRENDS query set   (b) MANUAL query set
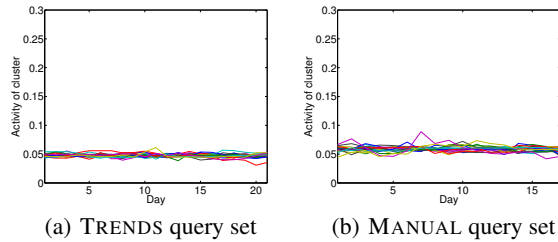
Figure 4: Temporal profile of latent clusters given by TENSORCOMPARE, for Google and Bing. The y-axis corresponds to the membership of the particular day to the cluster of interest. For both query sets, the temporal profile of all clusters is approximately constant over time. In particular, each value for TRENDS is $\approx 1/21$ and for MANUAL it is $\approx 1/17$. As stated in Section 3.1, this indicates that both Bing and Google returned persistent results, at least during the duration of our experiment. Due to this uniformity, we overplot all clusters, without making any distinctions.

The first, immediate, observation is that the latent clusters for both query sets behave very similarly. This fact is encouraging because it shows that our analysis can be applied to both head and trunk queries. In order to interpret the aforementioned plots, we consult Fig. 1. We observe that Google and Bing produce similar results. This is indicated by the fact that in Fig. 3, the majority of the points lie around the $(0.5, 0.5)$ point (we remind the reader that this point indicates almost exact similarity for the case of two search engines), showing near equal participation of Google and Bing to the majority of the latent clusters. This finding is quite surprising and is in sharp contrast with the past studies. We further observe that that there are somewhat more results unique to Google

than Bing since there are more clusters where Google has single participation.

Finally, with respect to the temporal variation of the results, as indicated by Fig. 4, the temporal profile of each cluster is almost uniform across time. This, consequently, means that for both search engines, either in cases where they agree or in cases where they produce somewhat distinct results, their behavior is stable over time, at least as observed during the duration of our study.

4.4 Results of CrossLearnCompare

We next present our analysis of the application of CROSSLEARN-COMPARE to the search results of two engines. To obtain feature space for our instances, we remove terms that are verbatim equal to or contain the query string and then take the 100 highest frequency words for each search engine. We use the union of these two bags of words as the feature space of the training and testing instances. Each such instance is, thus, the vector space representation of a result for a given date and position in the result-set. We use a binary representation, where 1 indicates that the corresponding word appears in the particular instance.

We train one-vs-all linear SVM classifiers for each query set, for each search engine. The performance of the two classifiers of CROSSLEARNCOMPARE for the two query sets is shown in Fig. 5; the measure of performance used is the standard Receiver Operating Characteristic (ROC) curve [7]. There are four curves on the same figure, showing the performance of predicting Bing using Google and vice versa, and for query sets TRENDS and MANUAL. Table 3 contains the Area Under the Curve (AUC) for the ROC curves shown in Fig. 5.
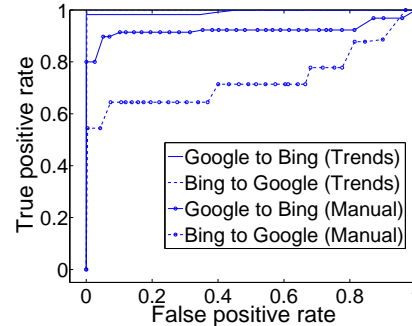
Figure 5: ROC curves produced by CrossLearnCompare (higher is better in terms of classification accuracy). If two search engines were completely mutually predictable, the ROC curve would be exactly on the $(0, 0) - (0, 1)$ and $(0, 1) - (1, 0)$ lines. Conversely, if two search engines were completely mutually unpredictable, the ROC curve would lie on the $(0, 0) - (1, 0)$ and $(1, 0) - (1, 1)$ lines. Finally, when the classifier is random, the curve would lie on the 45-degree line.

