

BrowseLine: 2D Timeline Visualization of Web Browsing Histories

Orland Hoerber

Department of Computer Science
Memorial University
St. John's, NL, Canada A1B 3X5
hoeber@cs.mun.ca

Joshua Gorner

Faculty of Engineering & Applied Science
Memorial University
St. John's, NL, Canada A1B 3X5
gorner@enr.mun.ca

Abstract

Re-finding previously viewed Web pages in browsing histories is often a difficult task, due to the incomplete and vague knowledge people have about the information they are seeking. In this paper, we present a visual interface for the task of re-finding Web pages within browsing histories. BrowseLine employs a novel two-dimensional timeline metaphor, allowing users to visually identify temporal patterns within their browsing histories. These visual patterns can be matched to the users' recollection of their browsing activities, allowing them to jump to a time interval in their browsing history for further investigation. Preliminary evaluations of BrowseLine have found that users can readily grasp the two-dimensional timeline representation, and can use the system effectively to re-find previously viewed pages.

1 Introduction

Maintaining and viewing Web browsing histories has been an important part of browsers as early as NCSA Mosaic [1]. Some studies have suggested that over 80% of the pages viewed by Web users are ones they have previously visited [4]. While many of these page views are of well-known or bookmarked sites, there is a clear need for people to be able to easily re-find a page they have previously viewed. The goal of this research is to support this activity in a visual manner.

The history mechanisms in modern Web browsers commonly employ a list-based representation of the information. These interfaces often include interactive features that allow the users to re-sort the pages based on title, URL, and view date. However, they have the same problems as most other list-based representations: users must evaluate the textual contents in a sequential manner.

An alternate method for users to re-find Web pages is to start typing the title or URL of the page in the browser

location bar. In most modern browsers, an auto-complete feature provides a list of pages from the history that match the terms entered. This lightweight search feature can be very effective when users are able to remember specific details about the title or URL of the page they are seeking. However, when users are unable to recall such information, this feature is of limited value.

Our preliminary studies have found that people who are seeking to re-find a previously viewed Web page use the history features provided by the Web browser as a last resort. More commonly, when they are unable to recall specific details about the page they are seeking, they attempt to re-trace the steps they took to arrive at the page. Examples include people re-constructing their Web search queries, or attempting to follow the links from some known starting point. Clearly, the success of these methods depends heavily on the ability of the user to recall the details of their past browsing activities.

In this work, we propose a visual representation of Web browsing histories, based on a novel two-dimensional timeline metaphor. Our fundamental hypothesis is that while people may not be able to remember the exact details of a Web page they wish to re-find, they may be able to remember how much other Web browsing they were doing at the time, and an approximate time interval during which they first viewed the page. For example, one might recall that they viewed a page sometime last week, in the afternoon, and that it was one of the pages found during a Web search activity. BrowseLine supports the visual identification of these temporal features of the browsing data, allowing the user to easily jump to the location in the browsing history to evaluate the viewed pages in more detail. See Figure 1 for a screenshot of the system. The details of the design and features of BrowseLine are outlined in this paper.

The remainder of the paper is organized as follows. Section 2 provides a brief overview of related work on visualizing browsing histories. A characterization of the data and the tasks is provided in Section 3. Section 4 outlines the key features and implementation details of BrowseLine.




Figure 1. A screenshot of BrowseLine. Note the main timeline view (left) and the timeline navigation view (right). The yellow indicator at the top of the timeline navigation view shows where/when the main timeline view is focused.

Section 5 elaborates on how different patterns of browsing produce different visual patterns in the system. The results of a preliminary evaluation are summarized in Section 6. The paper concludes with a summary of the contributions and an outline of future work in Section 7.

2 Related Work

A number of visual approaches to representing browsing histories have been documented in the literature. In this section, we highlight those techniques that have influenced the design of BrowseLine.

Given the connected nature of the Web, many systems that have endeavored to provide a graphical representation of browsing histories have done so using a tree-based representation [2, 3, 5, 6, 7, 9]. While trees are an effective means for showing the paths of browsing activities, it is not clear whether following these paths provides sufficient support for re-finding previously viewed pages. As such, we rejected the use of tree-based representations in favour of following a timeline metaphor.

