

Predicting Web Actions from HTML Content

Brian D. Davison*
Computer Science & Engineering Dept.
Lehigh University, Packard Lab
19 Memorial Drive West
Bethlehem, PA 18015 USA
davison@cse.lehigh.edu

ABSTRACT

Most proposed Web prefetching techniques make predictions based on the historical references to requested objects. In contrast, this paper examines the accuracy of predicting a user's next action based on analysis of the content of the pages requested recently by the user. Predictions are made using the similarity of a model of the user's interest to the text in and around the hypertext anchors of recently requested Web pages. This approach can make predictions of actions that have never been taken by the user and potentially make predictions that reflect current user interests. We evaluate this technique using data from a full-content log of Web activity and find that textual similarity-based predictions outperform simpler approaches.

Categories and Subject Descriptors

H.4.3 [Information Systems Applications]: Communications Applications—*Information browsers*; H.3.5 [Information Storage and Retrieval]: Online Information Services—*Web-based services*; H.3.1 [Information Storage and Retrieval]: Content Analysis and Indexing; H.5.2 [Information Interfaces and Presentation]: User Interfaces—*Evaluation/methodology, User-centered design*

General Terms

Algorithms, Measurement, Performance, Design, Experimentation, Human Factors

Keywords

WWW, prefetching, prediction, user modeling, textual similarity, information retrieval

1. INTRODUCTION

The World Wide Web has become the ultimate hypertext interface to billions of individually accessible documents.

*Work was performed while a graduate student in the Department of Computer Science at Rutgers University.

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, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

HT'02, June 11-15, 2002, College Park, Maryland, USA.
Copyright 2002 ACM 1-58113-477-0/02/0006 ...\$5.00.

Unfortunately, for most people, retrieval times for many of those documents are well above the threshold for perception, leading to impressions of the “World Wide Wait.” Web caching [22] is one approach that can effectively eliminate the retrieval time for many recently accessed documents, but it cannot help with documents that have never been visited in the past.

If future requests can be correctly predicted, those documents could potentially be pre-loaded into a local cache to provide fast access when requested by a user, even when they have not been accessed recently. This paper proposes and evaluates such a predictive approach. While conventional predictive techniques look at past actions as a basis to predict the future, we consider an approach based on the hypertext characteristics of the Web.

Recently we showed that Web pages are typically linked to pages with related textual content, and more specifically, that the anchor text was a reasonable descriptor for the page to which it pointed [20]. Many applications already take advantage of this Web characteristic (e.g., for indexing terms not on a target page [33, 8] or to extract high-quality descriptions of a page [2, 3]). In this work we will examine in detail another application of this result.

Our motivating goal is to accurately predict the next request that an individual user is likely to make on the WWW. Therefore, this paper examines the value of using Web page content to make predictions for what will be requested next. Many researchers have considered complex models for history-based predictions for pre-loading (e.g., [47, 55]), but relatively few have considered using anything other than simplistic approaches to the use of Web page content. Unlike many techniques that do examine content, our approach does not noticeably interfere with the user experience at all — it does not ask for a statement of interest, nor does it modify the pages presented to the user. Instead, it can be used to invisibly improve user-perceived performance.

One naive content-based approach is to pre-load all links on a page. A slightly more intelligent approach is to pre-load as time permits the links of a page, in HTML source order from first to last (corresponding generally to links visible from top to bottom). In this work we compare those approaches with an information retrieval-based one that ranks the list of links using a measure of textual similarity to the set of pages recently accessed by the user. In summary, we find that textual similarity-based predictions outperform the simpler approaches.




Figure 1: The distribution of the number of hyper-text (non-embedded resource) links per retrieved HTML page.

2. BACKGROUND

By transparently reducing network latencies and bandwidth demands, Web resource caching [22] has become a significant part of the infrastructure of the Web. Unfortunately, caching can only help when the objects requested are already present in the cache. Typically, caches contain objects that have been accessed in the past. Prefetching [47, 44], however, can be used to speculatively put content into the cache in advance of an actual request. One difficulty, however, is in knowing what to prefetch. Typical approaches have used Markovian techniques (e.g., [28, 55]) on the history of Web page references to recognize patterns of activity. Others prefetch bookmarked pages and often-requested objects (e.g., [50]). Still others prefetch links from the currently requested page [46, 13, 10, 57, 38, 42, 37, 51, 26, 34].

However, prefetching all of the links of the current page is not a viable option, given that the number of links per page can be quite large, and that a prefetching system typically has only a limited amount of time to prefetch before the user makes a new selection. While likely heavy-tailed [11, 16], this “thinking-time” is typically less than one minute [43]). Likewise, prefetching content that is never used exacts other costs as additional resources such as bandwidth are consumed. We do not consider the thinking time or prefetching costs further here (see [6, 39, 1, 60] for various utility theoretic approaches to balancing cost and benefits).

The difficulty of selecting a choice in content-based analysis is illustrated in Figure 1. It shows the distribution of the number of unique hypertext (non-embedded resource) links per retrieved HTML page (with a mean of 26 links and a median of 11) for the dataset used in this paper (described further in Section 5.1). Note that this histogram reflects the distribution of pages requested by users (which includes repeated retrievals), versus the more or less static distribution of pages on the Web (such as described by Bray [7]).

