The Design and Evaluation of a Flick Gesture for 'Back' and 'Forward' in Web Browsers

Michael Moyle **Andy Cockburn**

Human-Computer Interaction Lab
Department of Computer Science
University of Canterbury
Christchurch, New Zealand

{mjm184, andy}@cosc.canterbury.ac.nz

Abstract

Web navigation relies heavily on the use of the 'back' button to traverse pages. The traditional back button suffers from the distance and targeting issues that govern Fitts' Law. An alternative to the button approach is the use of marking menus—a gesture based technique shown to improve access times of commonly repeated tasks. This paper describes the implementation and evaluation of a gesture-based mechanism for issuing the back and forward commands in web navigation. Results show that subjects were able to navigate significantly faster when using gestures compared to the normal back button. Furthermore, the subjects were extremely enthusiastic about the technique, with many expressing their wish that "all browsers should support this". Subjective measures also showed significantly higher ratings for the gesture system over the back button. Finally, subjects found the 'flick' gesture easy to learn.

Keywords: Marking menus, gestures, web navigation, browser design, evaluation.

1 Introduction

Web navigation is an ever-increasing daily activity for millions of users. The introduction of improved web navigation techniques has the potential to yield vast productivity gains, as illustrated by Nielsen's (1993) statement: "The smallest of usability problems, when multiplied across thousands or millions of users, becomes a source of massive inefficiency and untold frustration". In this paper, we show that adding a simple gesture-based navigation facility to web browsers can significantly reduce the time taken to carry out one of the most common actions in web use: navigating back to previously visited pages.

The gestural input mechanism described in this paper is based on marking menus. Introduced by Kurtenbach & Buxton (1991), marking menus allow users to select items from a pie-menu (Callahan, Hopkins, Weiser & Shneiderman, 1988) before the menu is displayed. The user does so by gesturing with the cursor in the direction of the desired item. If the user hesitates in their selection, the full pie-menu is displayed.

The following fictitious scenario drives our evaluation of gesture-based navigation: Microsoft and Netscape have released new versions of their browsers that support gestural navigation for the back and forward actions. We wish to understand whether these new features increase the efficiency of navigation, whether users quickly learn to use them, and whether users appreciate them. Efficiency is evaluated in two experiments that represent common web browsing tasks. Subjects used Microsoft Internet Explorer for all evaluation tasks. The gesture features were seamlessly supported in the browser.

The next section describes related work. We then present the details of our implementation, which runs under any standard unaltered browser. The experimental design and results are then presented and discussed, followed by the conclusions.

2 Backround

2.1 Web Navigation

Several prior studies have shown that navigating back to previously visited pages is an extremely common activity in web use. In a study of three-weeks of client-side web logs, Catledge & Pitkow (1995) found that, on average, each user visited 58% of URLs more than once. In a more recent four-month study, Cockburn & McKenzie (2001) found a revisitation rate of 81%: four out of five page visits are to pages previously seen by the user.

Web browsing applications support many mechanisms for revisiting web pages, including 'favorites' (or bookmarks), history tools, and the back button. The back button is a dominant source of page requests, with Catledge & Pitkow (1995) and Tauscher & Greenberg (1997) reporting that it accounts for 41% and 30% of requests respectively. The forward button, in contrast, was lightly used, accounting for only 2%.

In recognising the importance of 'back', both Microsoft Internet Explorer and Netscape Navigator provide several shortcuts for issuing the command. Internet Explorer

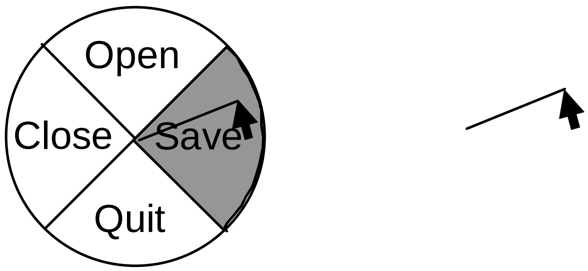

Figure 1: Pie-menu (left), and an equivalent selection using a marking menu (right).

supports two keyboard accelerators for back: Backspace and 'Alt+left-arrow'. Netscape Navigator also uses the 'Alt+left-arrow' key binding. The main limitation of key-bindings for web browsing is that the mouse is the main mechanism for accessing the links on the page (link selection accounts for approximately 50% of user actions (Tauscher & Greenberg, 1997)). Issuing commands through keyboard shortcuts therefore incurs an overhead in homing the hands between the mouse and keyboard.

Another shortcut for 'back' provided by both browsers is the context menu that can be popped up by pressing the right mouse button. Accessing the 'back' menu item incurs overheads in waiting for the menu to be posted and in the Fitts' Law (1954) limitations of pointing to the menu item.

Through our experience, studies, and subjects' comments, it appears that the back command is seldom issued through any interface mechanism other than the back button. For further analysis and discussion of the pros and cons of the current representation and behaviour of the back button, see Greenberg & Cockburn (1999).

2.2 Marking Menus

Marking menus are a specialisation of pie-menus. Pie menus (Callahan *et al.*, 1988) are pop-up menus that appear immediately under the user's cursor when the mouse button is pressed (see Figure 1, left). The user selects items by dragging the cursor into the appropriate segment of the pie. The motivation for pie-menus is to minimise Fitts' Law constraints on time-to-target—in theory, a movement of one-pixel is sufficient to reach any of the menu items, and further movements result in the target effectively becoming larger.

Marking menus (Kurtenbach & Buxton, 1991) extend the pie-menu concept by allowing users to select items before the menu appears (Figure 1, right). Expert users can select items with a rapid 'flick' in the appropriate direction. If the user hesitates in their gesture (a delay of more than approximately half a second) then the pie-menu is displayed to assist in learning the gestures.

In evaluating marking menus, Kurtenbach & Buxton (1994) found that the marking feature was heavily used once users learned the location of commands on the menu. Studies have also shown that performance with marking menus deteriorates as the number of items in the menu increases (Kurtenbach, Sellen & Buxton, 1993).

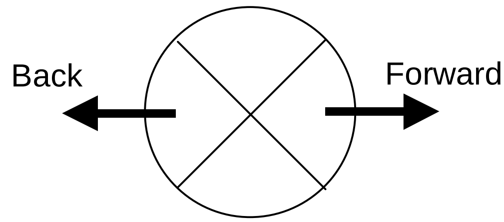

Figure 2: Gestures used in our marking system for web navigation.

Dulberg, Amant & Zettlemoyer (1999) compared simple 'flick' gestures with normal button clicks and keyboard shortcuts. In tasks involving flicking towards abstract targets, they showed the flick gesture to be 26% faster than button selection, but not reliably faster than key-bindings. Users also found the gestures easy to learn, with only 4 'errors' from 3300 trials—inaccurate flick directions were not considered to be errors. In their informal study (six subjects) of the flick gesture for redirecting keyboard focus to items on the Microsoft Windows desktop, subjects reported no problems with learning the technique, and five of the six participants said they would use it if available. Our evaluation furthers Dulberg *et al.*'s excellent study by rigorously analysing gesture controls in a realistic task.

Gesture controls for web navigation appeared in commercial browsers at approximately the same time as we started examining their efficiency. Opera 5.11[1], released in April 2001, provides facilities similar to those evaluated in this paper. The Mozilla Optimoz project[2] released gesture navigation late in 2001. Sensiva[3] is a commercial front-end to the Windows, Mac and Linux operating systems that provides around a dozen gesture-based shortcuts for commonly issued commands, including cut, copy, paste, and back. To our knowledge none of these systems have been formally evaluated.

3 Gesture Navigation Implementation

To support gesture navigation, we constructed a web-site where each page contained a Javascript program that interacted with the browser. The web-site www.cosc.canterbury.ac.nz/~andy/gestureSite/ demonstrates the Javascript implementation.

---

[1] www.opera.com.

[2] optimoz.mozdev.org.

[3] www.sensiva.com.

Using unaltered browsers, the script records information on the movement of the mouse whenever the mouse button is pressed. The script's execution is transparent to the user. When the script records a valid mark, it instructs the browser to execute the appropriate action (navigating back or forward). The pop-up pie-menu of traditional marking menus is not supported. Given the small size of the gestural input vocabulary (two commands) and the natural mapping between the gesture and the commands, it is unlikely that the gestures will be forgotten.

To issue a back or forward command, the user 'flicks' the mouse, with the left button held down, left (back) or right (forward), as shown in Figure 2. The requirements for a valid 'flick' are as follows.

Gesture size. To prevent the script from recognising a simple mouse click (such as a link selection) as a gesture command, the mouse coordinates must change by at least 35 pixels between the button being pressed and released. The selection of this value was guided by the median mark size of 48.4 pixels reported by Dulberg, *et al.* (1999).

We conducted a pilot study of 'naturalistic gestures' to determine appropriate values for gesture size and duration. Seven subjects, all computer science graduate students, were asked to generate fifty 'flick' gestures using the left mouse button in each of four directions (up, down, left and right). A pointer acceleration setting (control-display gain) of two-to-one was used in the study. The mean displacements on the x-axis for the left and right gestures were 137.5 (σ 23.7) and 149.1 (σ 32.7) pixels. Figure 3 shows the distribution of the length of flicks in the left and right directions, aggregated into ten-pixel boundaries. For the up and down gestures, the mean displacements on the y-axis were 136.1 (σ 47.4) and 123.7 (σ 51.5) pixels. Across all four directions, the mean and median gesture sizes were 136.6 and 139 pixels. Analysis of variance showed no significant difference between gesture distances in the four directions: $F(3,18)=1.09$, $p=0.38$.

The mean gesture size in our pilot study is substantially larger than the 48-pixel value reported by Dulberg, *et al.* (1999). This is most likely due to our use of two-to-one pointer acceleration, but Dulberg, *et al.* do not specify the level of pointer acceleration used in their study. With our display settings, we calculated that a mouse motion of approximately 7mm corresponded to the mean 136.6-pixel gesture.

Gesture duration. The left mouse button is frequently used to select and highlight text on web pages. To prevent the script from recognising these actions as gesture commands, gestures must be completed within 250ms.

The pilot study revealed mean gesture duration (time between the mouse button being depressed and released) in the left, right, up and down directions of 150.5 (σ 41.2), 159.9 (σ 48.9), 157.7 (σ 46.5) and 163.4 (σ 45.7) milliseconds, giving an overall mean of 157.9 (σ 43.3)

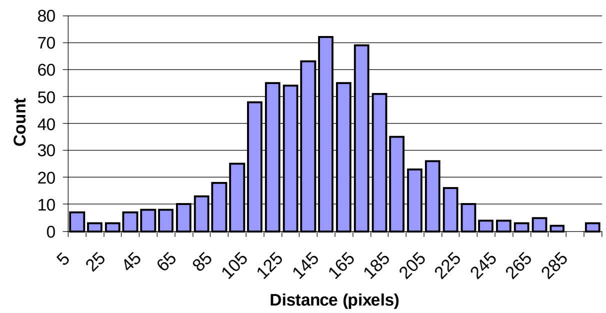

Figure 3: Distribution of 'flick' gesture distances, aggregated by ten pixels, measured on the x-coordinate for left and right gestures.

milliseconds. These means are not significantly different: $F(3,18)=1.5$, $p=0.24$.

Gesture direction. The marking menu is split into four regions (Figure 2). This allows the direction of the mark to be easily determined. The upper and lower regions are unused. To detect the back or forward gesture, the absolute mouse coordinate change between the button being pressed and released must be greater on the X-axis than on the Y-axis.

Page links. Gestures are not recognised when the mouse button is pressed while the cursor is over a page link. This property is almost certainly undesirable, but we were unable to control it within our implementation. We explained this feature to subjects as an unfixed usability bug.

3.1 Evaluation

The purpose of the evaluation was to determine the effectiveness of the gestural method for navigation in comparison to the normal back button. In particular, we wished to compare the efficiency of the two schemes, to observe whether users would be able to quickly learn to use the technique, and to measure the levels of subjective satisfaction with the two techniques.

All of the subjects participated in two web-browsing experiments. Both experiments were based around common types of backtracking behaviour in web-use: depth-first search and back, and breadth-first 'hub-and-spoke' (Catledge & Pitkow, 1995) browsing. All tasks in both experiments involved selecting a total of four page links and issuing the equivalent of four independent 'back' commands. In both experiments, a total of five different pages are displayed in the browser, and a total of nine page visits are made. In experiment one, four of the five different pages are displayed twice each. In experiment two, one page is visited four times, one twice, and the remaining three are visited once each.