The visual representation of timelines have been explored within the domains of medical data [12] and news

coverage [8], as well as in seminal works on visual design [16]. There are limited examples of using temporal representations for Web browsing histories. WebView [3] included a chronological ordering of the page views, in addition to a tree-based layout. Unfortunately, the temporal ordering in this system does not provide any visual cues with respect to the intervals between page views. In the History-Centric Browsing (HCB) system [14], two coordinated timelines are used: one shows only Web searches (which are used as guides); the second shows all the page views between these search activities. No evidence is given as to whether the intervals between Web searches is a useful method for re-finding Web pages.

A common feature among alternate browsing history representations is to include thumbnails of the viewed pages [2, 3, 6, 7, 9]. This feature allows the user to visually identify a page, taking advantage of the human capacity to quickly recognize previously viewed images [15, 17]. Other features that have appeared in these systems include filtering [7], domain stacking [7], graphical indications of the number of page views [3], and zoomable views of the information space [9]. We have included some aspects of all of these features in the design of BrowseLine.

3 Data and Task Characterizations

Before we provide more detail on the design of BrowseLine, it is useful to specify and characterize features of the data to be represented, as well as discuss the task of re-finding previously viewed Web pages. This analysis has guided the development of the two-dimensional timeline that forms the primary contribution of the system.

We characterize the data generated from the browsing activities of a single user as *bursty*. That is, people don't normally view Web pages in a regular pattern. Instead, they may view many pages over a short period of time (e.g., when searching the Web), or view pages once in a while over the course of the day (e.g., checking Web-based email), or view a small collection of pages at seemingly random times (e.g., browsing activities). As such, the data is often *sparse*, with numerous short intervals of browsing activities that are of varying density (from one page to tens of pages).

The user-oriented task that is the focus of this research is to re-find some previously viewed Web page. If the user is able to fully specify the page they are looking for, there is no need to use the browsing history; the user can simply enter the URL into the browser, select the page from their bookmarks, or navigate to the page from some known starting location. As such, our focus is on those situations where the user has incomplete or vague knowledge about the page they are seeking.

While the user may not remember the title of the page, the URL, nor the steps taken to get to the page, we postulate that they will often be able to recall vague temporal information as well as the approximate density of browsing activity when the page was viewed. Therefore, our goal in designing BrowseLine is to support users in re-finding Web pages given these vague and incomplete specifications of their information need.

4 Key Features of BrowseLine

4.1 Two-Dimensional Timeline

The primary design element in BrowseLine is the use of a novel two-dimensional timeline representation. Traditional timeline representations array the data points along a single dimension. While this can be very effective when there is a nearly continuous data stream over time (e.g., news articles [8]), there are difficulties when the data is sparse, yet bursty. With sparse data, one might choose to compress the timeline in order to pull the data points closer together. However, this becomes very problematic during the bursty intervals where there are a large number of data points to represent within a compressed time interval.

Our solution to this problem is a two-dimensional timeline, representing *macro-time* and *micro-time* as orthogo-

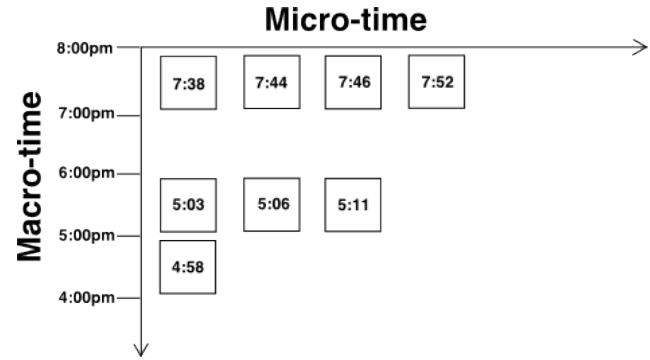


Figure 2. The two-dimensional timeline preserves the space-time relationship on the vertical dimension, but breaks it on the horizontal dimension to produce a dense packing of the data.

nal dimensions (see Figure 2). In the current implementation for representing Web browsing histories, the *micro-time* slots encompass one hour of data; however they could be arbitrarily large or small depending on the features of the data. The *macro-time* slots run in the vertical direction, with movements up representing positive movements in time. The *micro-time* slots run in the horizontal direction, with movements to the right representing positive movements in time.