The prefetching of content may cause problems, as not all content is cacheable (so prefetching it only wastes resources), and prefetching even cacheable content can abuse server and network resources, worsening the situation [21, 15]. An alternative is to do everything but prefetch [14] — that is, to resolve the DNS in advance, connect in advance to the Web server, and even warm up the Web server with a dummy request.




Figure 2: The distribution of the minimum distance, in terms of the number of requests back that the current page could have been prefetched.

Prefetching systems are difficult to evaluate, especially content-based prefetching ones (although see [18] for one approach) because even state-of-the-art proxy evaluation techniques (e.g., Polygraph [53] as used in caching competitions [54]) use artificial data sets without actual content. Even when using captured logs as the workload, real content will need to be retrieved, and will no longer be the same as the content seen when the logs were generated.

A more complex analysis is shown in Figure 2. It shows (by summing the fractions for columns 1 and 2-3) that approximately 46% of prefetchable pages (those that are considered likely to be cacheable) can be reached by examining the links of the current page and its two immediate predecessors in the current user’s request stream.

For comparison, the per-client average recurrence rate (that is, the rate at which requests are repeated [35]) for clients making more than 500 prefetchable requests was 49.2%, and for all clients was 22.5% (since some clients made very few requests). This means that if each client utilized an infinite cache that that passively recorded cacheable responses, the per client average hit rate on prefetchable content would be less than 25% (corresponding to the average recurrence rate). However, if the clients additionally prefetch all perceived prefetchable links, the hit rates on prefetchable content would increase to almost 75%, because the prefetchable pages that would not be in the cache (i.e. the misses) correspond at most to the pages represented by a distance of unknown in Figure 2. Such pages were never linked from any pages seen during our data collection. Note that these are the results of per-user analysis on `text/html` responses not satisfied by the browser cache. Additionally, a single shared cache (that stores all user requests) would have a recurrence rate of 50.9% for cacheable resources (47.3% for all resources). From this analysis we conclude that there is significant potential for content-based prediction of future Web page requests when caching is considered.

3. RELATED WORK

Chan [12] outlines a non-invasive learning approach to construct Web user profiles, that incorporates content, linkage, as well as other factors. It uses frequency of visitation, whether bookmarked, time spent on page, and the percentage of child links have been visited to estimate user interest in a page. Predictions of user interest in a page are made

by building a classifier with word phrases as features and labeled with the estimation of page interest. Thus, new pages can be classified as potentially being of interest from the examination of their contents. Chan’s goal is similar but not the same as ours. He is interested in modeling what the user might like to see, while we are concerned with what the user will do. (See Haffner *et al.* [36] for a theoretical discussion of the similarities of the two goals.) Unlike the approaches considered here, his method requires that anything that is to be considered for recommendation must already be prefetched (for analysis). Additionally, the methods he proposed were not incremental, nor did they change over time.

Chinen and Yamaguchi [13] describe and evaluate the performance of their Wcol prefetching proxy cache. Their system simply prefetches up to the first n embedded images and m linked documents. In a trace with a large number of clients, they found that if all targets were prefetched, a hit rate of 69% was achievable. Prefetching ten targets resulted in approximately 45%, and five targets corresponded to about 20% hit rate. In this paper we use a variant of their approach (termed original order) as one of the prediction methods that we will use for evaluation.

More sophisticated link ranking approaches are possible. Klemm’s WebCompanion [42] is a prefetching agent that parses pages retrieved, tracks costs to retrieve from each host, and prefetches the pages with the highest estimated round trip time. Average access speedups of over 50% with a network byte overhead (i.e., wasted bandwidth) of 150% are claimed from tests using an automated browser.

Ibrahim and Xu [37] describe a neural net approach to prefetching, using clicked links to update weights on anchor text keywords for future predictions. Thus they rank the links of a page by a score computed by an artificial neuron. System performance of approximately 60% hit ratio is claimed based on an artificial news-reading workload.

A number of others have considered the similar problem of Web page recommendation and Web page similarity. However, they (e.g., [25]) tend to use information that can only be acquired by the retrieval of additional pages, which would likely overwhelm the link bandwidth and negate the purpose of intelligent prefetching. Our approach, in contrast, uses only the content that the client has previously requested to base decisions on what to prefetch next.

Recently Yang *et al.* [59] considered various approaches to hypertext classification. Their results are mixed, finding that identification of hypertext regularities and appropriate representations are crucial to categorization performance. They note, however, that “algorithms focusing on automated discovery of the relevant parts of the hypertext neighborhood should have an edge over more naive approaches.” While their study examines a related question, unlike our work it treats all links (both incoming and outgoing) equally, and does not consider the use of anchor or surrounding text.

One of the difficulties in comparing work in the area of content-based prediction and prefetching is the inability to share data. Some, like Chinen and Yamaguchi, evaluate their work by using their system on live data. Even if the trace of resources requested were available, such traces deteriorate in value quickly because the references contained within them change or disappear. Others, such as Klemm as well as Ibrahim and Xu test their systems with relatively contrived methods, which make it difficult to ascertain general applicability. Our initial approach here has been to

capture a full-content Web trace so that analysis can be performed offline. However, because of privacy concerns, we (like other researchers) are unable to provide such data to others. Elsewhere [18] we have proposed a mechanism capable of comparing multiple prefetching proxy caches simultaneously. Such a system is under development [24], but is only applicable to fully realized prefetching systems.

