

The Impact of Fluid Documents on Reading and Browsing: An Observational Study

Polle T. Zellweger, Susan Harkness Regli*, Jock D. Mackinlay, Bay-Wei Chang

Xerox PARC, 3333 Coyote Hill Road, Palo Alto, CA 94304
{zellweger, mackinlay, bchang}@parc.xerox.com, shr1@cmu.edu

ABSTRACT

Fluid Documents incorporate additional information into a page by adjusting typography using interactive animation. One application is to support hypertext browsing by providing glosses for link anchors. This paper describes an observational study of the impact of Fluid Documents on reading and browsing. The study involved six conditions that differ along several dimensions, including the degree of typographic adjustment and the distance glosses are placed from anchors. Six subjects read and answered questions about two hypertext corpora while being monitored by an eyetracker. The eyetracking data revealed no substantial differences in eye behavior between conditions. Gloss placement was significant: subjects required less time to use nearby glosses. Finally, the reaction to the conditions was highly varied, with several conditions receiving both a best and worst rating on the subjective questionnaires. These results suggest implications for the design of dynamic reading environments.

Keywords

Fluid user interfaces, fluid documents, focus+context, hypertext navigation, on-line reading, eye tracking, studies of dynamic user interfaces

INTRODUCTION

Typographic conventions such as footnotes and sidebars are often used to keep the main body of a document clear and succinct while still allowing the reader to access additional details. In an electronic document, hypertext can be used to provide more details than can fit typographically on a page. However, locating details elsewhere makes them more difficult to compare with the source document. Furthermore, hypertext requires users to navigate while reading, which is known to be cognitively difficult [5].

The Fluid Documents project has been exploring how electronic documents can provide more details on a page by fluidly adjusting typography on demand [4, 8, 20, 21]. We have developed a range of techniques that vary by how radically they adjust document typography. One focus of

these techniques is to support hypertext browsing by providing additional information (termed a *gloss*) at a link anchor, in order to help readers in choosing among links and understanding the structure of a hypertext [20].

While placing glosses close to their anchors seems most beneficial, we were concerned that the necessary radical adjustments of typography might be disruptive. In addition, we wanted to find out how different techniques affect how glosses are used, and how readers react to the availability of glosses and the techniques used to display them.

To shed light on these questions, we carried out an observational study exploring the impact of Fluid Documents on reading and hypertext browsing. We compared three fluid techniques with two established techniques for incorporating details into a document, as well as a conventional hypertext condition that displayed no glosses. Subject comments conveyed a wide range of reactions to the fluid techniques: some subjects loved it for reading and browsing; some hated it. However, a visualization of eyetracking data collected during the experiment showed no substantial differences in eye behavior between conditions, indicating that the fluid techniques did not create visual disruptions. Furthermore, we compared the average length of gloss events. In conditions with glosses near the source anchor, subjects kept glosses open on average for significantly shorter intervals than in conditions with distant glosses. Finally, we observed some subjects actively using the novel capabilities of Fluid Documents to freeze and thaw glosses while browsing.

RELATED WORK

Providing information about links that can hint at the content of the destination has been proposed in hypertext [11] and in the general field of information navigation [7]. The Hyperties system [12] used gloss-like information that appeared at the bottom of the page to help readers choose which links to follow. More recently, World Wide Web browsers have begun to provide facilities that can be used to display glosses: Microsoft Internet Explorer displays popup ToolTips to show a link's "title" attribute [15]. In addition, researchers have modified WWW browsers to show information about links and their destinations [10, 19].

In addition to providing information about link destinations,

* Current address: Baker Hall 259, Carnegie Mellon University, 5000 Forbes Ave. Pittsburgh, PA 15213

supplementary material can eliminate the need to follow a link altogether and give readers the opportunity to read additional information in its related context. Classic and recent research in reading comprehension processes has upheld the result that the context in which a passage is encountered can have a significant effect on a reader's comprehension (e.g., [3, 9]). A potential hypertext strategy, then, would be to provide a way for additional information to be brought into view in proximity to the primary material so that the elaboration can be read in context.

Fluid Documents bring more detail into view while preserving the surrounding information. This shares characteristics with fisheye views [6], zooming user interfaces [2], and other focus+context interfaces [16].

FLUID LINKS

We have implemented a hypertext browser, the Fluid Links browser, to experiment with techniques for fluidly displaying glosses in textual documents. The Fluid Links browser displays documents in pages of text, surrounded by margins on all four sides. In this study, each hypertext document fit fully on the page, so no scrolling was required.

Underlined text indicates the presence of a gloss; when a reader places the mouse over that text, the gloss smoothly and quickly expands nearby while the surrounding information alters its typography to create the needed visual space. The technique uses perceptually based animation to provide a lightweight and natural feeling to readers as a gloss appears and disappears.

Fluid techniques

The Fluid Links browser can be configured to display glosses in several ways. We briefly describe the techniques here; more detail can be found in [4, 20, 21].