	TRENDS $\rightarrow$	TRENDS $\leftarrow$	MANUAL $\rightarrow$	MANUAL $\leftarrow$
Google- Bing	0.993	0.999	0.922	0.730

Table 3: Area Under the Curve (AUC) results for CROSSLEARNCOMPARE. The right arrow $\rightarrow$ indicates that we use the left search engine to predict the right one, and $\leftarrow$ the converse.

Firstly, we observe that the search results are mutually highly predictable for the TRENDS query set. This implies that the top results for these popular queries for Google and Bing are very similar. The same behavior continues to be observed for the MANUAL query set, albeit Google results are somewhat less predictable from Bing results.

4.5 Sensitivity Analysis

One might wonder how sensitive are our conclusions to the fact that we analyzed the top-10 search results. To this end, we apply TENSORCOMPARE and CROSSLEARNCOMPARE to the top-5, top-3, and top-1 search results, for both TRENDS and MANUAL query sets. Figures 6 and 7 show the results of this analysis for top-5 and top-1 for TENSORCOMPARE and CROSSLEARNCOMPARE respectively. The results for top-3 lie between top-5 and top-1 and have been omitted.
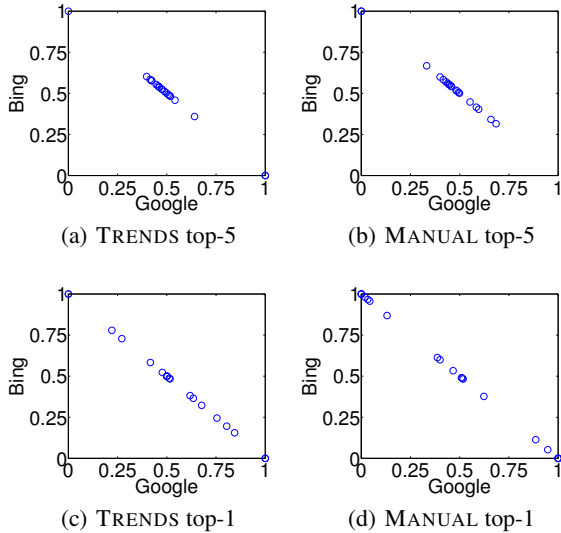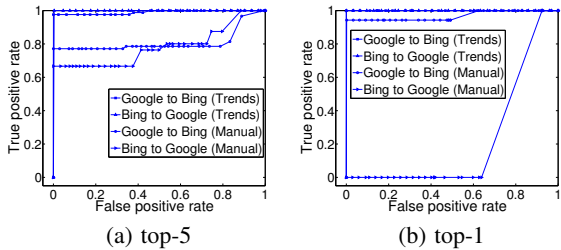

Figure 6: TENSORCOMPARE sensitivity

Figure 7: CROSSLEARNCOMPARE sensitivity

We see that our earlier findings are robust and consistent with the ones presented here. A few specific remarks follow:

- The two search engines continue to exhibit more similar results for the TRENDS query set (head queries) than the MANUAL set (trunk queries).

- Using top-5 as the cut-off, the similarity is slightly higher than using top-10. This indicates that it is more likely that the search engines will have an exclusive result below position 5.

- For the single top result, even though there is similarity, the top result is not necessarily the same (but the manual inspection reveals that the top result of one is almost always present in the top-5 of the other).

- The results of CROSSLEARNCOMPARE reinforce the findings of TENSORCOMPARE. For top-5, the classifier learned using the results of one search engine is able to quite accurately predict the results of the other. However, given the sensitivity of the top result to the position, the performance degrades for top-1.

5. SUMMARY AND FUTURE WORK

We introduced two novel tools for studying the similarity and distinctiveness of web results of search engines. Our main observation, stemming from our analysis, is that Google and Bing exhibited a significant degree of similarity in their search results in our data set. This observation is in sharp contrast to the prior published work where minimal overlap is reported. A fair interpretation of our observation is stating that the visual experience of users in Google and Bing is very similar for the queries we studied. This can be seen as an upper bound to the exact number of overlapping results, which the prior work is trying to estimate.

We can only speculate why there is greater convergence in the results produced by the two search engines. They include deployment of greater amount of resources by search engines to cover a larger fraction of the indexable Web, much more universal understanding of search engine technologies, and the use of similar features in ranking the search results.