The experimental procedure was consistent across all tasks in both experiments. First, the subjects were shown the

Start page
→ *follow link*
→ *follow link*
→ *follow link*
→ *follow link*
← back
← back
← back
← back
Start page

Figure 4: Depth-first navigation path used in experiment one.

Start page
→ *follow link to hub page (Figure 6)*
→ *follow link to spoke page*
←back to hub page
→ *follow link to spoke page*
← back to hub page
→ *follow link to spoke page*
← back to hub page
← back
Start page

Figure 5: Breadth-first 'hub-and-spoke' navigation path used in experiment two.

precise path through which they would navigate. They were required to rehearse the path at least twice (and more times if they wished). They were then asked to follow the path *as quickly as possible* using the normal back button. Having completed all of the tasks in experiments one and two using the back button, they repeated exactly the same paths for the two experiments using the gesture interface. We intended that the rehearsal of the navigation tasks prior to timing performance would minimise the impact of learning effects. Essentially, we were measuring expert performance of routine tasks.

3.2 Experiment one

Experiment one examined the effectiveness of the two interfaces in depth-first navigation. Figure 4 depicts the path, which involved following four links on subsequent pages then backtracking with four successive back commands. This task represents a directed search style of web-use: for instance, searching for a faculty member's web page starting from their university's home-page. The data from experiment one was analysed using a paired T-Test to compare task performance using the back button and the gesture system.

There are two reasons for hypothesising that the gesture system might not provide significant performance benefits in this experiment. First, the user need only make one movement to the back button, minimising the time-to-target overheads predicted by Fitts' Law. Having moved the cursor to the back button, the task can be completed by clicking four times with no additional cursor motion. Second, once the pointer is over the back button, the users can short-cut back to the top-level page using the back menu. Using the back menu adds one more cursor-positioning task in pointing to and selecting the desired menu item. It is unclear that the back-menu technique will be more efficient than issuing four discrete clicks of the back button without the overhead of cursor movement.

3.3 Experiment two

Experiment two examined the effectiveness of the two interfaces in breadth-first navigation, also called 'hub-and-spoke' navigation (Catledge & Pitkow, 1995). Hub-and-spoke navigation involves visiting a series of links (or 'spokes'), one at a time, off a central 'hub' page: for example, visiting the pages for several members of faculty, one at a time, by selecting a series of links on the 'Faculty' page.

Figure 5 depicts the navigation path used. It shows that, beginning at a 'start' page, the user follows a link to a main 'hub' page, then navigates to three 'spoke' links off that page, pressing 'back' to return to the hub each time. Finally, the user issues a back command to return to the start page.

When using the back button, the mouse-pointing requirements of this task are much higher than experiment one. After selecting each link to a spoke page the user must point to the back button, and then point to the next page link. Fitts' Law predicts that the back button technique will result in slower task performance than the gesture technique.

A second factor for 'amplitude' was introduced in this experiment to allow us to analyse the degradation of performance as the distance between the back button and the links on the hub-page increased. The links on the hub-page were vertically aligned immediately above one another at one of three corners of the web page: top-left, bottom-left, and bottom-right of the page (see Figure 6). The location of the links gives three levels of amplitude 'low', 'medium' and 'high' with mean distances between the back button and the group of links of 8.5cm, 14cm, and 20cm.

The performance data in experiment two was analysed using a two-factor analysis of variance with repeated measures. The factors were 'interface type' with two levels (back button and gesture system) and 'amplitude' with three levels (low, medium and high).

3.4 Subjective measures

Seven questions were presented to the subjects during the experiments. We wished to measure the subjects' satisfaction with the back button (both before and after using the gesture system) and with the efficiency and learnability of the gesture system. All questions were answered on a five-point Likert-scale from one (disagree) to five (agree). The questions are summarised in Table 1.

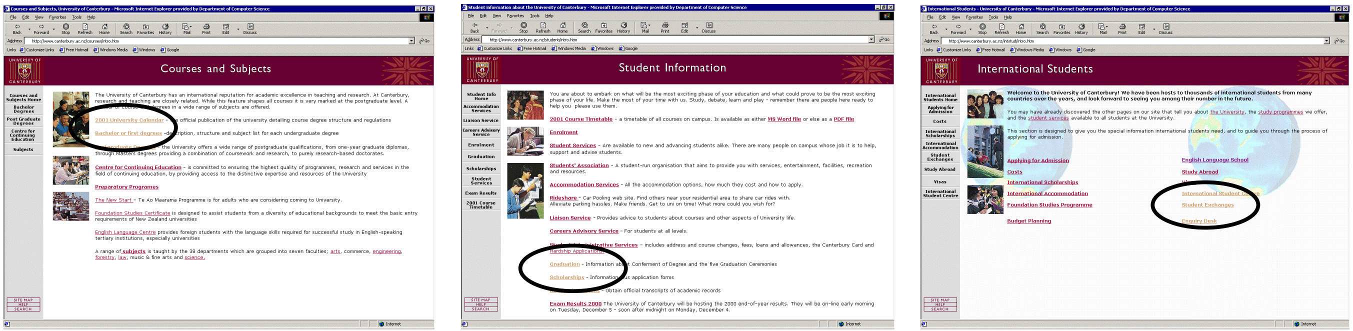

Figure 6: The three 'hub' pages used in experiment two. From left to right are the low, medium and high levels of amplitude, with the spoke links in the top-left, bottom-left and bottom-right of the page. The location of the spoke links is shown by the oval.

Question 1, "The back button is an effective means of navigation", was presented at the start of the experiment, before any tasks or training had been conducted. Question 2, "The back button allowed me to quickly navigate the pages", was presented after completing both experiments using the back button (before using the gesture system).

Subjects responded to Question 3, "The gesture system will allow me to navigate faster", after their first practice session with the gesture system.

Questions 4, "The gesture system did allow me to navigate faster" and 5/6, "The (back button/gesture system) is an effective means of page navigation" were presented after all tasks in both experiments were complete. Question 7, "The gesture system was easy to learn", completed the evaluation.

3.5 Subject Details and Training

The twenty subjects were all volunteer postgraduate Computer Science students familiar with web navigation. Each subject's training with the gesture system immediately followed solving experiment one using the back button. Training involved a brief (one or two minute) demonstration of the 'flick' gestures and explanation of the rules determining a valid mark. Each evaluation lasted approximately twenty minutes.

3.6 Apparatus

The experiment was conducted using Internet Explorer version 5.50 running on Windows 2000. The browser window was sized at 1152x864 pixels on a 17inch-monitor display running at a resolution of 1280x1024 pixels. The implementation of the gesture system was transparent to users, with many commenting on the seamless integration with the browser.

The Javascript implementation of the gesture system was embedded in every page viewed by the browser during the evaluation. The evaluation pages were placed on an IIS 5.0 web server, and the script was automatically added to every page served. All pages were held in the browser's cache prior to each experiment, ensuring that download speeds did not affect the results.

Task completion times were measured using a stopwatch.

The web pages used in the evaluation were selected from the University of Canterbury web-site. All pages had a consistent format, with a banner across the top, and an index menu on the left (see Figure 6). This format is familiar to the subjects who were all students at the university.

4 Results

As explained in the preceding section, to ensure that a learning effect did not confound the experiment, it was necessary for the subjects to perform their tasks in an expert manner, without instruction from the experimenter, and without hesitation in accessing each page. From our observation of the subjects' performance, we can confirm that the subjects' repeated rehearsal of the navigation paths was successful in encouraging expert performance. All but one of the subjects completed all of the tasks without hesitation, and mean task completion times were low. Across both interfaces and both experiments, the mean task completion time was 6.63 (σ 0.97) seconds. Considering that both experiments involved displaying a total of nine pages this task completion time may seem unrealistically low. Prior research, however, indicates that web browsing is a surprisingly rapid activity (Cockburn & McKenzie, 2001), with a high percentage of page visits lasting less than one second.

Although most subjects had little to no difficulty in learning to make valid gestures, one subject persisted in making slow and deliberate gestures, even when the importance of rapid flicks was stressed. His mean time for gesture tasks was three and a half standard deviations from the mean for the remaining nineteen subjects. All his data is removed from the evaluation. Almost all other subjects were extremely enthusiastic about the gesture system, and many expressed their desire that all browsers should support the technique.

4.1 Experiment one

Experiment one involved depth-first navigation through a series of links, followed by backtracking to the starting page.

The mean task completion times for the back button and gesture systems were 6.1 (σ 1.1) seconds and 5.4 (σ 0.96) seconds, showing a reduction of 11% in the mean task time when using the gesture system. This is a significant difference: two tailed T-Test, $t(18)=2.68$, $p<0.05$.

The reliability of the result (in favour of the gesture system) is surprising given that relatively little mouse-motion is necessary for the task, and that the back-menu can be used as a shortcut (one menu selection rather than four clicks of the button). Six of the subjects did use the back-menu as a shortcut for backtracking to the start page. Four of these were still faster using the gesture system. One subject incorrectly identified the title of the start page in the back menu, and had to issue one further back command.

Two subjects used the backspace key to issue the back command when using the normal browser. Both solved the task extremely rapidly (4.9 and 4.7 seconds). Interestingly, one of these subjects was faster still when using gestures (3.1 seconds).

4.2 Experiment two

The second experiment compared the efficiency of the two interfaces for breadth-first 'hub-and-spoke' browsing. Additionally, it examined the relative performance of the two interfaces as the distance (amplitude) between the back button and the links on the page increased.

The mean time for the back button was 7.44 (σ 1.12) seconds compared to a mean of 6.09 (σ 1.19) seconds for the gesture system, giving a significant main effect: $F(1, 18)=82.2$, $p<.001$. The gesture system reduced the mean task time by 18%.

The mean times for the three levels of amplitude were not significantly different ($F(2,36)=1.17$, $p=0.32$), with all means approximately 6.7 seconds. Considering that amplitude has no effect on task completion time in the gesture system, it is unsurprising that the means are not reliably different.

As expected, the interaction between factors 'interface type' and 'amplitude' is significant: $F(2,36)=6.78$, $p<0.01$. Figure 7 shows the cause of the interaction. As the amplitude increases, the mean completion times for the back button increase, but the task times for the gesture system remain relatively constant.

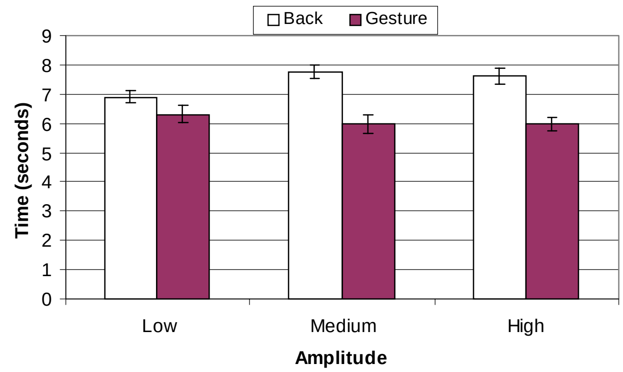

Figure 7: Mean task completion times in experiment two for the normal and gesture interfaces across the three levels of amplitude. Error bars show one standard error above and below the mean.

4.3 Comments and Subjective Measures

The subjects' comments go far beyond the quantitative results in demonstrating the effectiveness of the gesture system. Half of the subjects were extremely enthusiastic about the system, making comments such as "Fantastic", "That's amazing", "Really really nice" and "Brilliant!"

There were, however, some negative comments that frequently occurred. Six subjects complained that the implementation did not treat gestures as first class actions. This problem arose from two specific types of user action: first, when starting a gesture over a link on the page; and second, when starting a gesture over non-link text on the page. As mentioned in the implementation details section, gestures are not recognised when initiated over a link. Consequently, the subjects had to put some care into the location of their gestures over the browser window. This placement care is inconsistent with the informal and 'sloppy' flick gesture. Subjects felt that the gesture should work at any location over the page. If the gesture scheme was implemented in a commercial browser, it would be relatively straightforward to overcome this problem.

The second situation causing problems with gestures—initiating gestures over text—did not cause gesture recognition to fail, but did cause momentary concern for a few subjects. Normally, dragging the mouse over text causes text selection. In our implementation, mouse motion over text is overloaded: rapid brief mouse-down motion causes a gesture to be recognised, less rapid or longer mouse-down motion causes text selection. This overloading meant that when the user initiated a gesture over text, the text was highlighted as normal. The brief 'flash' of selected text prior to navigating back (or forward) was disconcerting to some of the subjects who complained of "having to find a blank bit of the page".

