

Smoothing Clickthrough Data for Web Search Ranking

Jianfeng Gao^{*}, Wei Yuan[#], Xiao Li^{*}, Kefeng Deng[§], Jian-Yun Nie[#]

^{*}Microsoft Research, Redmond, USA, Email: {jfgao; xiaol}@microsoft.com;

[§]Microsoft, Search Technology Center, Beijing, China, Email: kefengd@microsoft.com;

[#]University of Montreal, Canada, Email: {yuanwei; nie}@iro.umontreal.ca

ABSTRACT

Incorporating features extracted from clickthrough data (called *clickthrough features*) has been demonstrated to significantly improve the performance of ranking models for Web search applications. Such benefits, however, are severely limited by the data sparseness problem, i.e., many queries and documents have no or very few clicks. The ranker thus cannot rely strongly on clickthrough features for document ranking. This paper presents two smoothing methods to expand clickthrough data: query clustering via Random Walk on click graphs and a discounting method inspired by the Good-Turing estimator. Both methods are evaluated on real-world data in three Web search domains. Experimental results show that the ranking models trained on smoothed clickthrough features consistently outperform those trained on unsmoothed features. This study demonstrates both the importance and the benefits of dealing with the sparseness problem in clickthrough data.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval; I.2.6 [Artificial Intelligence]: Learning

General Terms

Algorithms, Experimentation

Keywords

Clickthrough Data, Smoothing, Random Walk, Discounting, Learning to Rank, Web Search

1. INTRODUCTION

We consider the task of ranking Web search results, i.e., a set of retrieved Web documents (URLs) are ordered by relevance to a query issued by a user. In this paper we assume that the task is performed using a ranking model (also called *ranker* for short) that is learned on labeled training data, i.e., human-judged query-document pairs. The ranking model is a function that maps the feature vector of a query-document pair to a real-valued relevance score. Such a learned ranking model is shown to be superior to classical retrieval models [6, 11] largely due to its ability to integrate both traditional criteria such as TF-IDF and BM25 values, and non-traditional features such as hyperlinks.

In general Web search, a document can be described by multiple text streams. Some of the most useful text streams for

Web search are (1) a *content stream* consisting of all the title and body texts in a page, (2) an *anchor stream* consisting of all the anchor texts of a page's incoming links, and (3) a *clickthrough stream* consisting of all the user queries that have click(s) on the document. Recent research shows that incorporating features extracted from the clickthrough stream (called *clickthrough features*) could significantly improve the performance of ranking models for Web search because the clickthrough stream can provide complementary information about a user's intention [1].

However, clickthrough data typically suffer from the sparseness problem. Two related aspects are involved. First, for a query, users only click on a very limited number of documents, thus the clicks are not complete. We refer to it as the *incomplete click* problem. Second, for many queries and documents, no click at all is made by users. We call this the *missing click* problem. As a consequence, the clickthrough streams for most of documents are either short or empty. Although one can use such raw text streams to extract some clickthrough features as in previous studies (e.g., [1, 6, 7]), their potential is severely limited because of the following reasons: First, with incomplete clicks, the click-related features that we can generate for a document-query pair are also incomplete and unreliable. Second, no clickthrough features can be generated for pairs without clicks. In the rankers used in most previous studies [1, 6, 7], this is equivalent to assigning zero values for clickthrough features. In ranker training, the zero-valued features make a categorical difference between the documents with and without clicks, and severely penalize the documents without clicks. However, in reality, the "true" difference between these documents may be much smaller because a document could be unclicked for a variety of reasons even if the document is relevant.

The missing click problem bears a strong resemblance to the problem of determining the frequency or probability of an unseen event, which has been well-studied in the context of estimating n -gram language models [8]. Various smoothing techniques have been proposed and successfully used to deal with this problem, including clustering (by grouping observations on similar n -grams) and discounting (by assigning some counts to unseen n -grams) [8, 14]. In the case of clickthrough data, we can consider a click for a document-query pair as an n -gram. Then clickthrough data can also be smoothed in two directions: by clustering similar queries or by assigning non-zero values to the clickthrough features of unclicked documents through discounting. In this paper, we propose to perform query clustering via Random Walk on click graphs, and a discounting method inspired by the Good-Turing estimator [13]. The Random Walk method is intended to address the incomplete click problem. In some particular settings, such as image retrieval [9] and query classification [21], it has been shown that expanding clicks to similar documents and queries via Random Walk can lead to significant improvements. However, to our knowledge, no