The *fluid interline* technique displays the gloss directly below the anchor (see Figure 1a on separate color plate). When the mouse moves over the anchor, the gloss grows from an invisible, tiny size to its full size, while intensifying into its final color. At the same time, the primary text moves apart to make space below the anchor into which the gloss can expand, using the top and bottom margin space as well as compressing the interline spacing throughout the page.

The *fluid margin* technique displays the gloss in the margin (see Figure 1b). When the mouse moves over the anchor, a line expands from the anchor to the margin, and the gloss grows into the margin.

The *fluid overlay* technique grows the gloss directly below the anchor, as in the interline technique (see Figure 1c). However, the primary text does not move. Instead, any overlapped lines of text fade to a light gray color, allowing the overlaying gloss to be readable.

The fluid techniques display gloss text in red so as to more clearly set them off from the primary text. Furthermore, readers can freeze glosses in place to leave them open on the

page while other glosses are examined, and then remove or “thaw” the glosses when the information is no longer useful.

Conventional techniques

The Fluid Document techniques introduce many dynamic elements: text moves around on the page, grows and shrinks, changes in color, and can overlap. While these animated effects are often seen in television commercials, they have not been common in “serious” reading interfaces.

To compare this activity with established techniques for displaying additional text, we augmented the Fluid Links browser to display glosses in two other ways. The first of these is familiar from paper documents, and was employed in the Hyperties system: displaying the gloss material in a *footnote* area at the bottom of the page (see Figure 1e).

The second technique shows the gloss material in a *popup* window that appears at the cursor, as in Microsoft Internet Explorer's implementation of link titles (see Figure 1d).

Because footnotes and popups were intended to mirror existing implementations, they do not support freezing or animated growing and shrinking—they simply appear and disappear immediately as the cursor dwells on the anchor and leaves the anchor. Furthermore, their text is shown in black.

The final version of the browser does not show glosses at all, mimicking a conventional hypertext or WWW browser.

Dimensions of variation

The various gloss techniques that the Fluid Links browser can exhibit allowed us to examine three major dimensions of variation.

Placement of gloss. Fluid Documents attempt to place glosses close to the anchor to minimize the reader's eye movement and allow easy comparison to the primary text. Fluid interline, fluid overlay, and (non-fluid) popups all display the glosses directly below the anchor. To avoid disrupting the primary text, the fluid margin technique places glosses in the side margins, still relatively close to the anchor. However, a typical margin is narrow and requires a significantly different aspect ratio for the gloss (tall and narrow vs. short and wide). Finally, footnote placement in the bottom margin has the advantages of a predictable location, normal aspect ratio, and possible familiarity from footnotes on printed pages.

Modifications to the primary material. Placing glosses near the anchor usually requires some kind of modification to the primary material to make the gloss material readable. Fluid interline moves the primary text apart, and fluid overlay fades part of the primary text and covers it with the gloss. The popup technique completely obscures part of the primary text, but this disruption is likely to be familiar to most users of modern computer interfaces. The fluid margin technique interferes minimally, by drawing a line in between the text lines; otherwise it simply uses existing white space to display the gloss. Finally, the footnote technique does not interfere with the primary text at all, with glosses appearing in a dedicated region of the page.

Transitions. We intend for the animation in the three fluid techniques to guide the reader's eye to a gloss as it is displayed, and then back again to the primary text afterwards. We believe this smooth movement mitigates the concomitant changes in the primary text's typography. The transitions for the non-fluid techniques were not animated.

These dimensions allowed us to explore three questions:

1. Do large fluid-style changes in the page (text moving apart or fading in color, lines shooting to the margin) impact eye behavior?
2. Do different techniques affect how readers use glosses?
3. Given the technique variations (in the use of animation, in ways they alter the primary text, and in where the gloss is displayed), do readers have preferences?

To answer these questions, we created an experimental situation in which a large amount of usage data could be collected about a set of readers, using a within-subjects design in which each subject saw and interacted with each of the different conditions.

METHODS

Task

Subjects were asked to browse through hypertext documents in order to answer a series of 18 questions.

Subjects

Eight study subjects were recruited from a pool of research interns at Xerox PARC, ranging in age from 17 to 40; data from two subjects had to be discarded due to equipment problems (see Table 1). Subjects regularly read documents and browsed the web on computer screens. No subject had any prior knowledge of the Fluid Documents interface.

gender	5 male, 1 female
education	2 recent high school grad, 1 undergrad, 3 grad student
field	3 computer science, 1 chemistry, 1 architecture & product design, 1 undecided

Table 1. Selected characteristics of the final six subjects.