Overcoming this problem in a commercial implementation of the gesture scheme requires thought. One option is to use the right mouse button for gestures. The right button would then be overloaded, but this is already the case, as the contents of this menu depends on the object that the cursor

Table 1: Mean and standard deviation of responses to 5-point Likert scale questions.

Question	Mean	S.D.
Q1. The back button is an effective means of page navigation.	2.95	1.13
Q2. The back button allowed me to quickly navigate pages.	3.26	1.19
Q3. The gesture system will allow me to navigate faster.	4.11	0.81
Q4. The gesture system did allow me to navigate faster.	4.63	0.76
Q5. The back button is an effective means of page navigation.	2.42	1.12
Q6. The gesture system is an effective means of page navigation.	4.26	0.81
Q7. The gesture system was easy to learn.	4.26	0.93

is over when the button is pressed. Disambiguating gestures from menu requests could be based on coordinate change per unit time. We could not use the right-button technique in our Javascript implementation because Internet Explorer always pops up the menu on a right-button release, regardless of the coordinate change between button press and release.

The subjects' responses to the Likert-scale questions showed several interesting effects. Mean responses to the questions are summarised in Table 1.

Before beginning either experiment the subjects were asked to rate the effectiveness of the back button for navigation (Q1). The mean response was 2.95 (σ 1.13). After using the gesture system the subjects were again asked the same question (Q5). The responses showed a significant decrease to 2.42 (σ 1.12): Wilcoxon Signed Ranks test, $N=8$, $z=2.45$, $p<0.01$.

Comparing the subjects' responses to Question 3 and Question 4 shows that the subjects found the gesture system faster than they expected it to be after their initial training. Mean responses increased from 4.11 (σ 0.81) to 4.63 (σ 0.76) giving a reliable difference: Wilcoxon Signed Ranks test, $N=11$, $z=2.4$, $p<0.01$.

Eighteen of the twenty subjects rated the efficiency of the gesture system more highly than the back button (questions 5 and 6). The remaining two subjects gave the same rating for both interfaces. The mean responses for the effectiveness of the back button and gesture system were 2.42 (σ 1.12) and 4.26 (σ 0.81), giving a reliable difference: Wilcoxon Signed Ranks test, $N=17$, $z=3.46$, $p<0.01$.

Finally, subjects rated the gesture system highly for ease of learning (Q7): mean 4.26 (σ 0.93).

5 Discussion

To summarise the results, the gesture system significantly reduced the mean task completion times in both experiments. In the depth-first browsing task of experiment one, the gesture system reduced mean task times by approximately 11%, and in the hub-and-spoke browsing of experiment two the reduction was approximately 18%. Considering that backtracking activity such as that in experiments one and two is completed hundreds of millions of times every day, it appears that gesture navigation has the potential to enhance the overall efficiency of web navigation.

The potential of gesture navigation is further demonstrated by the minimal training that our subjects required and by the enthusiastic comments and subjective measures provided by the subjects.

5.1 Experimental concerns

There are several potential confounding factors that might have affected our results. Overall, we do not believe that removing or controlling any of these factors would alter the primary result, which reveals the enhanced efficiency offered by gesture navigation.

5.1.1 Implementation limitations

The Javascript implementation of the gesture system was not optimal—subjects had to ensure that their 'flicks' were not initiated over a link, and some subjects were disturbed by text-selections that occurred while generating gestures.

This factor is likely to have adversely affected subjects' performance with the gesture system. As discussed in the Results section, overcoming this limitation should be relatively straightforward in a commercial system.

5.1.2 Experimental tasks

The tasks in both experiments were based on expert navigation through rehearsed page sequences. This experimental design decision was necessary to remove highly variable performance factors such as searching for links, navigating to incorrect links, and so on.

Prior studies have shown that users' patterns of web use are highly repetitive and that a high percentage of pages are visited for surprisingly short periods of time (Catledge & Pitkow, 1995; Cockburn & McKenzie, 2001; Tauscher & Greenberg, 1997). Although the tasks in our experiments are not indicative of all web-browsing activities, we believe that deploying gesture features in commercial browsers would yield similar performance gains to those shown in our results.

