Submitted to ACM *Transactions on Computer-Human Interaction,* September 20, 2004. Please do not cite or distribute.

High-Cost Banner Blindness: Ads Increase Perceived Workload, Hinder Visual Search, and Are Forgotten

MOIRA BURKE and ANTHONY HORNOF
University of Oregon

and

ERIK NILSEN and NICHOLAS GORMAN
Lewis & Clark College

The seeming contradiction between "banner blindness" and a general distaste for Web advertisements motivates a pair of experiments into the effect of banner ads on visual search. Experiment 1 measures perceived cognitive workload and search times for short words with two banners on the screen. Four kinds of banners were examined: (1) animated commercial, (2) static commercial, (3) cyan with flashing text, and (4) blank. Using NASA's Task Load Index, participants report increased workload under flashing text banners. Experiment 2 investigates search through news headlines at two levels of difficulty: exact matches and matches requiring semantic interpretation. Three kinds of banners were used: (1) animated commercial, (2) static commercial, and (3) gray rectangles. Results show both animated and static commercial banners decrease visual search speeds. Eye tracking data reveal people rarely look directly at banners. A *post hoc* memory test confirms low banner recall and, surprisingly, that animated banners are more difficult to remember than static look-alikes. Results have implications for cognitive modeling and Web design.

Categories and Subject Descriptors: H5.2 [**Information Interfaces and Presentation**]: User Interfaces – *Graphic user interfaces (GUI)*; Screen design; H1.2 [**Information Interfaces and Presentation**]: User/Machine Systems – *Human information processing*

General Terms: Design, Experimentation, Human Factors

Additional Key Words and Phrases: Animation, visual search, banner ads, eye tracking, memory, World Wide Web

---

1. INTRODUCTION

Web designers routinely animate advertisements in an attempt to make them more conspicuous. Yet few empirical studies explore the effect of animation on a concurrent visual task, and few if any examine it in an ecologically realistic context. Early research suggested that Web users are functionally blind to rectangular graphics that they perceive to be advertisements [Benway and Lane 1998], but more recent studies indicate that people do notice ads, dislike them, and that site credibility suffers [Fogg et al. 2001]. Ad-blocking software, such as the products summarized by Rowe et al. [2001], are popular. However, the HCI community lacks data conclusively demonstrating whether animated ads impede common visual Web tasks.

Animation ostensibly aids memory, but little evidence supports this theory when applied to Web advertising. Advertisers want to leave a lasting impression: favorable

---

This work was supported by National Science Foundation Grant IIS-0308244 to the University of Oregon with Anthony Hornof as the principle investigator.
Authors' addresses: Moira Burke, Anthony Hornof, Computer and Information Science Department, University of Oregon, Eugene, OR 97403; {moira, hornof} @cs.uoregon.edu; Erik Nilsen, Nicholas Gorman, Psychology Department, Lewis & Clark College, Portland, OR, 97219; nilsen@lclark.edu, ngorman@usc.edu

brand recall and viewer attitude matter as much as other effectiveness metrics (i.e., click-through) [Interactive Advertising Bureau 2001]. How animation affects subjective user experience and whether it aids memory have yet to be thoroughly investigated. Is animation beneficial enough to advertisers to outweigh its negative reputation? We seek to answer these questions through two experiments with animated banners: one that measures search time and participants' impressions of task workload, and another that tests ad recall. The research presented here confirms that animated banner ads interfere with common Web tasks and yet are no more memorable than static ads.

1.1 Animation and Attention Capture

Humans assign attentional priority via two biases: *stimulus-driven* (bottom-up) triggers that distract regardless of the task, and *goal-directed* (or top-down) strategies [Yantis 2000]. Separating the two has proven challenging even under highly-controlled experimental design [Bacon and Egeth 1994]. Web tasks are even more complicated because searchers' goals are not always well-defined, and ads with tempting words like "FREE" may attract attention through semantic appeal. However, traditional studies of top-down and bottom-up attention capture provide clues to how animation affects Web users. There are two prevailing schools of thought: the first contends that certain forms of animation always attract attention from the bottom up, whereas the second argues that people unintentionally create task-completion (top-down) strategies that make irrelevant stimuli relevant, and thus distracting.

Motion attracts attention, though researchers disagree on the extent to which this occurs automatically. Yantis and colleagues [Yantis and Jonides 1990, Jonides and Yantis 1988, Hillstrom and Yantis 1994] found that abruptly-appearing stimuli capture attention in a purely bottom-up manner. However, people do not involuntarily look at other forms of animation, such as oscillation and simple shape changes. Motion per se does not attract attention, but rather abrupt appearances create new perceptual objects and these new objects tend to attract attention. Mere luminance changes were ruled out as a confound [Enns et al. 2001]. Even when the abruptly-appearing items were known to be irrelevant to the task, people still had trouble ignoring them [Remington et al. 1992]. However, Franconeri and Simons [2003] refuted the "new perceptual object" theory, proposing instead that attention is delegated to events of behavioral urgency. They maintained that with more salient stimuli, other kinds of motion will capture attention as strongly as abrupt-onset. Looming and disocclusion (moving from behind another object) were highly distracting, while receding and unique-coloring were not. Combined, these experiments suggest that some kinds of motion will inherently attract attention regardless of the viewer's intent.

The second prevailing theory is that no stimulus truly captures attention automatically; instead, people adopt task-completion strategies that make them susceptible to certain stimulus properties [Folk et al. 1992, Bacon and Egeth 1994, Yantis and Egeth 1999]. For example, observers searching for a particular singleton (an object unique from its neighbors in one dimension, such as shape) will be delayed by an irrelevant singleton in another dimension (color, orientation) [Pashler 2001, Theeuwes 1991]. This effect was originally attributed to bottom-up capture. Bacon and Egeth [1994] instead concluded that, though the additional singleton was not informative of target location, searchers were in *singleton detection mode*, deliberately attending to anything that "popped out," because it required less effort than consciously filtering for the singleton in the relevant dimension. When the same target was no longer a singleton (i.e., there were multiple instances of the target object), the irrelevant singleton no longer distracted. More broadly, in the *contingent involuntary orienting hypothesis,* Folk et al. [1992] asserted that distractors sharing task-critical properties with the target (such as singleton status) will have an effect, while other highly-salient but task-irrelevant distractors (such as abrupt onsets) will not. Pashler [2001] supported the theory, with the surprising result that participants searched faster in the presence of irrelevant (flashing, twinkling, and shimmying) distractors. The attractiveness of certain stimuli may depend on the viewer's mindset rather than any intrinsic power of animation.

Though the debate over the relative influence of goal-directed and stimulus-driven attractors continues, there has been little effort to explore these issues in common Web tasks. Does a car suddenly appearing within a banner advertisement create a new perceptual object? If underlined blue text on a Web page is task-relevant to someone searching for a link, what if the blue text is within a banner ad? Would a lone ad be considered a singleton on a typical, cluttered Web page?

A Web-like environment has been used in at least two visual search studies. In Zhang [2000], participants searched and counted random text strings on a Web page with an animated distractor irrelevant to the task. Distractors included images (i.e., a blinking eye or a waving robot) and big letter strings that alternately loomed and receded (a motion shown to capture attention in some situations [Franconeri and Simons 2003] but not in others [Hillstrom and Yantis 1994]). Zhang found that both kinds of distractors slowed the primary search task, but the degree depended on the difficulty of the task. Participants were worse at counting both short and long strings in the presence of animation, but short-string tasks were more adversely affected by animation than long-string tasks. The results are difficult to interpret in part because speed and accuracy were not reported separately, but instead combined into a single metric of "performance," in which fast but inaccurate counts were potentially scored the same as slow but accurate ones. Static versions of the distractors were not tested, so the effect could be due to the presence of large graphic singletons rather than animation.

In Diaper and Waelend [2000], participants answered questions based on blocks of text adjacent to animated graphics. Two levels of text length (short, long) and three versions of graphics (none, static, animated) were tested. Participants also rated the complexity of the six conditions "at a glance." Animation did not have an effect on either search time or perceived complexity. Search times were greater for the longer text blocks¸ as commonly observed. The study concluded that the amount of text on a page contributes to task difficulty far more than animation does, but provided no statistically significant support for this conclusion. The results are difficult to interpret: Participants were not given incentive to search quickly and complexity was rated on an unmarked visual scale. The experiments of Zhang [2000] and Diaper and Waelend [2000] bridge traditional attention capture and Web research, but methodological issues hinder their usefulness.

1.2 Animation and Memory

Some studies have explored the memorability of banner ads. Whether or not animated ads capture attention in a strict sense, they may imprint some features strongly enough to achieve subsequent recognition. Bayles [2002] addressed this issue by posting static and animated versions of two novel banners on a modified Library of Congress Web page. After four information-gathering tasks, participants were presented with a surprise recall test: to draw the layout and contents of the page from memory. They were also given a page of twelve ads and asked to select the ones they had seen. No correlation was found between animation and recall, and more than half the participants did not remember the presence of the banners at all. Animation also did not affect recognition. One detail in the design of the experiment is that the ten distractor banners in the recognition task included several that were very similar to the two banners shown in the information-gathering task. It is not clear whether people failed to recognize the banners altogether or just specific design details. Furthermore, only two banners were used in the experiment. Without a larger pool of banner designs, Bayles's results cannot be easily generalized.