Materials

Two sets of linked documents and questions were constructed from materials used to conduct previous experiments: the Spoofer Corpus described how to identify someone sending forged email [14], and the Wine Corpus gave information about wine [13]. Each corpus consisted of 30-40 pages of interlinked material. No table of contents was provided. To make substantially the same corpus available to subjects even in the No Gloss condition, glosses were constructed in one of two forms. The text of *complete glosses* was repeated on the destination page for all conditions; these glosses functioned as annotations to the primary text and were identified with a distinctive end marker (||). *Leading glosses*, available only in gloss conditions, gave a short summary of the information that would be found on the destination page, along with a double arrow (>>) end marker.

To encourage subjects to read and engage fully with the material, answers to questions were located sometimes in the primary text and sometimes in the text of a complete gloss.

Conditions

Subjects interacted with the documents in six conditions: Fluid Interline, Fluid Margin, Fluid Overlay, Pop-up, Footnote, and No Gloss. These conditions employ the corresponding gloss techniques described earlier. The order of conditions among the six subjects was balanced such that subjects saw the conditions in different ordinal positions and no two adjacent condition pairs were repeated.

In gloss conditions, dwelling the mouse cursor over an anchor for 200 msec triggered the display of the gloss. (We chose a relatively short dwell as a starting point to investigate easy gloss access in a reading environment.) Once triggered, the animation to a readable gloss in fluid conditions took 300 msec (including all text movement, lines to margin, etc.).

Measures

Application logs. Logs of application activity recorded all user and application events to the nearest millisecond.

Eye logs. We used the ISCAN RK-426PC pupil/corneal reflection tracking system to track the eye movement of subjects as they read and searched for information on screen. Eye logs written by the tracking system captured the point of regard of the eye on the display at 60 Hz. Similar systems have been used in previous user interface studies [1].

Think-aloud protocol. Subjects were asked to speak aloud as they navigated the documents; specifically, each was asked to talk about "what you are looking for as you read a particular document, what links you are planning to explore and why, and what information you are trying to find if you explore a particular link." In a conventional think-aloud protocol for reading research, subjects would have been asked to speak continuously and read everything out loud; we varied the method by allowing silent reading, so as to minimize the disruption of reading patterns for the eyetracking data.

Video of screen, eye movement, and think-aloud protocol. The video signal to the subject's display was captured and composed with a video signal from the eyetracker that showed the real-time point of regard of the eye by a small white annular marker. The audio of the think-aloud protocol was included in this recording as well.

Subjective response. Subjects completed a questionnaire collecting subjective responses about effect of gloss placement, effect of gloss movement, usefulness of the glosses, and ratings for each condition. They could interact with each interface as they filled out the questionnaire. In addition, we conducted a semi-structured interview with each subject after the questionnaire was complete.

Procedure

The eyetracking equipment was calibrated at the start of the experiment session and re-calibrated two or three times throughout the experiment to adjust for expected drift.

The subject began by using the experimental browser to read a page of text that contained no link anchors, for baseline measurement of reading. The following steps were then repeated for each of the six conditions: First, the subject used the browser to find the answer to one question about the Spoofer Corpus. This pre-instruction phase of each condition was designed to capture the subject's untrained reaction to a condition's display changes. Next, subjects were instructed in the operation of the condition, including (if permitted) how to freeze and thaw glosses, and that multiple glosses could be frozen in view at once. Subjects were encouraged to try these actions during the instruction period. Finally, the subject used the browser to find the answers to two questions about the Wine Corpus. If the subject took longer than five minutes to answer any question, hints were provided at increasing frequency until the subject found the location of the answer. This was intended to ensure that no subject fell significantly behind others in their experience with either corpus.

After all conditions had been encountered, the subject answered a subjective questionnaire, and engaged in a semi-structured interview about subjective responses to the variations among interface behaviors. The entire experimental session lasted 2-3 hours.

ANALYSIS

Gloss event table

We were most interested in the instances that subjects caused a gloss to be displayed, which we call *gloss events* (see Figure 2). A gloss event begins the moment when the gloss display is triggered, after the mouse has dwelled sufficiently on the anchor. It extends through the time the gloss grows to full size (in the fluid conditions) and until either (a) the gloss begins to shrink back down or (b) a new page is visited.

Application logs were processed to construct a table of gloss events. 889 gloss events were recorded for the six subjects in all conditions.

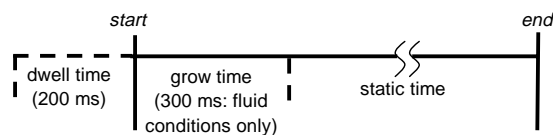


Figure 2. Gloss event.

Log simulation and visualization

Application and eye logs were temporally and spatially synchronized using a custom simulator matched with the videotape record. We also instrumented the simulator to generate a visualization, which we call a *gaze region*, to show the impact of a gloss event on the point of regard of the eye. In particular, we were concerned that fluid animations that affect large portions of the screen (e.g., moving all of the text apart in the fluid interline condition) might draw the subject's eyes to screen locations distant from the gloss or its related primary text. Even a momentary glance could indicate that

the animation was distracting the subject from the reading task. Because the gaze region visualization fills in the entire area that the eye has observed, even brief glances to distant parts of the screen are quite obvious.

The gaze region visualization uses a black line to trace the path of the point of regard during a gloss event. (Because we were interested in possibly rapid behavior of the eye, point-of-regard data was used rather than an analysis of eye fixation.) The temporal aspect of the path is described by dividing the first eight seconds of the gloss event into one-second intervals. Each interval is summarized by a convex hull (i.e., a rubber-banded line enclosing all points during that interval) that is overlapped and shaded from red to blue to show the temporal behavior of the eye in a static visualization (see Figure 3 on color plate). Most gloss events last less than eight seconds (see Table 2); later eye movement is shown only by the black line. Gaze region visualizations were generated for all 889 gloss events.

Pruned gloss event table

Initial examination of the raw gloss events revealed that there were many very short events (80 under 300 msec, 141 under 500 msec, 253 under 1 sec). Analysis of the raw application logs and the videotapes yielded the explanation that these short events were a result of the lightweight mouseover UI, in which the system interprets mouse dwell as an implicit request to view a gloss. One way such an "inadvertent" gloss event can occur is if the subject clicks on an anchor to follow the link, but happens to dwell on the anchor long enough to begin the gloss event. Subjects who track their reading location with the mouse can also inadvertently trigger short gloss events that they immediately cancel by moving the mouse away.

Another kind of inadvertent gloss event can occur if the subject is attending to one region of the screen and their mouse happens to dwell on an anchor in a different region of the screen. In this case, it is possible that the subject will not notice the gloss event.

Because we wanted to focus on how subjects used glosses they intentionally opened and (to some extent) read, we pruned the gloss event table of inadvertent gloss events. We carried this out in two steps. First, to eliminate "unnoticed" inadvertent gloss events, two experimenters independently coded all 889 gaze region visualizations to determine whether the eye went into the gloss at all. 156 gloss events in which the eye did not enter the gloss were pruned. (121 of these were less than 1 second long; the maximum of the remaining 35 was a Footnote gloss event lasting 5.318 seconds). Two such gloss events were retained because those glosses had been intentionally frozen and were thus not inadvertent. Second, to eliminate "short" inadvertent gloss events, we calculated a figure for each gloss event based on subject reading speed, gloss word count, gloss event static time, and an assumption that subjects need only skim a gloss to derive value from it. A cutoff time of 1 second could conservatively

account for these figures, so all gloss events whose length was shorter than 1 second were pruned, removing 133 additional gloss events.

One final adjustment was made to the gloss event data. Glosses in animated (i.e., fluid) conditions were not available for attention to anything other than their existence (i.e., no reading or skimming activity) for the first 300 msec after the gloss event, because they were moving and growing. We thus subtracted 300 msec from those raw recorded durations to yield a *static time*, which could then be compared on an equal footing with the unaltered durations in the non-animated conditions. Pruning adjusted gloss events under 1 second in length removed 26 more events, yielding a final count of 574 gloss events.

Transcription

The audio recording of the think-aloud protocols and the subjective discussions for three out of six subjects were transcribed in detail with the goal of understanding self-reported variations in preference and usage. The resulting transcripts were mined for comments that clustered around common issues, such as distraction, disturbing text (broken into obscuring and moving text), reactions to animations, use of screen space and spatial arranging, locality of gloss to anchor text, and the amount of visible information on a page. Twenty-eight clusters arose out of the data; some quotations were relevant to more than one cluster. Related statements from different subjects were then analyzed for areas of strong agreement or disagreement. On the basis of these relations we were able to describe several dimensions along which user preference varied.

OBSERVATIONS

Impact on eye behavior

After comparing the 889 gloss events with videotape, simulation, and visualization, we did not discern any substantial differences in eye behavior between conditions except those due to gloss location and shape. Furthermore, pre-instruction gloss events, which were the first ones seen by a subject, looked just like post-instruction events. Figure 3b (on color plate) is an example of an interline gloss where the eye moves to the gloss during the grow even though the rest of the text on the page is also moving.

Gloss usage

Analysis of gloss event length

To test the hypothesis that the condition would have a significant effect on the gloss event length, repeated measures ANOVA tests were conducted for both the full gloss event table and the pruned gloss event table (analyses were performed on a log₁₀ transformation of the data to satisfy the assumption of a normal distribution). In order to account for the unequal number of measures in each condition, the ANOVA was completed using 30 mean calculations as the

dependent variable of interest:¹ the mean was computed for each subject in each condition, as an estimate of how long a subject keeps a gloss open in a given condition. To determine whether a condition had a significant effect on the number of gloss events, a similar repeated measures ANOVA was conducted for the count of gloss events. This analysis revealed no main effects for count (total sample: $F(4,20)=1.435$, $p=.2590$; pruned sample: $F(4,20)=1.510$, $p=.2371$).

Main effects for condition were found in the full gloss event table, $F(4,20)=3.469$, $p=.0262$. The observed effect was even stronger in the pruned gloss event table, $F(4,20)=11.374$, $p=.0001$. The Fluid Margin and Footnote conditions share the quality that they are more distant from the anchor text than Fluid Overlay, Popup, and Fluid Interline, which are local to the anchor text. A contrast analysis [17] of the pruned gloss event table revealed a significant difference between the distant Fluid Margin and Footnote conditions versus the local Fluid Overlay, Popup Gloss, and Fluid Interline conditions, $F(1,20)=42.015$, $p=.0001$.

	Total Sample	Pruned Sample
Fluid Margin	.426 (.107)	.565 (.126)
Footnote	.214 (.115)	.496 (.072)
Fluid Overlay	.297 (.218)	.324 (.136)
Popup	.195 (.184)	.246 (.158)
Fluid Interline	.251 (.074)	.239 (.075)

Table 2. Means (S.D.s) for length of gloss opens (log₁₀ seconds) for total and pruned event tables (N=30).

These results suggest that subjects were able to read local glosses more quickly than distant ones. However, it should be noted that some subjects reported that text was harder to read in the Fluid Margin condition due to its narrow formatting (see Figure 3c). Moreover, subjects may be motivated to dismiss local glosses more quickly to remove them from their region of focus, while distant glosses may require less urgency since they are not in the way.

Attention to glosses

Analysis of the gaze region visualizations revealed that when a gloss opened on a page, the reader did not always look at the content of the gloss; footnote glosses in particular tended not to be noticed. Footnote gloss events were both most numerous and longest among the events that were pruned. Figure 3a illustrates an instance in which the eye never moved anywhere near the footnote gloss text during the gloss event, thus indicating that the gloss was apparently not noticed at all.

Subjective data supported the observation that footnotes tended not to be seen: responses to the questionnaire item “I

¹ We would like to acknowledge Dr. Joel Greenhouse for valuable assistance in determining how to analyze the gloss event tables.

found the placement of glosses (bottom of the page, side margin, between lines, etc.) in this condition to be hard to find or easy to find” indicated that footnote glosses were hardest to find.

In subjective discussion transcripts, subjects also reported that footnote glosses required extensive eye movement: “the one irritant is it requires a fairly drastic eye movement”; “you really have to jump your eyes down. I don’t care so much about that, but there’s something nice about having the information appear near the link”; and “I tend to be very focused at where I am, and this thing happening at the peripheral vision was not very fun.”

Value of glosses

Figure 4 (on color plate) shows the aggregated time that each subject spent answering all three questions in each question set in a given condition, with the total time subdivided into time when different numbers of glosses were open. No subject had more than five glosses open at one time. To facilitate comparison of conditions across question sets, conditions are shown in the same order in each group, while the order of subjects varies.

Glosses were strongly but not universally valued during hypertext browsing. Across all subjects, questions, and gloss conditions, subjects had at least one gloss open for 26.6% of the time; per-subject gloss usage ranged from a minimum of 22.4% to a maximum of 32.1%. Since all of the information to answer the questions could be accessed via hypertext jumps rather than by viewing glosses, this number suggests that subjects found them useful. Responses to the subjective questionnaires support this observation: subjects reported that having information in the glosses was helpful, clear, and helped them follow useful links. One subject stated that “after using the others, it (the No Gloss condition)’s like surfing blind.”

User preference

User preference varied widely along several dimensions. The only two areas of strong agreement supported the original dilemma which Fluid Documents were designed to address; namely, that it is desirable to keep the gloss local to the anchor text and that the gloss should not disturb the primary text. For some readers, however, “disturbing” the text in a negative way meant moving the primary text in any way; for others it meant occluding the primary text so they couldn’t read it. This reflected one of the most extreme observed variations: between readers who had a strong need for the primary text to remain completely static (and thus occluding the text would be preferable to moving it) and readers who genuinely appreciated the ability to vary the spatial arrangement of the primary text while seeing as much information as possible (thus moving the text would be preferable to occluding it).

Reactions to the animation of text in the Fluid interfaces varied from extreme dislike to extreme appreciation. In the words of one subject, the practice of animating the text was

“bad, because again I can’t read things while it’s animated so it’s a waste of time and a distraction” and “irritating”; for another, the animation was “aesthetic,” “useful,” and (discussing the interline condition) “Really really cool. . . it’s very intuitive for me. The whole thing of the text parting and new thing coming on is a great metaphor for deeper information.” Preference for animation was related overall to the desire for a static versus a variable page layout, the time it takes for the animation to occur, the control it exerts over their focus, and the aesthetic of the text shifting to accommodate more visible information.

Preference for speed of the gloss appearance was not uniform, even in terms of what one reader might have desired. For one reader popup was too fast: “I tend to move my mouse around a lot. And even though this has a slight, you know, pause before it brings it on, I find the popup message far too demanding of attention.” For another the fast speed of the popup appearance was good, but might be affected by task: “when I’m doing something like this [the experimental task], it’s nice for it to be fast because I’m using it all the time but if I’m actually reading and mousing I might want something . . . that’s a little slower, so that if I just go over a thing it doesn’t activate.”