In the future, we would like to explore the feasibility of building a search engine that uses signals not used by Google and Bing and yet produces useful results. Such signals could possibly come from social networks. It will also be interesting to explore the practicality of explicitly designing-in diversity into search engines [24].

APPENDIX

Proof of Lemma 1

Figure 8: The two slices of $\mathbf{X}$.

PROOF. Consider a tensor $\underline{\mathbf{X}}$ with dimensions $I \times J \times 2$ (in our case, the first mode corresponds to queries, the second to results, and the third to search engines). Assume that $\underline{\mathbf{X}}$ is rank one, which means that there is one component in its PARAFAC decomposition. In the frontal slice corresponding to the first search engine (Slice 1 in Fig. 8), we have Q queries and T results forming a perfect block, which we assume to be filled with 1's. The second slice, which corresponds to the second search engine, has a block that spans only a fraction of the queries and results of Slice 1.

We assume that the components $\mathbf{a}$, $\mathbf{b}$ of the PARAFAC decomposition are normalized by their ℓ_2 norm, and the scaling is absorbed in $\mathbf{c}$. We further assume that the components are non-negative.

Let $\hat{\mathbf{a}}$, $\hat{\mathbf{b}}$, $\hat{\mathbf{c}}$ be the optimal solution. An upper bound $\mathbf{a}, \mathbf{b}, \mathbf{c}$ to the optimal is the following: The first Q elements of $\mathbf{a}$ will be equal to $\frac{1}{\sqrt{Q}}$ (the rest are zero), and the first T elements of $\mathbf{b}$ will equal $\frac{1}{\sqrt{T}}$. This implies that the coefficients of $\mathbf{c} = \begin{bmatrix} c_1 & c_2 \end{bmatrix}^2$, which multiply $\mathbf{a}\mathbf{b}^T$ in order to approximate the respective slices of $\underline{\mathbf{X}}$, will be proportional to the respective densities of the blocks in either slice, i.e. $c_1 \propto d_1$ and $c_2 \propto d_2$ Making this uniformity assumption for the non-zero elements of $\mathbf{a}, \mathbf{b}$ allows us to bound the ratio of the coefficients of $\hat{\mathbf{c}}$ by the ratio of the densities of the blocks in each slice. More specifically, we have

$$\frac{\hat{c}_1}{\hat{c}_2} \leq \frac{d_1}{d_2} = \frac{QT}{p_1 Q p_2 T} = \frac{1}{p_1 p_2}.$$

Hence, $\hat{c}_2 \leq p_1 p_2 \hat{c}_1$. If we substitute $y = \hat{c}_2$ and $x = \hat{c}_1$, as they correspond in Fig. 1, then we have shown the desired upper bound. □

A. REFERENCES

[1] D. Antoniades, I. Polakis, G. Kontaxis, E. Athanasopoulos, S. Ioannidis, E. P. Markatos, and T. Karagiannis. we.b: The web of short URLs. In *20th international conference on World Wide Web*, pages 715–724. ACM, 2011.

[2] B. W. Bader and T. G. Kolda. Matlab tensor toolbox version 2.2. *Albuquerque, NM, USA: Sandia National Laboratories*, 2007.

[3] J. Bar-Ilan. Search engine ability to cope with the changing web. In *Web dynamics*, pages 195–215. Springer, 2004.

[4] Z. Bar-Yossef, I. Keidar, and U. Schonfeld. Do not crawl in the DUST: different urls with similar text. *ACM Transactions on the Web*, 3(1):3, 2009.

[5] K. Bharat and A. Broder. A technique for measuring the relative size and overlap of public web search engines. *Computer Networks and ISDN Systems*, 30(1):379–388, 1998.

[6] A. Broder. A taxonomy of web search. *ACM Sigir forum*, 36(2):3–10, 2002.

[7] C. D. Brown and H. T. Davis. Receiver operating characteristics curves and related decision measures: A tutorial. *Chemometrics and Intelligent Laboratory Systems*, 80(1):24–38, 2006.