4. CONTENT-BASED PREDICTION

In many user interfaces, multiple user actions may appear atomic and thus may be inaccessible for analysis. For those environments, prediction methods are often limited to those that can recognize and extrapolate from patterns found in a history log (e.g., UNIX shell commands [23] or telecommunication switch alarms [58]). In contrast, the Web provides easy access to the content being provided to the user. Moreover, this content prescribes boundaries within which the user is likely to act. By modifying the browser, Tauscher and Greenberg [56] collected data showing that users predominantly (approximately 80% of the time) choose a link from the page being shown to them (as opposed to the using a bookmark or typing a URL by hand). We have seen similar results from analysis of proxy and server logs (measuring the presence of referrers for HTML pages [19]). Therefore, the set of links shown to the user is a significant guide to what the user is likely to request next.

In addition, the content provides possible next steps that historically-based prediction approaches cannot. A user model based on past experience cannot predict an action previously never taken. Thus, a content-based approach can be a source of predictions for new situations with which history cannot help.

While a user does provide some thinking time between requests in which prefetching may occur, it is limited, and so there may be insufficient time to prefetch all links. Moreover, even if there were sufficient time, the user and network provider may find the extra bandwidth usage to be undesirable. Thus, the approach is typically not to predict all links, but instead to assign some weight to each link and use only the highest scoring links for prediction (and subsequent prefetching).

A prefetching system, whether integrated into a browser, or operating at a proxy, has access to the content of the pages served to the user. Given a particular page, the external links that it contains can be extracted. The anchor and surrounding text can be used as a description of the page (as shown in [20] and used in [2, 3]). Given a set of links and their descriptions, the problem is then how to rank them.

Ideally, we would have a model of a user’s interest and the current context in which the user is found (e.g., as in [9]). One approach is to explicitly ask the user about their current interest, as done in WebWatcher [4, 40] and in AntWorld [41] or to rate pages (as in Syskill and Webert [48, 49] and AntWorld). We could also consider modifying the content seen by the user to give hints or suggestions of recommended or preloaded content (as in Letizia [45, 46], WebWatcher, QuIC [29], and common uses of WBI [5]).

However, our preference is to build a user model as unobtrusively as possible (as in [12]) to perform prefetching behind the scenes. For the purposes of this study, this means we cannot ask the user explicitly of his or her interest, nor change the content, but it does enable us to be able to work

offline with logs of standard usage. Therefore, we will not consider obtrusive approaches further here.

Instead we will use the textual contents of recently requested pages as a guide to the current interests of the user. This method allows us to model a user’s changing interests without mind-reading or explicit questions, and is intended to combine with other sources of predictions to generate an adaptive Web prefetching approach [17]. How well it performs is the subject of the next section.

5. EXPERIMENTS

As mentioned earlier in Section 2, experimental evaluation of a content-based prefetching system is difficult. Rather than deploying and testing a complete prefetching system, we have elected to test the content-based prediction methods described above in an offline manner. For comparison, we will use a baseline of the accuracy of a random ranking.

5.1 Workload

We implemented a custom Java-based non-caching HTTP/1.0-compliant proxy to monitor and record the full HTML content and all headers of user requests and responses. We captured approximately 135,000 requests, of which just over a third generated HTML responses. (The rest correspond to embedded resources such as images and non-data responses such as errors or content moved.) Users were predominantly computer science students and staff from Rutgers University. More than fifty distinct clients used the service from August 1998 to March 1999.

While covering close to seven months, this trace is rather small, and predominantly reflects the traffic of university users. However, with concerns over privacy, the choice was made to recruit explicit volunteers instead of surreptitiously recording the activity of all users of an existing cache or packet snooping on a network link (e.g., as in [31, 30]).

This log does not distinguish users — it only distinguishes clients based on IP address, and so if multiple users of this proxy operated browsers on the same machine (a possibility under UNIX) or behind the same proxy, those users’ requests could be interleaved. Likewise, users with dynamic IP addresses are seen as distinct users on each session and multiple users (if assigned the same IP address at different times) may seem like the same user returning over multiple sessions. We believe the likelihood of multiple users per IP (whether by proxy, or the re-use of dynamic addresses) happening is small for this data. More importantly, though, the log does not reflect the requests made by users that are satisfied by browser caches. Users of this proxy were directed to use the proxy, and not told to disable browser caching.

5.2 Methodology

The data was recorded in several segments, generally corresponding to separate runs of the proxy. The first segment, representing about 5% of the log, was used for exploratory analysis and was manually examined. As we describe specific adaptations made to the data, it is from the measurement and analysis of this initial segment. Overall results shown in graphs will reflect measurements over the entire dataset.

5.2.1 Data preparation

Our custom proxy attempted to record all HTTP [32] headers of requests and responses, and the contents of all




Figure 3: Sequence of HTML documents requested. The content of recently requested pages are used as a model of the user’s interest to rank the links of the current page to determine what to prefetch.

HTML responses. These were written in chronological order, thus interleaving requests when multiple users were active. Since we are interested in predictions from sequences of requests for individual users, we extracted the HTML records and sorted by client and timestamp. From these sorted access requests, we built sets of up to five HTML responses. (While a larger number of responses might have provided more data, we wanted to model the current user interest which might be diluted with data from further in the past.) Thus, for example, responses d_1, d_2, d_3, d_4, d_5 would be the first set, from which we would attempt to predict d_6 . Likewise, d_2, d_3, d_4, d_5, d_6 corresponded to the second set, from which we would attempt to predict d_7 . However, occasionally error messages (e.g., HTTP response codes 403, 404), redirections (e.g., 301, 302), and content unchanged (code 304) ended up being included, as Web servers sometimes described these as of type `text/html` as well.