Within the *macro-time* dimension, the spacing of *micro-time* slots preserves the relationship between vertical space and time. That is, a pair of time slots that are four hours from one another are presented four vertical space units away. If there are no data elements between the pair, the spacing is preserved. However, within the *micro-time* dimension, the relationship between horizontal space and time is not preserved. That is, the page views within the one-hour *micro-time* slot are pushed together and presented in a sequence rather than a true timeline. The result is essentially a temporal histogram.

Preserving the space-time relationship in the *macro-time* dimension, but breaking it in the *micro-time* dimension, follows from our hypothesis that a user's ability to recall when they performed their browsing activities will only be accurate at the *macro-time* level of detail. That is, people will be able to recall general notions of the time during which they were browsing the Web (e.g., "sometime in the afternoon", "around 10:00 PM", or "shortly after work"), but will not be able to recall the exact time of a page view nor the time between page views. As such, we preserve the space-time relationship at the level of detail that the users will be able to recall (*macro-time*), and make more efficient use of the space at the level of detail that they will not be able to accu-

rately remember (*micro-time*).

4.2 Domain Stacking

Even though choosing not to preserve the space-time relationship in the horizontal dimension results in a more efficient use of horizontal space, vertical scrolling may still be required when someone views a large number of pages during the course of a single *micro-time* slot. To address this problem, we have included a method for performing domain stacking.

For each unique domain viewed during a given *micro-time* slot, a single object is included on the horizontal timeline. When subsequent pages are viewed within this same domain, they are stacked on top of the previous page views. Different shades of blue on a perceptually ordered scale (light to dark) are used to represent how many pages are included in each stack. The stacks include the domain icon provided by the Web page, as well as the domain name (which may be trimmed due to space limitations). Clicking on the stack reveals the details of the individual page views (as described in the following section).

While this domain stacking has the negative effect of taking pages out of the order in which they were viewed, we don't foresee this being a problem for users. We hypothesize that when people are recalling their browsing histories, the interactions within a single Web site are combined into a single cognitive cluster. Therefore, rather than preserving the true ordering of page views, stacking the domains upon one another more closely matches the user's mental model of their browsing activities.

4.3 Thumbnail and Meta-Data

When a domain stack is clicked, a list of the pages contained within the stack are displayed in a drop-down window (see Figure 3). For each page viewed, a thumbnail of the page is shown, along with meta-data that includes the title, URL, view date, and view time. This list supports both the visual recognition of a previously viewed page, as well as the browsing of the textual meta-data. Clicking on the thumbnail opens the target page in the browser.

The ability to browse the pages within a domain stack is expected to be more efficient than browsing within the complete list of page views. Because many of the pages will have similar visual layouts (having come from the same domain), the user will be able to easily identify whether this is the domain they are seeking, and quickly reject it if this is not the case. Further, due to the relatively small number of page views within the domain stack, the user will be able to quickly scan the list, either finding the page they were looking for, or moving on to investigate another domain stack or another location in the timeline.




Figure 3. A view of the thumbnails and meta-data, after clicking on a domain stack.

4.4 Timeline Navigation View

The timeline navigation view is an abstract, zoomed-out view of the two-dimensional timeline, providing a coordinated view of the information space [13]. Approximately four days can be viewed at once within this view, as opposed to 10 hours in the main timeline representation. Small rectangular icons are used to represent the domain stacks within each *micro-time* slot; the colour of these icons is coordinated with the colours of the domain stacks, providing orientation cues to the viewer.

A yellow indicator within the timeline navigation view shows where in the *macro-time* dimension the main timeline view is focused. Clicking or dragging this indicator allows the user to easily jump to a location within the timeline that is of interest. Further, as users scroll the main timeline view in the normal manner, the indicator is dynamically updated.

4.5 Domain Filtering

A final feature that we believe will support the users in re-finding previously viewed pages is the domain filter. As users type portions of a domain name in the filter textbox, a visual indication is provided to highlight those domains that match the filter. The matched domain stacks are rendered with a dark border, and all other domain stacks are faded away. This occurs in both the main timeline view as well as the overview. The resulting effect is that the matched domains "pop out" of the display due to the high luminance contrast between foreground and background; the opposite effect occurs with the unmatched domains.