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

-
1. Message Web Design - Home
www.message.uk.com
 2. **MSN Web Messenger**
webmessenger.msn.com
2008-11-01 15:01:15
 3. High School Baseball Web
www.hsbaseballweb.com/message_boards.htm
 4. Send a Wireless Web Message
messaging.sprintpcs.com
 5. SprintPCS 2Way SMS
messaging.sprintpcs.com/sml/guestcompose.do
 6. Email on the Web
webmail.netzero.net
 7. USA MOBILITY
www.arch.com/message
 8. Message Boards - rootsweb.com
boards.rootsweb.com
 9. **Yahoo! Messenger - Chat, Instant message**
messenger.yahoo.com
2008-11-01 15:00:59
 10. Yahoo! Message Boards - Home
messages.yahoo.com
-

Figure 1. The query session for the query “web message”. Marked in bold are the links the user clicked on.

study has been carried out on general Web search applications showing a similar improvement. Our experiments show that the expanded clickthrough data is noisy, and it should be used with caution. Effective improvement is possible only when we extract those features that are robust to noise for ranking. Notice that documents and queries with no click cannot be enriched through Random Walk.

Thus, inspired by the Good-Turing method [13, 20], we present a discounting method to estimate the values of the clickthrough features for the documents without clicks.

Our experiments will show that both smoothing techniques can significantly improve the retrieval effectiveness compared to the utilization of raw clickthrough data. In particular, the simple discounting method will prove to be effective on all the three test datasets. This series of experiments strongly indicate that sparseness is a crucial problem in clickthrough data, and an appropriate solution to this problem allows us to better take advantage of clickthrough data.

In the rest of the paper, Section 2 describes background information on clickthrough data and rankers. Section 3 presents two smoothing techniques. Section 4 presents experiments. Related work and conclusions are presented in Sections 5 and 6.

2. BACKGROUND

In this section, we first describe the clickthrough data we use and the way a Web document is represented by a clickthrough stream. Then, we present the clickthrough features to be incorporated in ranking models. Finally, we review the ranking model used in our experiments. Notice that we focus on clickthrough features in this paper. The features extracted from other text streams will remain unchanged and be used in the same manner as before.

2.1 Clickthrough Streams for Documents

Clickthrough data used in this study consists of a set of query sessions that were extracted from one-year log files of a com-

msn web	0.6675749
webmessenger	0.6621253
msn online	0.6403270
windows web messenger	0.6321526
talking to friends on msn	0.6130790
school msn	0.5994550
msn anywhere	0.5667575
web message msn com	0.5476839
msn messenger	0.5313351
hotmail web chat	0.5231608
messenger web version	0.5013624
instant messenger msn	0.4550409
browser based messenger	0.3814714
im messenger sign in	0.2997275
msn web browser download	0.0926431
msn passport	0.0035466
download msn messenger 6	0.0027844
install msn toolbar	0.0027248
msn people	0.0025993
...	...

Figure 2. Fragments of the clickthrough stream for the link *http://webmessenger.msn.com*

mercial Web search engine. A query session contains a query issued by a user and a rank list of (top-10) links browsed by the same user (with or without click). Following the notations in [18], a query session is represented by a triplet (q, r, c) consisting of the query q , the ranking r presented to the user, and the set c of links (documents) the user clicked on. Figure 1 shows a query session for the query “web message”. The documents #2 and #9 are clicked by the user, and the dates and times of the two clicks are also recorded.

Previous work has utilized clickthrough data as implicit feedback for Web search ranking in two different ways. The first approach is to derive training data from clickthrough data directly [18, 19, 26]. In particular, [19] argued that relative preferences derived from clicks are reasonably accurate. For example, in Figure 1, the document #2 is assumed to be more relevant to the query “web message” than #1 because #2 is clicked, and #1, though ranked higher than #2, is not clicked. By doing so, one could derive a large amount of preference pairs. Then a ranking algorithm, such as LambdaRank [6], can be trained on such preference pairs.

The second category of work is to derive features from the clickthrough data and incorporate them into a ranking model [1, 28]. Our approach belongs to this category. The method is based on the assumption that all the queries that have clicks on a document form a description of the document from users’ perspective. One can see an example of such a clickthrough stream in Figures 2 for the document “webmessenger.msn.com”. It consists of all the queries that have one or more clicks on the document. In Figure 2, each line in a clickthrough stream consists of a query and a clickthrough score $Score(d, q)$, which can be considered as the importance of the query q in describing the document d , similarly to the TF-IDF scores. The score can be derived from raw click information recorded in log files heuristically. In our experiments, one of the simplest functions that work well across all data sets is:

