White Rose Research Online

eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk
Universities of Leeds, Sheffield and York

Deposited via The University of Leeds.

White Rose Research Online URL for this paper:
https://eprints.whiterose.ac.uk/id/eprint/74330/

Article:
Do, TV and Ruddle, RA (2012) The design of a visual history tool to help users refind information within a website. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 7224. 459 - 462 . ISSN: 0302-9743

https://doi.org/10.1007/978-3-642-28997-2_41

Reuse
See Attached

Takedown
If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

The Design of a Visual History Tool to help Users Refind Information within a Website

Abstract. On the WWW users frequently revisit information they have previously seen, but "keeping found things found" is difficult when the information has not been visited frequently or recently, even if a user knows which website contained the information. This paper describes the design of a tool to help users refind information within a given website. The tool encodes data about a user's interest in webpages (measured by dwell time), the frequency and recency of visits, and navigational associations between pages, and presents navigation histories in list- and graph-based forms. A design evaluation, conducted using examples of refinding failures from a previous user study, shows that list-based presentation helps when users have only visited a small number of pages on a website (this accounted 50+ % of refinding failures). With the tool's filtering capabilities, graph-based presentation is likely to greatly increase success during revisiting when a website has often been used.

Keywords: Web History, Navigation, Revisiting, Visualization, Information Retrieval.

1 Introduction

Finding information on websites and intranets is central to many people's everyday life. Almost half of the webpages people look at are ones they have found before [29], but the task of refinding information is typically only straightforward if a page has been looked at either frequently or very recently [11]. If the opposite is true (e.g., a page that was looked at once, a week or more ago) then refinding fails on 20% of occasions, even when the user does remember which website contains the page [2].

This paper presents the design of a tool to help users refind information on a given website, via list- and graph-based presentation of their web navigation history. List-based presentation is suitable if a user has previously only visited a small number of pages on a website, whereas the graph-based approach provides richer information and scales better to revisiting on websites that the user often uses. Both types of presentation facilitate easy browsing of a user's navigation history, taking advantage of users' tendency to want to browse rather than search if they can recall previously visiting a webpage [26].

The remainder of the present paper is divided into three main parts. First, related research is summarized. Next, the tool design is described, focusing particularly on how history information is encoded for list- and graph-based presentation. Then a design evaluation is conducted, to assess how the tool is likely to have simplified the task of refinding, using examples of refinding failures from a previous user study.

2 Related Work

Ever since the early years of the World Wide Web, the extent to which webpages are revisited has been examined [5]. Subsequent studies [2, 7, 24, 29] reveal that more than one third of visits are revisits to webpages that have been previously seen. Some research [17, 25] explored how people revisit webpages in general, and other research has investigated specific aspects of revisiting such as the types of cue that are useful, and how users' web navigation history could be filtered, ordered and presented.

Useful cues for revisiting task were also examined in some research. Information such as a webpage's title, URL, and thumbnail was evaluated in [18]. Anchor text (the text of an hyperlink) was addressed in [8]. Navigational paths were considered in [20]. Thumbnails were proved to be most important cue for user recognition while anchor texts and navigation paths are useful contextual information. They are particularly helpful when people revisit a webpage by making a series of small "orienteering" steps [25].

Dwell time on a webpage is also an important factor for revisiting because previous research [12, 30] has shown that, during search and navigation activities, a dwell time of 30 seconds is indicative of a page that is of interest to a participant (i.e., informational rather than just navigational). A three month weblog revealed that more than two thirds of visits were navigational based on that criteria [2], so using this feature as a filter could substantially reduce the amount of history information that needs to be considered during revisiting.

Beside information about the webpages themselves, some tools use frequency to make suggestions to users. Browsers display a list of top frequently visited webpages either in a dropdown list (Firefox, Internet Explorer (IE)) or in a new blank tab (Google new tab, Chrome). The history list of web browsers (IE, Firefox, Chrome) also allows previously visited pages to be sorted by site name, frequency or recency.

The tools that have been developed to support revisiting adopt a list-based approach, spatial metaphor, or graph-based approach to present users' navigation history. A list-based approach is found in URL auto-completion [19], bookmarks/favourites [1], history lists [24], Google History [14], Infoaxe [16], WebMynd [28], and CWH [31]. Typically, these tools display a small icon of the website (a favicon) or thumbnail of a webpage, and text for the URL, title, frequency of visits, date of the last visit, and a synopsis of the webpage's content.

Spatial metaphors such as a bookcase or desktop are used by WebBook and Web Forager [4]. Data mountain [22] exploits the use of spatial memory for data management, because users may arrange thumbnails of webpages arbitrarily on an inclined 2D plane like a mountain in 3D environment. DeckScape [3] presents a set of web pages as if they are a deck of cards, of which only one is visible at a time.

Webmap [9] and PadPrints [15] were two of the first graph-based tools. They used spanning trees to present web browsing history. The former uses color-coding to distinguish normal nodes (blue), current node (red), and mouse over node (pink). Nodes are labeled by order of visit. Links between nodes are also color-coded: black for intra-site, green for inter-site, and dashed for movements not part of the spanning tree. A list view is also used to show title and URL for each node of the tree. The later

improve the former by using thumbnails of webpages as nodes. One of the main problems of these two tools is the tree sizes grow rapidly.

Domain Tree Browser [13] solved PadPrints' limitation by dividing the presentation into three panes: domain pane, tree pane, and webpage pane. Color-coding is used to highlight current domain and current page. Thumbnail sizes reflect the frequency of visits to a webpage. With this approach, each domain (website) is represented by a separate tree to deal with scale problem. However the way a new node is added to a tree might confuse the users because that node is always attached to the last visited node regardless of user navigational paths.

Nestor Navigation [10] and WebNet [6] draw 2D directed graphs of a user's navigational paths. The graph will be drawn as a straight line of connected nodes if the user does not backtrack. Otherwise, if the user backtracks and follows new directions, the graph will branch. Loops can also happen. Depending on a user's preferences, node sizes can represent frequency, distance, or recency. The drawback of these two tools is the graph will become very busy and complicated over time.

To simplify the complexity of the directed graph, SessionGraphs [21] visualizes a user's path as an animated, interactive graph for a session. The user needs to define a task, and then graphs for sessions are attached to the task. Each session can be view in another pane for detailed examination. Color-coding is used to mark currently view pages, selected pages, and mouse over page. Node sizes grow logarithmically with the number of visits to pages.

Inspired partly by the methods of presentation used in Google search and Google maps, which are familiar to millions of WWW users, the designs described in this paper employ list-based and graph-based approaches. However, our designs are novel because of the way that the information for display is chosen and encoded for presentation to users.

3 Tool Design

This section describes the design of a tool that aims to make it straightforward for users to return to any page they have previously visited within a given website. The following sections describe the tool's architecture, how a user's navigation history is presented in a list-based form (see Fig. 2) and a twin-view (list and graph; see Fig. 4) form, and how users interact. The list-based form adopts the style of existing search results listings, and the twin-view form is inspired by the approach used by applications such as Google maps.

3.1 Tool Architecture

An add-on for Mozilla Firefox has been developed to capture an individual's web browsing history, and has been tested and used on Windows, Mac and Linux computers. To improve performance of data queries, data from each day are stored in a separate file (daily file), the name of which is generated automatically based on the date. A daily file stores webpage information (thumbnail, URL, title, and hyperlinks

with their anchor texts) and navigational information (visited time, URL and ID of referrer, anchor texts used to access the page, and dwell time). An index file is also created and, each unique URL in the individual's history, stores the page title, number of visits (frequency), day count (how many different days a webpage has been visited), and pointers to the corresponding daily files. All the data files are stored in XML format in the user's login profile.

The tool's user interface has three panes: a top pane for search and filtering a history, the left pane for list-based presentation, and the right one for the graph view. Users can interactively modify the area used by the left and right panes (See Fig. 4).

For list-based presentation an XPath is performed on the index file to obtain pointers to the relevant daily files and, from them, retrieve the history information of the webpages that are to be presented. For graph-based presentation, a graph of all a user's navigation within the website is created from webpage "referrer" information (i.e., from the user's navigational path, which takes account of branches and multiple browser tabs), a nominal root is added to linking to the start node of every paths, and then a spanning tree is derived. The link weight calculation is described below.

3.2 Presentation of a User's Navigation History

Each visited webpage has a number of attributes that may be visually encoded, regardless of the presentation form. The most important attributes are: interest, frequency, recency and association.

Interest: The time which users spend on a page reveals indicates the extent to which they are interested in it. Previous research [12, 30] has shown that, during search and navigation activities, a dwell time of 30 seconds is indicative of a page that is of interest to a participant (i.e., informational rather than just navigational). Another study [2] reported that only 31% of visited webpages were informational and, when people tried to revisit a webpage which had not been visited recently or frequently, that webpage often was informational. Therefore, distinguishing informational from navigational webpages will help users find the information they are looking for.

Frequency: Frequency can be calculated in terms of the number of times a webpage was visited, or the number of sessions or days on which it was visited. Although people may not remember the exact number of visits, they can typically recall that they have been there only once vs. twice vs. more often. This attribute may be used for sorting or filtering a user's navigation history, but it can also be encoded within a visualization tool so that frequently visited pages can be seen in the context of all of a user's navigation.

Recency: This information concerns either the last time a page was visited or the pattern of all visits over a period of time. For histories that cover several years, people can generally only pinpoint when a particular event occurred to the accuracy of several months [27].

Association: Showing webpages that were visited around the time, or during the same session, that a given webpage was visited provides context and may help users find the information they are seeking. This is particularly true when users revisit a webpage by making a series of small "orienteering" steps [25].

List-based presentation.
Because Google search has become so familiar with internet users, our design adopts the style of Google search results to present a webpage's title, URL, description, frequency ("You've visited this webpage X times") and recency ("Last visited …"). A basic listing is enriched by adding a small thumbnail, using the same approach as [31]. This conveys the layout of a webpage, which proved to be useful for users' recognition [18, 31] and is becoming increasingly common with search engines such as Google. Considering the tradeoff between the amount of recognizable detail vs thumbnail size (which dictates how many thumbnails can be displayed at the same time), a thumbnail size of 235x148 has been chosen [18, 31].

Two new features are also included. First, to provide association data, the entry for each page can be expanded to list as bullet points the title of every webpage that was navigated to directly from that page. This provides users with contextual information, to improve the precision with which they need to recall a webpage. Second, the title of a webpage is color-coded to indicate whether it is navigational (bold blue) or informational (bold maroon). An example of one list-based entry is shown in Fig. 1, and a complete list for a website is shown in Fig. 2.

Fig. 1. List-based presentation of a user's navigation history

Graph-based presentation .
Graph-based presentation uses a tree, because trees are the type of graph that is most familiar to people in general, trees are often used to reflect the primary structure of a website [23], have been used in many previous studies (e.g., [9, 13, 15]), and are fast to compute. To calculate the tree, the weight of each link in a user's overall navigation graph (see above) is set to $(\text{numVisitsToNode} / \text{numVisitsViaThatLink})^2$ and then Dijkstra's shortest path algorithm is applied. This weight formula helps to ensure that the tree includes the links that the user most often traversed, and this information is encoded by the thickness of each edge in the tree.

To keep the graph-based and the list-based presentation consistent, each node of the tree is represented by a webpage thumbnail and, using the same colors as list-based presentation, the border of navigational and informational thumbnails are blue and maroon, respectively. Similar to WebNet[6] and SessionGraphs [21], frequency of revisits to a page is encoded by node size. The default size of a node is 238x151 pixels (size of the standard thumbnail for the list-based plus 3 pixels for border). If the frequency of visit to a node is less than 6, its width and height are scaled by (1 + (Frequency-1)/10), otherwise its size is the same as node with frequency of 5.

Association is encoded by the tree's links and the ordering of each node's children, which are displayed in order of when they were last visited (most recent at the bottom). However, the disadvantage of a tree is that only a subset of a user's

navigation is presented. To address this, users can expand the tree under a thumbnail to display all pages that were navigated to from the focus page as its descendents.

The tree is presented in a horizontal orthogonal layout. By default, the whole tree is zoomed to fit the display area, but users may also zoom and pan. As with Google maps, list-based presentation is added to the left pane of the dialog to display detailed information about each webpage in the tree (see an example in Fig. 4).

3.3 Interaction

With list- and graph-based presentation, users can define keywords, which will highlight matching pages in the list and graph, but still display non-matching pages to provide contextual information.

Filtering helps users narrow down a results set, so they can find what they want faster. Filters are provided for dwell-time, frequency, day count and recency. With list-based presentation, only pages which meet the criteria are displayed. However with graph-based presentation, the whole tree is recalculated which means, for example, that users can use the filter to display the history tree for a specific browsing session rather than all the navigation they have ever performed on a website.

Highlighting is done by changing the color of list entries or making the border of tree thumbnails bold. With list-based presentation mouseover functionality displays a large thumbnail in the right pane of the tool (see Fig. 2), and with graph-based presentation a dialog box shows larger thumbnail, title, URL, description, frequency, and the last visit of a page. Mouse clicking opens a webpage in the web browser.

Association is shown by the links between webpages in the graph-based presentation, and the ability to “expand” the information shown for each webpage in list-based presentation and as descendents of each webpage in graph-based presentation. Details are described above, and the metaphor is similar to expanding folders in desktop user interfaces (e.g., Windows Explorer).

4 Evaluation

This section describes the evaluation of the list-based and twin-view designs, with weblogs data from a previous user study [2] in which 12 participants recorded their web navigation for three months. During the study each participant took part in a controlled laboratory experiment and was asked to revisit a total of 48 webpages that they had previously visited neither frequently (on only one day) nor recently (between 1 week and 1 month ago). The study was designed to both quantify and characterize the revisiting that takes place during web navigation.

Analysis of the weblogs showed that revisits accounted for 36% of participants navigation, and 18% of those revisits were to pages that participants had previously visited neither frequently nor recently. In the laboratory experiment participants failed to find 20% of the pages they were asked to revisit, despite having access to all of the profile and history information they used during everyday web navigation (for the experiment, they logged into their own account). A major category of failure was

when participants correctly recalled that a target page belonged to a particular website, but could not find the page when they browsed/searched that website again. The following sections analyze how refinding would have been made easier if participants had had the tool described in this paper.

4.1 All Failures

This section summarizes how the tool would have made revisiting easier for every "known website" failure that occurred in the previous user study's experiment (see Table 1). Before the experiment, participants had visited from 2 to 71 different pages on the websites. However, two characteristics that were typical of the failures were that participants: (a) only looked at a few pages on the site before thinking that they could not find target pages and giving up, and (b) looked at as many new pages as pages that they had visited before.

Table 1. Number of pages visited before and/or during a revisiting experiment [2], for each "known website" failure that occurred. The history tree columns show the depth and maximum width at any depth of a participant's navigation tree of the website before the experiment.

Failure	Number of pages visited			History tree	
	Only before experiment	Before & during experiment	Only during experiment	Depth	Width
1	1	1	2	2	1
2	2	1	3	2	2
3	2	1	0	1	3
4	2	2	1	2	3
5	2	2	3	2	2
6	8	1	2	3	7
7	9	2	2	7	4
8	1	5	1	2	4
9	9	3	1	6	3
10	28	3	2	4	13
11	31	3	3	9	14
12	60	1	2	3	49
13	46	10	6	7	16
14	67	4	4	7	22

Trees for graph-based presentation have been created for all 14 failures, and the two last columns of Table 1 show the depth and width of each tree. Most of the trees have a low aspect ratio, but for Failure 12 it is very large (1:16) because this participant

mostly navigated to the website from Google search and forums. Failures 7 and 14 (the tree with the largest number of pages) are used as case studies below to illustrate list- and graph-based presentation, respectively.

4.2 List-based Presentation

Failure 7 is used to illustrate list-based presentation because it is neither trivial (the participant had previously visited 11 pages on the global (nintendo.com) and local website (nintendo.co.xx) of Nintendo) nor overly complex. When the participant was asked to find a page about NINTENDO 3DS hardware again, he went to Google search to navigate to Nintendo's global website, but the required information was on the local website. He browsed to three other webpages (one was visited before and the other two not), several times went back and forth trying carefully to locate the information, but could not and then gave up.

Fig. 2. List-based presentation of Failure 7. By using the tool's mouseover functionality a user can display a detailed thumbnail (right pane) of every webpage in the list (left pane).

If the participant had had the tool described in this paper, he would have just needed to type Nintendo in the Search website textbox and, in the worst case, would have needed to scan through 11 webpages. By moving the mouse over the list presentation, a full thumbnail of each webpage would have been displayed in the right hand pane (see Fig. 2), making it likely that he would have quickly found the webpage he was trying to revisit. The situation would have been even better if he used the tool's filters to display just informational webpages (dwell time $\geq$ 30s), because the list would only have shown three webpages (the target, and two others).

With other cases, even with a larger number of visited webpages, other filters may help. One thing that is certain is users would be more likely to focus their effort on the pages they had previously visited, and not waste time by unintentionally looking at new webpages.

4.3 Graph-based Presentation

Case 14 was selected to evaluate the graph-based presentation because it has more webpages than any other failure. Before the experiment, this participant had made 119 visits to 71 different webpages on Apple's website. When trying to refind information about an application which helps make presentation easy, she went to the correct Apple website thanks to URL auto-completion of Firefox browser. Then she browsed to seven other webpages (visiting some of them several times), including four that she had not visited before.

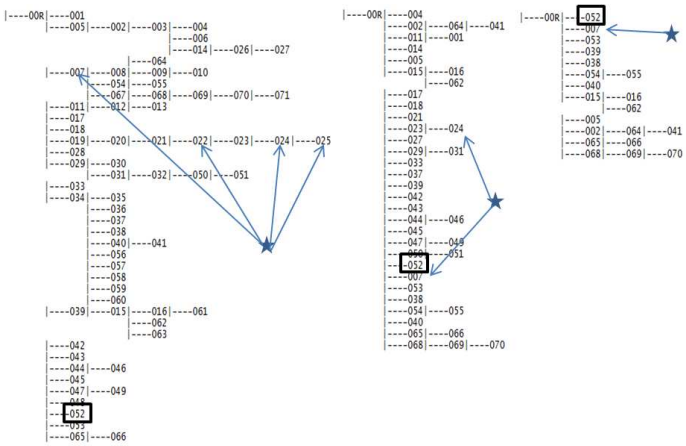

(a) (b) (c)

Fig. 3. Case 14's trees for (a) all pages, (b) filtered by dwell time, and (c) filtered by dwell time & recency

Fig. 3 displays three trees constructed for Case 14 in text-based form: (a) all webpages, (b) filtered by dwell time, and (c) filtered by dwell time and recency. The node in a box was the target webpage and nodes connected to a star were webpages visited by the participant during the experiment (Note that node numbers are not orders of visited time). All trees were at the same scale. It is easy to notice that sizes of the trees reduce substantially after each filter.

The trees also provide better overview of the visited website. If a participant remembers when travelling to the target webpage, he made a navigational path within the website; he should focus on branches with more depth. If he came there directly from other website he should pay attention on first level nodes only. In this case, if the participant had remembered that she reached to the target webpage from another website within a month and spent at least 30 seconds there, she would have had to examine only a dozen of webpages (see Fig. 4).

In other cases, if a particular starting of a navigational path, a special period of time, or estimated number of visits is vaguely recallable, other filters would be useful in making the tree more compact. With that the information space gets smaller.

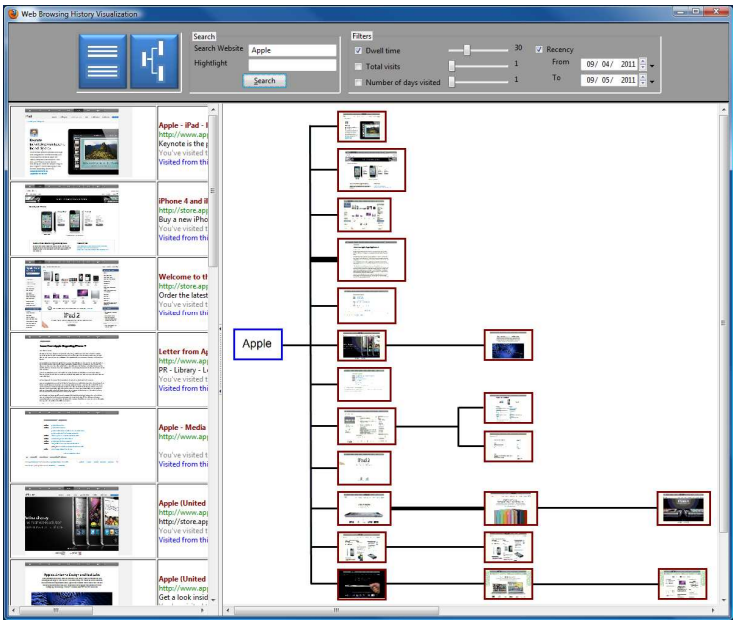

Fig. 4. Twin-view (list & graph) presentation for Case 14 with dwell time and recency filters