Since we had stored the content, whenever we encountered a 304 content unchanged response for content that we had previously recorded, we used our copy of the earlier content. Requests for URLs generating the remaining 304 responses could be predicted, but not used as content (since that content must have been sent before starting our measurements we couldn’t have it) for use in future predictions.

We discarded those responses with HTTP error codes (i.e., response codes ≥ 400), but retained redirections (e.g., response codes 301, 302) since links were made to one URL, which would then redirect elsewhere. In this case, the original log might have included retrieval events $d_1, d_2, d_3, d_4, d_5, r_6, d_6, d_7$ (where r_6 corresponds to a redirection to d_6) but we would instead generate the two cases d_1, d_2, d_3, d_4, d_5 predicting r_6 and d_2, d_3, d_4, d_5, d_6 predicting d_7 . We use r_6 as a target of prediction since any link followed (if present at all) would have been made to r_6 , but there is no usable content in r_6 so we don’t use it for text analysis. Unfortunately, as mentioned earlier, we only wanted HTML responses, but occasionally we would see redirections for non-HTML content, such as banner advertising that would use one or more redirections. We manually filtered URLs for some of the more popular advertising hosting services to minimize this effect.

The pages were parsed with the Perl `HTML::TreeBuilder` library. Text extraction from the HTML pages was performed using custom code that down-cased all terms and replaced all punctuation with whitespace so that all terms are made strictly of alphanumerics. Content text of the page includes title but not META tags nor alt text for images.

URLs were parsed and extracted using the Perl URI:URL library plus custom code to standardize the URL format (down-casing host, dropping #, etc.) to maximize matching of equivalent URLs. The basic representation of each textual item was bag-of-words with term frequency.

5.2.2 Prediction methods

We present a total of four methods to rank the URLs of the current page for prediction. Two are simple methods: the baseline *random* ordering, and *original rank ordering* (i.e., first-occurrence order in page). The second two rank each link based on the similarity of the link text with the combined non-HTML text of the preceding pages (as illustrated in Figure 3. To measure the similarity between two text documents (D_1, D_2), we use a very simple metric¹:

$$\text{TF}(w, D_i) = \text{number of times term } w \text{ appears in } D_i$$

$$\text{Text-Sim}(D_1, D_2) = \sum_{\text{all } w} \text{TF}(w, D_1) * \text{TF}(w, D_2)$$

The first method (termed *similarity* in our tests) uses the highest similarity score (in the case of multiple instances of the same link to a URL) of the anchor and surrounding text to the preceding pages. The second (termed *cumulative*) is identical, except that it sums the similarity scores when there are multiple identical links (giving extra weight to such a URL).

Links recognized include the typical `a href`, but also `area href`, and `frame src`. We did not parse embedded JavaScript to find additional links, although in reviewing the initial data segment we did find JavaScript with such links.

Since links are typically presented within some context, in addition to anchor text we tested the addition of terms before and after the anchor. For example, if the HTML content of a page were:

```
The <a href="http://www.acm.org/acm1/">ACM</a>
Conference will take you beyond cyberspace.
```

then 0 additional terms would provide just {ACM}. If we permit 5 additional terms on each side, then we would get the set {The, ACM, Conference, will, take, you, beyond} to be used for comparison against the text of the preceding pages. Note that text around an anchor may be used by more than one link, in the case that the window around an anchor overlapped that of another anchor. In experiments on the first data segment, we tested the use of up to 20, 10, 5, and 0 additional terms and found that in most cases, when more terms were used, performance was slightly better. Thus we use as many as twenty additional terms before and after the anchor text. However, the text around a URL was never permitted to extend into another URL's anchor text.

Furthermore, we noticed that when exploring pages within a site, certain elements of that site were often repeated on every page. Examples would include a copyright notice with a link to a use license page. In this case, such repeated elements would get undue weight over non-repeated text.

¹Other variations of this similarity metric (e.g., to include factors for document length) were tested on the first data segment but they performed similarly or worse. We make no claims as to the optimality of this similarity metric for this task.




Figure 4: Fraction of pages that are uncacheable, found as a link from previously requested page (potentially predictable), or not found as a link from previous page.


As a result, a link containing the repeated text (the link to a license, in the case above) would be ranked highly, even if irrelevant to the true content of the pages. To combat this effect, we tested a variation that no longer used all text. Instead we used the first page, plus the difference to the second page, plus the differences between pages 2 and 3, etc. In this way we hoped to eliminate significant emphasis on repeated structure in the pages. This was performed by a Perl version of `htmldiff` [27].

We used the well-known Porter [52] stemming (since it helped slightly), but did not use stop word elimination (either slightly worse or no change). The non-HTML text (including title, keywords, and meta description) of the preceding pages was combined to serve as a “document” representing the current user interest against which we would measure similarity. However, since parts of the current page are also represented in the anchor texts pointing to the target documents, we tested the use of the current page in the model for user interest, and found that it also decreased performance on the initial segment slightly, and so we do not include it. Thus in summary, we use the non-repeating text of the preceding four pages as our model of current user interest against which we measure similarity, as illustrated in Figure 3.