Preference for degree of control varied: in general readers didn’t want the interface to be too automated and demanding of attention, but some appreciated automated assistance like the guidance the animation can provide to the eye as it moves to find the gloss text. Figure 3c presents a good example of the amount of control the animated interface can exert on the eye movement of a reader. When the gloss appeared, the brightest red gaze region shows that the eye was immediately drawn to the gloss text. This effect may be perceived as either positive or negative, depending on whether the reader likes the interface to exert control over the focus of the eye.

Preference for speed and control seemed to be related partially to the way a reader uses the mouse while reading (e.g., following the reading location, doodling elsewhere, etc.), which differed noticeably among our subjects. Some of these behaviors can interact poorly with the implicit interpretation of mouse dwell as a command, especially when the primary text begins to shift inadvertently as a result.

Freezing and thawing behavior

The ability to freeze glosses in view is a novel functionality offered by the fluid interfaces. To give subjects a potential motivation for exercising this new behavior during the short experiment, the materials were designed to require subjects to compare information in two different glosses to answer four of the questions. All subjects were guaranteed to see at least one of these questions in a fluid condition, where freezing the glosses to view their contents simultaneously could be accomplished. Two of the six subjects nonetheless never froze glosses, preferring to view the individual glosses in succession, often repeatedly, to compare their contents. One

other subject froze glosses only during active comparison (when the condition permitted) to answer a focused question.

One subject froze glosses even before any comparison questions were asked, but only opened more than one gloss at a time during active comparison. This subject began to open glosses in a frozen state, which could be done by right-clicking on the link anchor. This usage style removes much of the implicit grow/shrink behavior of the interfaces, which is based on mouse position and movement. It has the dual advantages of opening the gloss more quickly, without any dwell time, and tacking the gloss down so that it will not vanish prematurely even if the mouse is moved. It thus allows for more dynamic, less deliberate, mouse movement.

Two subjects adopted a freezing strategy that involved multiple open glosses without a comparison prompt. One subject developed a breadth-first approach to reading and hypertext browsing, choosing to simultaneously open and freeze many or all of the glosses on a page shortly after arrival, and then to examine those glosses and possibly their destinations as appropriate as a later step. Two other subjects offered variants of the breadth-first approach during the subjective discussion.

CONCLUSION

This study suggests that care must be taken when designing dynamic documents for reading and browsing. Subject preferences were complex and intense. Subjects were even sensitive to subtle distinctions such as whether a gloss occludes or disturbs a document (popup or overlay). We were gratified to see in the eyetracker data that dynamic adjustments to document typography did not cause the eye's point of regard to shift wildly across the page. We were also gratified to see that gloss placement, a central issue for the Fluid Documents project, had significant impact. However, we caution the reader not to interpret this result as biased for one group over the other. When glosses must be processed quickly, typographic adjustments should be used to put them close to their source anchor. Otherwise, placing glosses outside the primary text can reduce negative reactions to typographic adjustments. Interaction details about how glosses are invoked, frozen, and removed are also important in the design of dynamic documents. Lightweight interaction, in particular, may lead to inadvertent invocations of glosses.

Clearly, the increasing use of computer-based documents is changing how we read. Our subjects ignored footnotes, even though they are the conventional way to put details in paper documents. Perhaps, as one subject suggested, Web browsers are training us to ignore text popping up at the bottom of the screen. We fully expect that changes in reading styles will continue. For example, frozen glosses were used by our subjects to compare glosses and to search the hypertext in a breadth-first order. We have high expectations that future authors and readers will want to use dynamic documents as long as issues such as the ones raised by this observational study are addressed.

ACKNOWLEDGMENTS

Cliff McKnight and Andrew Dillon graciously supplied materials from previous hypertext studies [13, 18] for use in constructing the study materials. We'd like to thank David Fleet and his team for last minute video magic to recover data. The ISCAN eyetracker was made available through the Office of Naval Research contract N00014-96-C-0097 with Peter Pirolli and Stuart Card.

REFERENCES

1. A. Aaltonen, A. Hyrskykari, K. Raiha. 101 spots, or how do users read menus? *Proc CHI'98*, 132-139.
2. B. Bederson, J. Hollan. Pad++: A zooming graphical interface for exploring alternate interface physics. *Proc. UIST'94*, 17-26.
3. J. Bransford, M. Johnson. Considerations of some problems of comprehension. In W.G. Chase (ed.), *Visual information processing*. New York: Academic Press, 1973, 383-438.
4. B. Chang, J. Mackinlay, P. Zellweger, T. Igarashi. A negotiation architecture for fluid documents. *Proc. UIST'98*, 123-132.
5. J. Conklin. Hypertext: an introduction and survey. *IEEE Computer*, Sept 1987, 17-41.
6. G. Furnas. Generalized fisheye views. *Proc. CHI'86*, 16-23.
7. G. Furnas. Effective view navigation. *Proc. CHI'97*, 367-374.
8. T. Igarashi, J. Mackinlay, B. Chang, P. Zellweger. Fluid visualization of spreadsheet structures. *Proc. Visual Languages'98*.
9. W. Kintsch. *Comprehension: A paradigm for cognition*. Cambridge: Cambridge University Press, 1998.
10. T. Kopetzky, M. Mühlhäuser. Visual preview for link traversal on the World Wide Web. *Proc WWW8*, 1999.
11. G. Landow. Relationally encoded links and the rhetoric of hypertext. *Proc. Hypertext'87*, 331-338.
12. G. Marchionini, B. Shneiderman. Finding facts vs. browsing knowledge in hypertext systems. *IEEE Computer*, Jan 1988, 70-80.
13. C. McKnight, A. Dillon, J. Richardson. A comparison of linear and hypertext formats in information retrieval. In R. McAlesse, C. Green (eds.), *Hypertext: State of the art*. Oxford: Intellect Books, Ltd., 1990, 10-19.
14. C. Neuwirth, J. Morris, S. Regli, R. Chandhok, G. Wenger. Envisioning communication: Task-tailorable representations of communication in asynchronous work. *Proc. CSCW'98*, 265-274.
15. J. Nielsen. Jakob Nielsen's Alertbox for January 11, 1998. <http://www.useit.com/alertbox/980111.html>
16. G. Robertson, S. Card, J. Mackinlay. Information visualization using 3D interactive animation. *CACM*, 36 (4), 1993, 57-71.
17. R. Rosenthal, R.L. Rosnow. *Contrast Analysis: Focused comparisons in the analysis of variance*. Cambridge: Cambridge University Press, 1985.
18. J. Rouet, J. Levonen, A. Dillon, R. Spiro, eds. *Hypertext and Cognition*. Lawrence Erlbaum Associates, 1996.
19. D. Stanyer, R. Procter. Improving Web usability with the link lens. *Proc. WWW8*, 1999.
20. P. Zellweger, B. Chang, J. Mackinlay. Fluid links for informed and incremental link transitions. *Proc. Hypertext'98*, 50-57.
21. P. Zellweger, B. Chang, J. Mackinlay. Fluid links for informed and incremental hypertext browsing. *CHI'99 Extended Abstracts (Video Program, 6:38 minutes)*, 7-8.

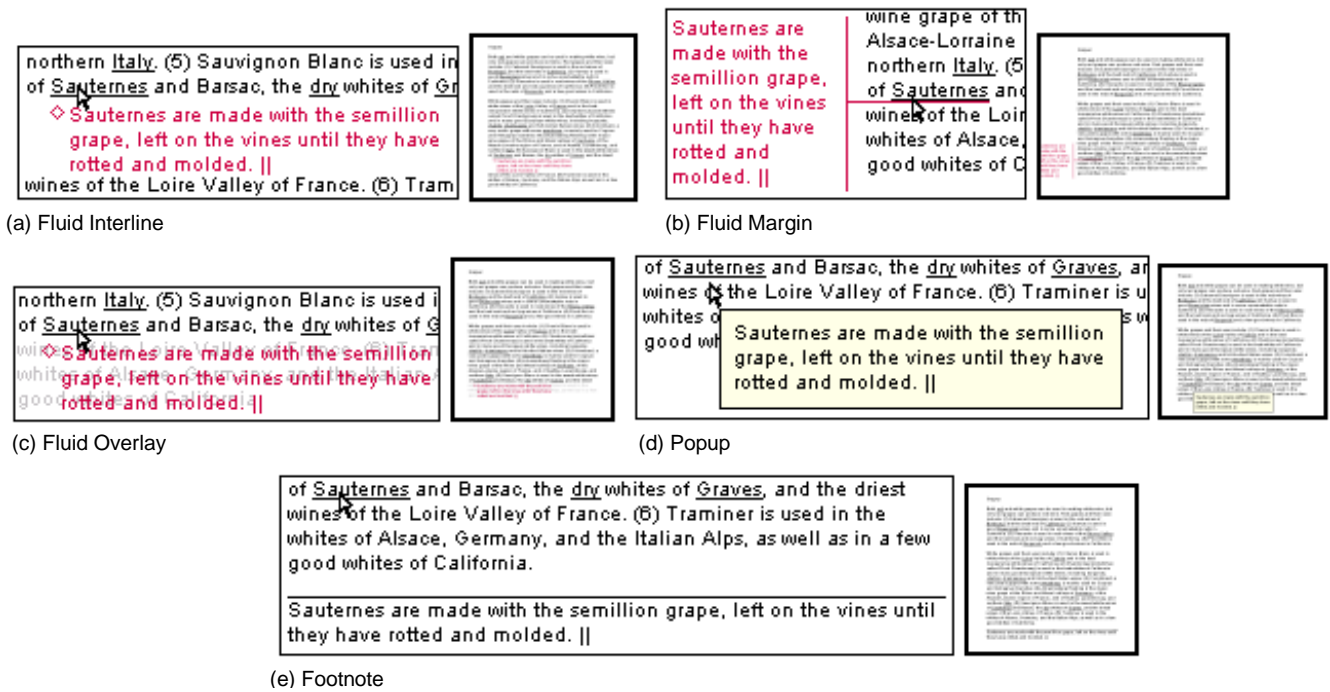


Figure 1. Fluid and conventional glosses, along with thumbnails of the pages of text from which they are excerpted. Font sizes have been greatly reduced for this figure; in the study, the primary text size was chosen to have a visual size of 12 points.

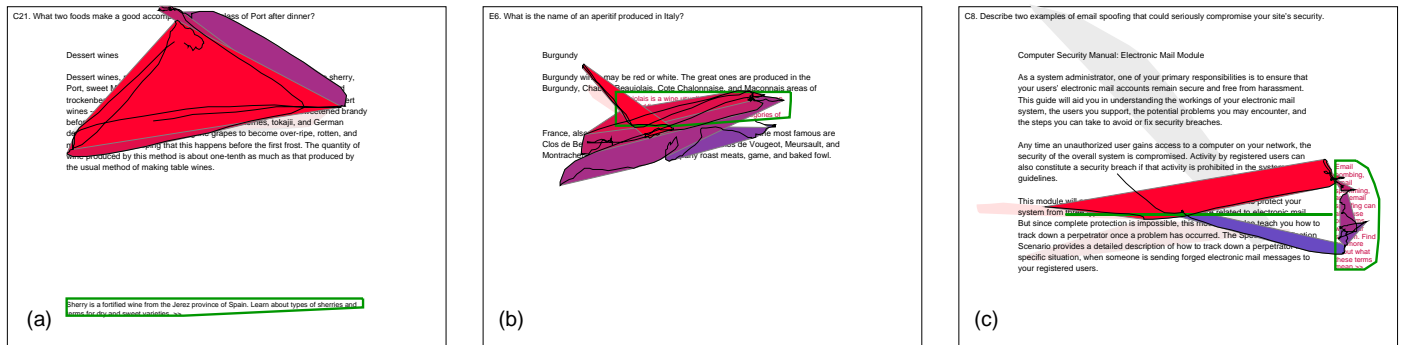


Figure 3. Gaze region visualizations of gloss events. The bold green line outlines the gloss, while the thin black line traces the eye's point of regard within the gaze region. The gaze region during the first second of each event is shown in bright red and may occlude later regions; gaze regions for successive seconds are shown in darker colors. The pale colors behind the text show gaze regions for three seconds before the event. (a) shows a Footnote event that the subject does not see. (b) shows a Fluid Interline event that the subject sees and reads. (c) shows a Fluid Margin event that the subject reads with seeming difficulty (note twists in black line).

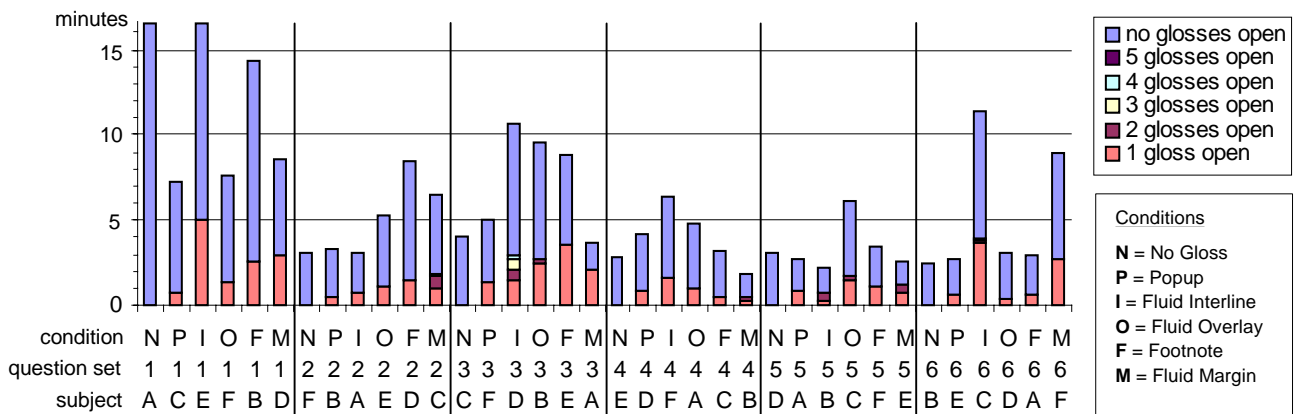


Figure 4. Gloss usage (pruned static time) per subject in each question set.