5 Conclusions and Future Work

Refinding information is often difficult when the information has not been visited frequently or recently, even if a user knows which website contained the information. This paper describes the design of a visual history tool to assist users in refinding information within a given website. The tool encodes and presents important data (interest, frequency, recency, and associations) from a user's Web history in a list-based or graph-based form. Drawing on the well-known presentation style of Google search, the tool enhances the list-based form by adding a small thumbnail (like [31]), color-codes titles of webpages to indicate whether they were navigational or informational, and provides contextual information by adding the navigational associations between pages. The graph-based form has the same approach as Domain Tree Browser [13] in constructing a tree for visited pages of a website. However, our tool exploits a minimal spanning tree to make the solution more scalable. An important distinction is the tree is constructed from a user's navigational paths (which deal with multiple tabs and sessions) rather than the sequence of visits over time. Links of the tree are also weighted to help ensure that the tree includes the links that the user most often traversed.

An evaluation has been conducted using examples of refinding failures from a previous user study [2]. It showed that list-based presentation is suitable for situations where the user has only visited a small number of pages on a website (this accounted for more than half of the failures in the previous study). Graph-based presentation, coupled with the filtering capabilities of the tool, is likely to greatly increase success

during revisiting when a website has often been used. One thing that is certain is that users would be more likely to focus their effort on the pages they had previously visited, and not waste time by unintentionally looking at new webpages. For the future work, a field study is planned to evaluate how people would actually use the tool and to compare its effectiveness with current history mechanisms such as Google history or Firefox, Chrome, and IE history list.

References

1. Abrams, D., R. Baecker, and M. Chignell, Information archiving with bookmarks: personal Web space construction and organization, in Proceedings of the SIGCHI conference on Human factors in computing systems. 1998, ACM Press/Addison-Wesley Publishing Company: Los Angeles, California, United States. p. 41-48.
2. Anonymous, The underlying causes of failure when users attempt to refind information on the Web. 2011.
3. Brown, M.H. and R.A. Shillner, DeckScape: an experimental Web browser. Computer Networks and ISDN Systems, 1995. **27**(6): p. 1097-1104.
4. Card, S.K., G.G. Robertson, and W. York, The WebBook and the Web Forager: an information workspace for the World-Wide Web, in Proceedings of the SIGCHI conference on Human factors in computing systems: common ground. 1996, ACM: Vancouver, British Columbia, Canada. p. 111-117.
5. Catledge, L.D. and J.E. Pitkow, Characterizing browsing strategies in the World-Wide Web. Computer Networks and ISDN Systems, 1995. **27**(6): p. 1065-1073.
6. Cockburn, A. and S. Jones, Which way now? Analysing and easing inadequacies in WWW navigation. International Journal of Human-Computer Studies, 1996. **45**(1): p. 105-129.
7. Cockburn, A. and B. McKenzie, What do web users do? An empirical analysis of web use. International Journal of Human-Computer Studies, 2001. **54**(6): p. 903-922.
8. Dai, N. and B.D. Davison. Mining anchor text trends for retrieval. in Proceedings of the 32nd European Conference on Information Retrieval. 2010. Milton Keynes, UK: Springer Berlin / Heidelberg.
9. Dömel, P., WebMap: a graphical hypertext navigation tool. Computer Networks and ISDN System, 1995. **28**(1-2): p. 85-97.
10. Eklund, J., J. Sawers, and R. Zeiliger. NESTOR navigator: A tool for the collaborative construction of knowledge through constructive navigation. in The 5th Australian World Wide Web Conference (AusWeb 99). 1999. Lismore.
11. Elsweiler, D., M. Baillie, and I. Ruthven, What makes re-finding information difficult? a study of email re-finding, in Proceedings of the 33rd European Conference on Information Retrieval. 2011, Springer-Verlag: Dublin, Ireland.
12. Fox, S., et al., Evaluating implicit measures to improve web search. ACM Transactions on Information Systems (TOIS), 2005. **23**(2): p. 147-168.
13. Gandhi, R., et al. Domain name based visualization of Web histories in a zoomable user interface. in Database and Expert Systems Applications, 2000. Proceedings. 11th International Workshop on. 2000.
14. Google. Web History. [cited 2009 20/10]; Available from: http://www.google.com/history/.
15. Hightower, R.R., et al., Graphical multiscale Web histories: a study of padprints, in Proceedings of the ninth ACM conference on Hypertext and hypermedia : links, objects,