5.3 Evaluation

Given the ranking methods described above, we will evaluate their predictive accuracy (the fraction of times that they include the correct next request in their prediction) over the entire dataset. In most cases, however, we will scale the results to show the fraction predicted correctly out of those possible to get correct (defined below). We will consider variations in the number of predictions by evaluating algorithms that use their best 1, 3, or 5 predictions. Additionally, we will consider the case in which the clients have an infinite cache into which predictions are prefetched. With such a cache, success is measured not by whether the immediate prediction was correct, but whether the requested object is present in the cache (a test arguably closer to real world considerations).


The ultimate goal of this effort is to use the best prediction method to provide suggestions for prefetching. It is generally not possible to prefetch the results of form entry, and is potentially undesirable to prefetch dynamically gen-



(a) Top one prediction



(b) Top three predictions



(c) Top five predictions

Figure 5: Overall predictive accuracy (within potentially predictable set) of content prediction methods when considering links from the most recently requested page.

erated pages, so we would like to not allow predictions for such pages (i.e., all responses generated by POSTs or with URLs containing “cgi” or “?”). Those responses represent 28% of all HTML collected. Additionally, we can recognize those cases where the next cacheable URL requested was not present in the preceding page, making a correct prediction an impossibility in another 47%. The remaining 25% have the potential to be predicted correctly by choosing from the links of the current page (see Figure 4). Approximately 2.8%

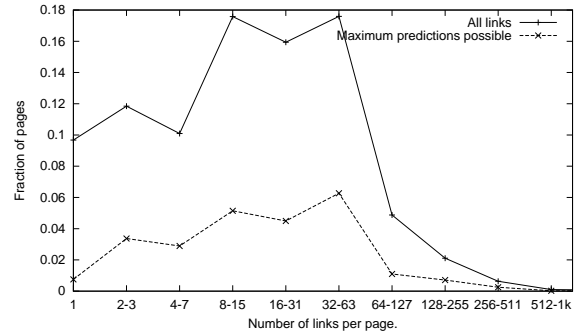


Figure 6: Potential for prediction when considering links from the most recently requested page.

of those pages (.7% of the total) had only one URL, and so will be predicted correctly under all prediction mechanisms.

Typically a prefetching system has time to prefetch more than one page, and so we consider the case in which we can select the top three and five most-likely URLs and mark the set as “correct” if one of them is the requested page. Under this arrangement, the set of potentially predictable pages stays the same (25%), but the number of default or “free” wins grows since there will be more cases in which the number of distinct links available is within the size of the predicted set.

6. EXPERIMENTAL RESULTS

6.1 Overall results


Figure 5 shows the overall accuracy for each of the content prediction algorithms, as a fraction of predictions that are possible to get correct from the last page of links. It depicts performance for the cases in which only the highest-scoring prediction is evaluated as well as when the top-3- and top-5-highest-scoring predictions are used. The “free” cases are also plotted to show a lower bound (recall that the free cases are pages in which the number of links in the page is no larger than the number of predictions permitted, thus automatically correct).

It can be seen that the similarity and cumulative ranking methods outperform the original rank ordering, and all three outperform the baseline random order. The cumulative approach edges out the similarity approach for top-one, but similarity is best for top-three and top-five.


6.2 Distributions of performance

This paper examines the predictability of web page requests based on the links of the current page. In this section we consider the distribution of successful predictions with respect to the number of links on the current page. Figure 6 compares the overall distribution of the number of links per page with the distribution of cases in which the predicted page is found in the links of the last page.


Figure 7 shows the distribution of the number of correct predictions per page size for each method plus the distribution of cases that can possibly be predicted correctly. Note that for clarity we have omitted the cumulative approach from these and further graphs, as its performance is closely tied with that of the similarity method. For most points of



(a) Top one prediction



(b) Top three predictions



(c) Top five predictions

Figure 7: Performance distribution for content prediction algorithms when allowing guesses from the most recently requested page.

Figure 7(a), in which a single prediction is permitted, the performance of the various approaches are almost indistinguishable, but with more allowed predictions, more variation becomes visible.

6.3 Results with infinite cache

Since the system prefetching Web pages would be putting them into a cache, we consider briefly here the performance of prediction approaches in such a context. Figure 8 shows a




Figure 8: Predictive accuracy (as a fraction of all pages) of content prediction methods with an infinite cache compared to no prediction.




Figure 9: Performance distribution for original order with and without an infinite cache.

summary of the predictive accuracy of two ranking methods assuming the use of an infinite cache, and compares them to an infinite cache without prefetching. Figures 9 and 10 provide performance distributions of the original rank orderings and similarity orderings respectively. In these figures, note that the top curve represents the total page distribution, since we are no longer limited to the pages potentially reached from the last ones retrieved. The original (non-cached) performance is shown (at the bottom), while the performance using an infinite cache is much higher. However, the cache does not do all the work, as the performance of an infinite cache alone is visibly below the cache+prefetching lines.

Figure 11 compares the cached performance of both the original rank ordering and similarity based ordering (we omit the top-three curves for clarity). The performance of cached original ranking and similarity ranking are nearly identical in some points, but similarity eventually rises above the original ordering (when five predictions are allowed).