4.6 Implementation

The BrowseLine prototype has been implemented as an extension to the Firefox Web browser. The interface is generated using XUL (XML user interface language), CSS

(cascading style sheets), and JavaScript [10]. All data except for the thumbnails are retrieved from the internal Firefox browsing history management system.

5 Visual Patterns of Browsing

Our design goal for including the zoomed-out timeline navigation view within BrowseLine was to support users in visually identifying patterns within their browsing activities. We have categorized three such patterns that may be common in users' browsing habits. These are depicted in Figure 4.

A *high-density burst of browsing activity over a short period of time* will be represented by a horizontal line of domain stacks in the timeline navigation. The type of Web use activity that will produce such a pattern includes Web searching and browsing over multiple domains. If the user is wishing to re-find a page they viewed as part of a Web search, this visual pattern will be of interest.

A *medium-density burst of browsing activity over an extended period of time* will appear as a vertical line of domain stacks. Within this burst, the pattern will include multiple horizontal domain stacks. Such patterns will be produced by users making extended use of an online service, such as Web mail, online gaming, and Web video services. If the user wishes to re-find a page from an online video service, this visual pattern will be of interest.

A *low-density burst of browsing activity over a short period of time* will appear as a single domain stack without any others nearby. Such patterns may appear when the user follows a link from an external application such as an instant messaging client or a desktop email client, or when they manually type a URL into the browser. If the user wishes to re-find Web pages that were viewed independent of any other browsing activities, this visual pattern will be of interest.

In each of these cases, once the user visually identifies a pattern that represents their recollection of the browsing activity that coincided with the page they wish to re-find, they can jump to that location in the timeline by clicking in the vicinity of the pattern. Then, within the main timeline view, the user may browse the domain stacks, view the thumbnails and meta-data to find the desired page.

6 Preliminary Evaluations

Preliminary evaluations of the BrowseLine prototype were conducted with four graduate students in an informal setting. The students were asked to install and use the system for a period of one day, and then participate in a focus group discussion about the system. Our primary purpose in conducting this focus group was to determine

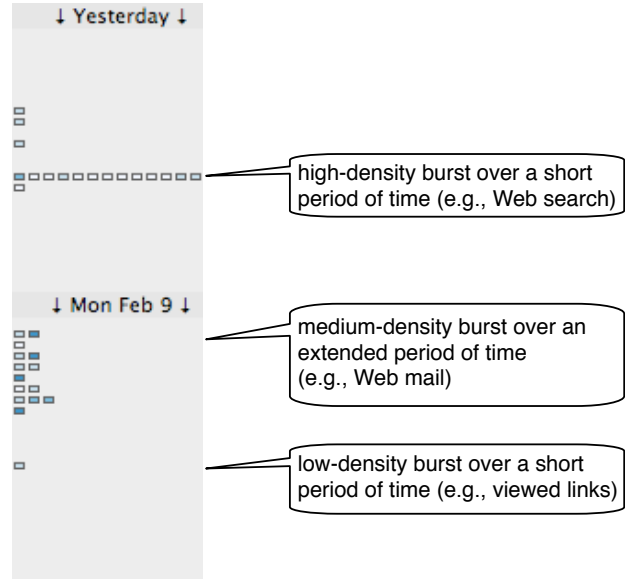


Figure 4. Different visual patterns in the timeline navigation view indicated different patterns of Web browsing activity.

whether the participants were able to grasp the nature of the two-dimensional timeline metaphor, and whether they found BrowseLine to be effective for re-finding previously viewed Web pages.

While there was discussion about a number of minor bugs the participants found, the general sense of the system was positive. The participants didn't have a problem with the two-dimensional timeline representation, nor the fact that the *micro-time* dimension broke the space-time correspondence. They all liked the ability to visually identify patterns of browsing activities. A number of additional features and changes were recommended, all of which will be considered for future versions of the system. Half the participants indicated that they are planning on continuing to use BrowseLine; the other half indicated that their reason for not wanting to continue to use it was due to the load times when the browsing history is long. We are currently working to resolve the issue with the load time.