time and space---structure in hypermedia systems: links, objects, time and space---structure in hypermedia systems. 1998, ACM: Pittsburgh, Pennsylvania, United States. p. 58 - 65.
16. Infoaxe. A Search Engine for Your Web Memory. [cited 2009 20/10]; Available from: http://www.infoaxe.com/.
17. Jones, W., H. Bruce, and S. Dumais, Keeping found things found on the web, in Proceedings of the tenth international conference on Information and knowledge management. 2001, ACM: Atlanta, Georgia, USA. p. 119 - 126
18. Kaasten, S., S. Greenberg, and C. Edwards, How People Recognize Previously Seen Web Pages from Titles, URLs and Thumbnails, in Proceedings of Human Computer Interaction 2001. p. 247-265.
19. Kellar, M., C. Watters, and M. Shepherd, The impact of task on the usage of web browser navigation mechanisms, in Proceedings of Graphics Interface 2006. 2006, Canadian Information Processing Society: Quebec, Canada. p. 235 - 242.
20. Li, J. and Y. Zhao, PathRank: Web Page Retrieval with Navigation Path, in Proceedings of the 31st European Conference on Information Retrieval. 2009, Springer-Verlag: Toulouse, France.
21. Mayer, M. and B.B. Bederson, Browsing Icons: A Task-Based Approach for a Visual Web History, in HCIL Technical Report. 2001.
22. Robertson, G., et al., Data mountain: using spatial memory for document management, in Proceedings of the 11th annual ACM symposium on User interface software and technology. 1998, ACM: San Francisco, California, United States. p. 153 - 162
23. Ruddle, R.A., How do people find information on a familiar website?, in Proceedings of the 23rd British HCI Group Annual Conference on People and Computers: Celebrating People and Technology. 2009, British Computer Society: Cambridge, United Kingdom.
24. Tauscher, L. and S. Greenberg, How people revisit web pages: empirical findings and implications for the design of history systems. International Journal of Human Computer Studies, 1997. **47**(1): p. 97-137.
25. Teevan, J., et al., The perfect search engine is not enough: a study of orienteering behavior in directed search, in Proceedings of the SIGCHI conference on Human factors in computing systems. 2004, ACM: Vienna, Austria. p. 415 - 422.
26. Vallet, D., I. Cantador, and J.M. Jose, Personalizing Web Search with Folksonomy-Based User and Document Profiles, in Advances in Information Retrieval. 2010, Springer Berlin / Heidelberg. p. 420-431.
27. Wagenaar, W.A., My memory: A study of autobiographical memory over six years. Cognitive Psychology, 1986. **18**(2): p. 225-252.
28. WebMynd. Webmynd - A Better Way to Search. [cited 2009 20/10]; Available from: http://www.webmynd.com.
29. Weinreich, H., et al., Off the beaten tracks: exploring three aspects of web navigation, in Proceedings of the 15th international conference on World Wide Web. 2006, ACM: Edinburgh, Scotland. p. 133 - 142.
30. White, R.W. and J. Huang, Assessing the scenic route: measuring the value of search trails in web logs, in Proceeding of the 33rd international ACM SIGIR conference on Research and development in information retrieval. 2010, ACM: Geneva, Switzerland. p. 587-594.
31. Won, S.S., J. Jin, and J.I. Hong, Contextual web history: using visual and contextual cues to improve web browser history, in Proceedings of the 27th international conference on Human factors in computing systems. 2009, ACM: Boston, MA, USA. p. 1457-1466