7. DISCUSSION

Recall that the non-cached results presented are in terms of potentially predictable pages. This is, in fact, a relatively small fraction of all requests generated. For example, the best predictive method was shown to be the similarity ranking when predictions of the top 5 links are permitted, achieving an accuracy of approximately 55% (a relative im-




Figure 10: Performance distribution for similarity ranking with and without an infinite cache.




Figure 11: Comparison of performance distributions for original order and similarity with an infinite cache.

provement of close to 30% over random link selection). However, this represents success in only 14% of all cases, since only a quarter of all pages are possible to predict in this manner. In the case of an infinite cache we are able to provide hits for 64% of all requests (a 40% improvement over a system without prefetching). Actual performance is likely to lie in between, as caches are finite and prefetched content can expire before it is needed.

Perhaps more importantly, this analysis is of a relatively small trace, and may not be representative of users in non-academic environments. However, more extensive analysis would require capturing content from a larger, more general population (such as that of an ISP), and is likely to raise significant privacy concerns.

The prefetching examined here is only for HTML resources. However, HTML resources represent only a fraction of all requests made by users. Most are embedded resources (such as images, sounds, or Java applets or the result of JavaScript that are automatically retrieved by the browser). We made the choice to ignore such resources as most of them are easy to predict by examining the HTML of the page in which they are embedded. This choice has a drawback, though, as users do indeed make requests for non-HTML resources, such as PDF and PostScript files, plain text resources, and the downloading of programs, albeit much less frequently.

In practice, a prefetching system will have time to fetch a variable number of resources. We have examined only three points on the range — allowing one, three, and five predictions. We believe that the typical number will fall within this range, however, as it will likely be useful to prefetch the embedded resources of prefetched pages, and some resources and pages will be cached from previous retrievals.

While extensive, these tests are not comprehensive — we have not attempted to disprove the possibility of other algorithms (text-based or otherwise) outperforming those presented here. In particular, we believe that systems with a stronger model of the user’s interest (e.g., AntWorld or WebWatcher) could provide for better prediction, but alternately may lose when the user’s interest shifts (as is often the case when surfing the Web). For comparison, recall that WebWatcher [4, 40] compares the user’s stated goal, the given page, and each link within it to other pages requested by previous users, their stated goals, and the links they selected. Using assessed similarity on this information in a restricted domain, it picked three most likely links, from which the selected link was present 48.9% of the time. In a second, smaller trial, WebWatcher scored just 42.9% (as compared to humans at 47.5% on the same task). Our most comparable experiments have shown a lower accuracy (approximately 40%), but additionally they have been used on a general task in an unlimited domain, with no user goals nor past users for guidance. Thus, these experiments provide evidence of the applicability of content-based methods for predicting future Web pages access.

8. CONCLUSION

This paper has examined the performance of Web request prediction on the basis of the content of the recently requested Web pages. We have compared text-similarity-based ranking methods to simple original link ordering and a baseline random ordering and found that similarity-based rankings performed 29% better than random link selection for prediction, and 40% better than no prefetching in a system with an infinite cache. In general, textual similarity-based rankings outperformed the simpler methods examined in terms of accuracy, but in the context of an infinite cache, the predictive performance of similarity and original rankings are fairly similar.

Most proposed Web prefetching techniques make predictions based on historical data. History-based prefetching systems can do well, when they have a model for the pages that a user is visiting. However, since users often request unseen (e.g., new) resources (perhaps 40% of the time [56]), history may not always be able to provide suggestions, or may have insufficient evidence for a strong prediction. In such cases, content-based approaches, such as those presented here, are ideal alternatives, since they can make predictions of actions that have never been taken by the user and potentially make predictions that are more likely to reflect current user interests.

Acknowledgments

Thanks are due to Haym Hirsh and William Pottenger for comments on an earlier version of this paper. This work has been supported in part by NSF grant ANI 9903052, and by DARPA under Order Number F425 (via ONR contract N6600197C8534).