5.1.3 Browser preference

All subjects used Microsoft Internet Explorer for the experiment. Approximately half of the subjects, however, used Netscape Navigator as their normal browser.

We do not believe that using Internet Explorer rather than Netscape Navigator affected the results because the location and behaviour of the back button is similar in both interfaces.

5.1.4 Subject pool

All of the subjects were post-graduate Computer Science students. Although this subject pool has more experience with mouse-use than most, we do not believe that their motor skills in generating the 'flick' gesture will be significantly different to other user groups.

5.1.5 Training

Although most users learnt the marking menu concept with ease, two subjects had problems during initial training. Rather than 'flicking' to make a gesture, they used a fairly precise but slow movement that took more than 250ms to complete. Once the importance of making rapid gestures was stressed, these subjects' performance dramatically increased.

In a commercial deployment of gesture navigation there may be problems with users learning the new features because they have no visual representation in the interface and because direct training is impractical. It seems likely however that awareness of the gesture features would propagate through informal discussions and observations of colleagues' web use.

5.1.6 Measurement tool

All tasks were timed using a stopwatch rather than software logging. There are obvious inaccuracies associated with stopwatch use. Initial trials of the experiment used server logs to measure navigation times, but these provided page access times at one-second granularity, which is a coarser measure than can be achieved using a stopwatch.

6 Conclusions

Marking menus provide a gestural means of issuing common commands in a quick and easy manner. This paper described the evaluation of a gesture system for navigating the web. The gesture system allows users to issue the frequently used back and forward commands with a simple 'flick' gesture.

Two experiments compared the time taken to traverse through a series of web pages using the standard back button and using the gesture system. The first experiment involved a depth-first traversal of pages, followed by backtracking to the starting page. The second experiment involved a breadth-first traversal, followed by returning to the starting page. Results show that the gesture system significantly reduced the time taken to complete these tasks, with a mean task time reduction of 18% for the breadth-first navigation task. Users' subjective ratings showed a strong preference for the gesture system, and their comments were extremely positive.

7 References

Callahan, J., Hopkins, D., Weiser, M., and Shneiderman, B. An Empirical Comparison of Pie versus Linear Menus. In *Proceedings of ACM CHI'88 Conference on Human Factors in Computing Systems*,. 1988. 95–100.

Catledge, L. and Pitkow, J. Characterizing Browsing Strategies in the World Wide Web. In *Computer Systems and ISDN Systems: Proceedings of the Third International World Wide Web Conference*. 10–14 April, Darmstadt, Germany, 1995, Vol. 27, 1065–1073.

Cockburn, A. and McKenzie, B. What Do Web Users Do? An Empirical Analysis of Web Use. *International Journal of Human-Computer Studies 54*, 6, 903–922, 2001.

Dulberg, M.S., Amant, R.S., and Zettlemoyer, L. An Imprecise Mouse Gesture for the Fast Activation of Controls. In *Proceedings of INTERACT'99*. 1999. 375–382.

Fitts, PM. The Information Capacity of the Human Motor System in Controlling the Amplitude of Movement. *Journal of Experimental Psychology 47*, 381–391, 1954.

Greenberg, S. and Cockburn, A. Getting Back to Back: Alternative Behaviors for a Web Browser's Back Button. In *Proceedings of the Fifth Conference on Human Factors on the Web*. 1999. http://zing.ncsl.nist.gov/hfweb.

Kurtenbach, G. and Buxton, W. Issues in Combining Marking and Direct Manipulation Techniques. In *Proceedings of ACM UIST'91*. 1991. 137–144.

Kurtenbach, G. and Buxton, W. User Learning and Performance with Marking Menus. In *Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systems*. 1994. 258–264.

Kurtenbach, G., Sellen, A. and Buxton, W. An Empirical Evaluation of Some Articulatory and Cognitive Aspects of Marking Menus. *Human-Computer Interaction 8*, 1, 1–23, 1993.

Nielsen, J., Usability Engineering, London: Academic Press, 1993.

Tauscher, L. and Greenberg, S. How people revisit web pages: Empirical findings and implications for the design of history systems. *International Journal of Human Computer Studies, Special Issue on World Wide Web Usability* 47, 1, 97–138, 1997.