[8] E. C. Chi and T. G. Kolda. On tensors, sparsity, and nonnegative factorizations. *SIAM Journal on Matrix Analysis and Applications*, 33(4):1272–1299, 2012.

[9] H. Chu and M. Rosenthal. Search engines for the world wide web: A comparative study and evaluation methodology. In *American Society for Information Science*, volume 33, pages 127–135, 1996.

[10] W. Ding and G. Marchionini. A comparative study of web search service performance. In *ASIS Annual Meeting*, volume 33, pages 136–42. ERIC, 1996.

[11] E. Enge, S. Spencer, J. Stricchiola, and R. Fishkin. *The art of SEO*. O'Reilly, 2012.

[12] Federal Communications Commission. Editorializing by broadcast licensees. Washington, DC: GPO, 1949.

[13] S. Gauch and G. Wang. Information fusion with profusion. In *1st World Conference of the Web Society*, 1996.

[14] Z. Guan and E. Cutrell. An eye tracking study of the effect of target rank on web search. In *SIGCHI conference on Human factors in computing systems*, pages 417–420. ACM, 2007.

[15] A. Gulli and A. Signorini. The indexable web is more than 11.5 billion pages. In *14th international conference on World Wide Web*, pages 902–903. ACM, 2005.

[16] A. Hannak, P. Sapiezynski, A. Molavi Kakhki, B. Krishnamurthy, D. Lazer, A. Mislove, and C. Wilson. Measuring personalization of web search. In *22nd international conference on World Wide Web*, pages 527–538. ACM, 2013.

[17] R. A. Harshman. Foundations of the parafac procedure: models and conditions for an" explanatory" multimodal factor analysis. Technical report, UCLA, 1970.

[18] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.

[19] F. W. Lancaster and E. G. Fayen. *Information Retrieval On-Line*. Melville Publishing Co., 1973.

[20] S. Lawrence and C. L. Giles. Searching the world wide web. *Science*, 280(5360):98–100, 1998.

[21] S. H. Lee, S. J. Kim, and S. H. Hong. On URL normalization. In *Computational Science and Its Applications–ICCSA 2005*, pages 1076–1085. Springer, 2005.

[22] T. Lei, R. Cai, J.-M. Yang, Y. Ke, X. Fan, and L. Zhang. A pattern tree-based approach to learning URL normalization rules. In *19th international conference on World Wide Web*, pages 611–620. ACM, 2010.

[23] D. Lewandowski. *Web search engine research*. Emerald Group Publishing, 2012.

[24] V. Maltese, F. Giunchiglia, K. Denecke, P. Lewis, C. Wallner, A. Baldry, and D. Madalli. *On the interdisciplinary foundations of diversity*. University of Trento, 2009.

[25] M.-C. Marcos and C. González-Caro. Comportamiento de los usuarios en la página de resultados de los buscadores. un estudio basado en eye tracking. *El profesional de la información*, 19(4):348–358, 2010.

[26] A. Pirkola. The effectiveness of web search engines to index new sites from different countries. *Information Research: An International Electronic Journal*, 14(2), 2009.

[27] K. Purcell, J. Brenner, and L. Rainie. *Search engine use 2012*. Pew Internet & American Life Project, 2012.

[28] E. Selberg and O. Etzioni. Multi-service search and comparison using the metacrawler. In *4th international conference on World Wide Web*, 1995.

[29] A. Spink, B. J. Jansen, C. Blakely, and S. Koshman. A study of results overlap and uniqueness among major web search engines. *Information Processing & Management*, 42(5):1379–1391, 2006.

[30] A. Spink, B. J. Jansen, and C. Wang. Comparison of major web search engine overlap: 2005 and 2007. In *14th Australasian World Wide Web Conference*, 2008.

[31] N. J. Stroud and A. Muddiman. Exposure to news and diverse views in the internet age. *ISJLP*, 8:605, 2012.

[32] D. Wilkinson and M. Thelwall. Search markets and search results: The case of Bing. *Library & Information Science Research*, 35(4):318–325, 2013.