7 Conclusions & Future Work

In this work, we have introduced BrowseLine as a method for supporting users in the task of re-finding Web pages in their browsing histories. The primary contribution of this work is the use of a two-dimensional timeline to represent data at *macro-time* and *micro-time* levels of detail. A timeline navigation view provides a zoomed-out rep-

resentation of the temporal characteristics of the browsing history, supporting the visual recognition of patterns that match the users' recollection of their browsing activities. Once an approximate area in the timeline is identified, users can quickly jump to this location and evaluate the contents of the domain stacks as they seek to re-find a Web page.

Although our preliminary investigations have provided positive feedback, there remain a number of unverified hypothesis that guided the design of BrowseLine. A comprehensive user evaluation in a controlled laboratory setting is currently in the planning stages. Field trials and real-world user evaluations are also under consideration.

Other future work for BrowseLine includes merging the the timeline navigation view and the main timeline representation using a focus+context technique similar to WaveLens [11], making improvements and refinements to the interface, and releasing the BrowseLine Firefox extension to the public. We are also investigating applying the two-dimensional timeline representation to other domains with a strong temporal aspect, such as medical informatics, news feeds, email, and blog entries.

Acknowledgements

This research has been made possible through the first author's Discovery Grant and the second authors Undergraduate Student Research Award, both provided by the Natural Science and Engineering Research Council of Canada (NSERC).

References

- [1] M. Andreessen. NCSA mosaic technical summary. Technical report, National Center for Supercomputing Applications, February 1993.
- [2] E. Z. Ayers and J. T. Stasko. Using graphic history in browsing the World Wide Web. In *Proceedings of the Fourth International World Wide Web Conference*, pages 11–14, 1995.
- [3] A. Cockburn, S. Greenberg, B. McKenzie, M. Jasonsmith, and S. Kaasten. WebView: A graphical aid for revisiting Web pages. In *Proceedings of the OZCHI'99 Australian Conference on Human Computer Interaction*, 1999.
- [4] A. Cockburn and B. McKenzie. What do Web users do? an empirical analysis of Web use. *International Journal of Human-Computer Studies*, 54(6):903–922, 2001.
- [5] P. Dömel. WebMap: a graphical hypertext navigation tool. In *Proceedings of the Second International World Wide Web Conference*, pages 785–789, 1994.
- [6] E. Frécon and G. Smith. WebPath - a three-dimensional Web history. In *Proceedings of the 1998 IEEE Symposium on Information Visualization*, pages 3–10, 1998.
- [7] R. Gandhi, G. Kumar, B. Bederson, and B. Shneiderman. Domain name based visualization of Web histories in a zoomable user interface. In *Proceedings of the 11th International Workshop on Database and Expert Systems Applications*, pages 591–598, 2000.
- [8] S. Havre, E. Hetzler, P. Witney, and L. Nowell. ThemeRiver: Visualizing thematic changes in large document collections. *IEEE Transactions on Visualization and Computer Graphics*, 8(1):9–20, 2002.
- [9] R. R. Hightower, L. T. Ring, J. I. Helfman, B. B. Bederson, and J. D. Hollan. PadPrints: Graphical multiscale Web histories. In *Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology*, pages 121–122, 1998.
- [10] Mozilla Foundation. The joy of XUL. Online (https://developer.mozilla.org/en/The_Joy_of_XUL), September 2007.
- [11] T. Paek, S. Dumais, and R. Logan. Wavelens: A new view onto internet search results. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, 2004.
- [12] C. Plaisant, B. Milash, A. Rose, S. Widoff, and B. Shneiderman. Life Lines: Visualizing personal histories. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 221–227, 1996.
- [13] J. C. Roberts. State of the art: Coordinated & multiple views in exploratory visualization. In *Fifth International Conference on Coordinated and Multiple Views in Exploratory Visualization*, pages 61–71, 2007.
- [14] Y. Shirai, Y. Yamamoto, and K. Nakakoji. A history-centric approach for enhancing Web browsing experiences. In *CHI '06 Extended Abstracts on Human Factors in Computing Systems*, pages 1319–1324, 2006.
- [15] L. Standing, J. Conezio, and R. N. Haber. Perception and memory for pictures: Single-trial learning of 2500 visual stimuli. *Psychonomic Science*, 19(2):73–74, 1970.
- [16] E. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, 2nd edition, 2001.
- [17] C. Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann, 2nd edition, 2004.