9. REFERENCES

- [1] D. W. Albrecht, I. Zukerman, and A. E. Nicholson. Pre-sending documents on the WWW: a comparative study. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI-99)*, volume 2, pages 1274–1279, Stockholm, Sweden, 1999. Morgan Kaufmann.
- [2] E. Amitay. Using common hypertext links to identify the best phrasal description of target web documents. In *Proceedings of the SIGIR'98 Post-Conference Workshop on Hypertext Information Retrieval for the Web*, Melbourne, Australia, 1998.
- [3] E. Amitay and C. Paris. Automatically Summarising Web Sites - Is There A Way Around It?. In *Proceedings of the Ninth ACM International Conference on Information and Knowledge Management (CIKM 2000)*, Washington, DC, Nov. 2000.
- [4] R. Armstrong, D. Freitag, T. Joachims, and T. Mitchell. WebWatcher: A Learning Apprentice for the World Wide Web. In *Proceedings of the AAAI Spring Symposium on Information Gathering from Distributed, Heterogeneous Environments*, Stanford University, Mar. 1995. AAAI Press.
- [5] R. Barrett, P. P. Maglio, and D. C. Kellem. How to personalize the Web. In *Proceedings of the ACM SIGCHI'97 Conference on Human Factors in Computing Systems*, pages 75–82, Atlanta, GA, Mar. 1997. ACM Press.
- [6] A. Bestavros. Speculative Data Dissemination and Service to Reduce Server Load, Network Traffic and Service Time for Distributed Information Systems. In *Proceedings of the International Conference on Data Engineering (ICDE'96)*, New Orleans, LA, Mar. 1996.
- [7] T. Bray. Measuring the Web. In *Proceedings of the Fifth International World Wide Web Conference*, Paris, France, May 1996.
- [8] S. Brin and L. Page. The anatomy of a large-scale hypertextual Web search engine. In *Proceedings of the Seventh International World Wide Web Conference*, Brisbane, Australia, Apr. 1998.
- [9] P. J. Brown and G. J. F. Jones. Context-aware retrieval: exploring a new environment for information retrieval and information filtering. *Personal and Ubiquitous Computing*, 5(4):253–263, 2001.
- [10] CacheFlow Inc. Active caching technology. <http://www.cacheflow.com/technology/whitepapers/active.cfm>, 2002.
- [11] L. D. Catledge and J. E. Pitkow. Characterizing Browsing Strategies in the World Wide Web. *Computer Networks and ISDN Systems*, 26(6):1065–1073, 1995.
- [12] P. K. Chan. A non-invasive learning approach to building Web user profiles. In *Proceedings of WebKDD'99*, pages 7–12, San Diego, Aug. 1999. Revised version in Springer LNCS Vol. 1836.
- [13] K.-i. Chinen and S. Yamaguchi. An interactive prefetching proxy server for improvement of WWW latency. In *Proceedings of the Seventh Annual Conference of the Internet Society (INET'97)*, Kuala Lumpur, June 1997.
- [14] E. Cohen and H. Kaplan. Prefetching the means for document transfer: A new approach for reducing Web latency. In *Proceedings of IEEE INFOCOM*, Tel Aviv, Israel, Mar. 2000.
- [15] M. Crovella and P. Barford. The network effects of prefetching. In *Proceedings of IEEE INFOCOM*, 1998.
- [16] C. R. Cunha, A. Bestavros, and M. E. Crovella. Characteristics of WWW Client-based Traces. Technical Report TR-95-010, Computer Science Department, Boston University, July 1995.
- [17] B. D. Davison. Adaptive Web Prefetching. In *Proceedings of the 2nd Workshop on Adaptive Systems and User Modeling on the WWW*, pages 105–106, Toronto, May 1999. Position paper. Proceedings published as Computing Science Report 99-07, Dept. of Mathematics and Computing Science, Eindhoven University of Technology.
- [18] B. D. Davison. Simultaneous Proxy Evaluation. In *Proceedings of the Fourth International Web Caching Workshop (WCW99)*, pages 170–178, San Diego, CA, Mar. 1999.
- [19] B. D. Davison. Web traffic logs: An imperfect resource for evaluation. In *Proceedings of the Ninth Annual Conference of the Internet Society (INET'99)*, June 1999.
- [20] B. D. Davison. Topical Locality in the Web. In *Proceedings of the 23rd Annual International Conference on Research and Development in Information Retrieval (SIGIR 2000)*, Athens, Greece, July 2000.
- [21] B. D. Davison. Assertion: Prefetching with GET is not good. In *Proceedings of the Sixth International Workshop on Web Caching and Content Distribution (WCW'01)*, Boston, MA, June 2001.
- [22] B. D. Davison. A Web caching primer. *IEEE Internet Computing*, 5(4):38–45, July/August 2001.
- [23] B. D. Davison and H. Hirsh. Predicting sequences of user actions. In *Predicting the Future: AI Approaches to Time-Series Problems*, pages 5–12, Madison, WI, July 1998. AAAI Press. Proceedings of AAAI-98/ICML-98 Workshop, published as Technical Report WS-98-07.
- [24] B. D. Davison and C. Krishnan. ROPE: The Rutgers Online Proxy Evaluator. Technical Report DCS-TR-445, Department of Computer Science, Rutgers University, 2001.
- [25] J. Dean and M. R. Henzinger. Finding related pages in the World Wide Web. In *Proceedings of the Eighth International World Wide Web Conference*, pages 389–401, Toronto, Canada, May 1999.
- [26] Deerfield Corporation. InterQuick for WinGate. <http://interquick.deerfield.com/>, 2002.
- [27] F. Douglis and T. Ball. Tracking and viewing changes on the Web. In *Proceedings of the USENIX Technical Conference*, pages 165–176, San Diego, Jan. 1996.
- [28] D. Duchamp. Prefetching hyperlinks. In *Proceedings of the Second USENIX Symposium on Internet Technologies and Systems (USITS '99)*, Boulder, CO, Oct. 1999.
- [29] S. R. El-Beltagy, W. Hall, D. De Roure, and L. Carr. Linking in context. In *Proceedings of the Twelfth ACM Conference on Hypertext and Hypermedia (HT'01)*, Aarhus, Denmark, Aug. 2001.

- [30] A. Feldmann. Continuous online extraction of HTTP traces from packet traces. In *World Wide Web Consortium Workshop on Web Characterization*, Cambridge, MA, Nov. 1998. Position paper.
- [31] A. Feldmann, R. Cáceres, F. Douglis, G. Glass, and M. Rabinovich. Performance of Web Proxy Caching in Heterogeneous Bandwidth Environments. In *Proceedings of IEEE INFOCOM*, pages 106–116, New York, Mar. 1999.
- [32] R. T. Fielding, J. Gettys, J. C. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee. Hypertext Transfer Protocol — HTTP/1.1. RFC 2616, <http://ftp.isi.edu/in-notes/rfc2616.txt>, June 1999.
- [33] Google Inc. Google home page. <http://www.google.com/>, 2002.
- [34] Goto Software Corporation. WebEarly3 product page. <http://www.goto-software.com/WE3/WEstart.htm>, 2002.
- [35] S. Greenberg. *The Computer User as Toolsmith: The Use, Reuse, and Organization of Computer-based Tools*. Cambridge University Press, New York, NY, 1993.
- [36] E. G. Haffner, U. Roth, A. H. II, T. Engel, and C. Meinel. What do hyperlink-proposals and request-prediction have in common? In *Proceedings of the International Conference on Advances in Information Systems (ADVIS)*, pages 275–282, 2000. Springer LNCS 1909.
- [37] T. I. Ibrahim and C.-Z. Xu. Neural Net Based Pre-fetching to Tolerate WWW latency. In *Proceedings of the 20th International Conference on Distributed Computing Systems (ICDCS2000)*, Apr. 2000.
- [38] Imsisoft. Imsisoft home page. <http://www.imsisoft.com/>, 2002.
- [39] Z. Jiang and L. Kleinrock. An adaptive network prefetch scheme. *IEEE Journal on Selected Areas in Communications*, 16(3):358–368, Apr. 1998.
- [40] T. Joachims, D. Freitag, and T. Mitchell. WebWatcher: A Tour Guide for the World Wide Web. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, pages 770–775. Morgan Kaufmann, Aug. 1997.
- [41] P. B. Kantor, E. Boros, B. Melamed, V. Menkov, B. Shapira, and D. J. Neu. Capturing human intelligence in the net. *Communications of the ACM*, 43(8):112–115, Aug. 2000.
- [42] R. P. Klemm. WebCompanion: A friendly client-side Web prefetching agent. *IEEE Transactions on Knowledge and Data Engineering*, 11(4):577–594, July/August 1999.
- [43] B. Krishnamurthy and J. Rexford. *Web Protocols and Practice : HTTP 1.1, Networking Protocols, Caching, and Traffic Measurement*. Addison-Wesley, 2001.
- [44] T. M. Kroeger, D. D. E. Long, and J. C. Mogul. Exploring the Bounds of Web Latency Reduction from Caching and Prefetching. In *Proceedings of the USENIX Symposium on Internet Technologies and Systems (USITS '97)*, Dec. 1997.
- [45] H. Lieberman. Letizia: An Agent That Assists Web Browsing. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'95)*, pages 924–929, Montreal, Aug. 1995.
- [46] H. Lieberman. Autonomous Interface Agents. In *Proceedings of the ACM SIGCHI'97 Conference on Human Factors in Computing Systems*, Atlanta, GA, Mar. 1997.
- [47] V. N. Padmanabhan and J. C. Mogul. Using predictive prefetching to improve World Wide Web latency. *Computer Communication Review*, 26(3):22–36, July 1996. Proceedings of SIGCOMM '96.
- [48] M. Pazzani, J. Muramatsu, and D. Billsus. Syskill & Webert: Identifying Interesting Web Sites. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, Portland, OR, 1996.
- [49] M. J. Pazzani and D. Billsus. Learning and Revising User Profiles: The identification of interesting web sites. *Machine Learning*, 27:313–331, 1997.
- [50] PeakSoft Corporation. PeakJet 2000 web page. <http://www.peaksoft.com/peakjet2.html>, 2002.
- [51] J. E. Pitkow and P. L. Pirolli. Life, Death, and Lawfulness on the Electronic Frontier. In *ACM Conference on Human Factors in Computing Systems*, Atlanta, GA, Mar. 1997.
- [52] M. F. Porter. An algorithm for suffix stripping. In K. Sparck Jones and P. Willet, editors, *Readings in Information Retrieval*. Morgan Kaufmann, San Francisco, 1997. Originally published in *Program*, 14(3):130-137 (1980).
- [53] A. Rousskov. Web Polygraph: Proxy performance benchmark. Available at <http://www.web-polygraph.org/>, 2002.
- [54] A. Rousskov, M. Weaver, and D. Wessels. The fourth cache-off. Raw data and independent analysis at <http://www.measurement-factory.com/results/>, Dec. 2001.
- [55] R. R. Sarukkai. Link prediction and path analysis using Markov chains. In *Proceedings of the Ninth International World Wide Web Conference*, Amsterdam, May 2000.
- [56] L. Tauscher and S. Greenberg. How People Revisit Web Pages: Empirical Findings and Implications for the Design of History Systems. *International Journal of Human Computer Studies*, 47(1):97–138, 1997.
- [57] Web 3000 Inc. NetSonic Internet Accelerator web page. <http://www.web3000.com/>, 2002.
- [58] G. M. Weiss. Predicting telecommunication equipment failures from sequences of network alarms. In W. Kloesgen and J. Zytkow, editors, *Handbook of Data Mining and Knowledge Discovery*. Oxford University Press, 2001.
- [59] Y. Yang, S. Slattery, and R. Ghani. A study of approaches to hypertext categorization. *Journal of Intelligent Information Systems*, 18(2/3):219–241, 2002.
- [60] I. Zukerman, D. W. Albrecht, and A. E. Nicholson. Predicting users' requests on the WWW. In *Proceedings of the Seventh International Conference on User Modeling (UM-99)*, pages 275–284, Banff, Canada, June 1999.