Pagendarm and Schaumburg [2001] suggest that banners are more memorable to casual browsers than goal-driven searchers. A group of "aimless browsing" participants explored a 55-page Web site "according to their own interest," and a group of "goal-directed" participants navigated the site to answer a list of questions. Both groups were exposed to 16 animated ads. Aimless browsers recalled motifs and products in 3 to 17 times as many ads as did goal-directed searchers, but the performance was still very low, with details recalled from an average of just one or two ads. Participants also rated their confidence recognizing the banner ads when presented with them again; confidence was higher for aimless browsers. A more direct recognition test including distinct banners that participants had not seen was not administered.

The two previous studies examined banner recognition in a secondary task immediately following a primary browsing task but, because of details in the experimental designs, in both cases it is difficult to conclude whether people retained knowledge of the ads that appeared during the primary task.

2. EXPERIMENTS

The present experiments test whether standard animated banner ads affect Web users' (a) visual search speed, (b) perceived workload, and (c) memory.

As Pagendarm and Schaumburg [2001] propose that "banner blindness" may occur especially in a goal-directed task, this is the context in which we focused our examination of the phenomenon. Zhang [2000] observed that simpler tasks were more adversely affected by animation, and so we assessed three levels of task complexity. We also sought to extend Bayles's [2002] analysis of banner memory; in that study, only 2 different banners were shown during the primary task, so we increased the pool to 100 banners. Lastly, unlike the Diaper and Waelend [2000] study, when we failed to find a significant effect for search time in our first experiment, we looked further, exploring other search tasks and distractor types.

In Experiment 1, participants searched for short words while two banners appeared within the search area. Banners included (a) animated commercial, (b) static versions of the commercial banners, (c) novel cyan banners that flashed big text, and (d) blank (invisible) banners. After the timed trials, participants ranked their impressions of workload for each type of banner. In accordance with the contingent involuntary orienting hypothesis [Folk et al. 1992], we expected search times to be greatest for the "big text" banners because their text was similar to the search target. We also predicted that the animated versions of commercial banners would increase search times more than their static counterparts. Though irrelevant to the task, the animated banners contained graphics that appeared abruptly or grew in size, dynamic events that have been shown to capture attention [Hillstrom and Yantis 1994, Franconeri and Simons 2003]. Finally, we expected participants to report greater workload under the big text and animated banner conditions: if these banners do capture attention, they should require more effort to ignore.

Experiment 2 extended the first study with eye tracking and a more ecologically valid task: searching for news headlines. Participants performed two kinds of searches: *exact*, in which the target headline text was known, and *semantic*, in which the first few sentences of a full story appeared and the best-matching headline had to be found. A *post hoc* recognition test determined which banners participants remembered having seen. Consistent with previous findings that animation affects simpler tasks more than harder ones [Zhang 2000, Diaper and Waelend 2000], we expected animated banners would prolong both the exact and semantic search tasks, but to different degrees. We also predicted banner memory would be low, in accordance with Bayles [2002] and Pagendarm and Schaumburg [2001]. The eye tracking would reveal under what conditions participants looked at the banners.

3. EXPERIMENT 1

The primary aim of Experiment 1 was to determine whether a simple visual search task, finding a single word, would be affected by animated ads. Confounds such as those introduced by reading were removed from the experimental design.

3.1 Method

Participants. Twelve adults (six female) with a mean age of 27 participated in the experiment for compensation. All were experienced with graphical user interfaces and had normal or corrected-to-normal vision.

Materials. Each experimental trial presented one target object amid 19 distractors. Targets and distractors were capitalized four-letter words in 18-point Helvetica enclosed in rectangles with 1-point borders. Roughly 700 words were used. They were collected from the first two levels of Wired.com, filtered for profanity, and limited to one or two syllables to facilitate vocal repetition during the search.

Two banners appeared among the targets and distractors. Figure 1 presents the four banner types tested: (a) blank, (b) animated commercial, (c) static commercial, and (d) flashing text. One hundred ten animated banners were selected from reputable commercial Web sites including the New York Times, AOL, and Alta Vista. Static banners consisted of a representative frame from each animated banner. Flashing text banners were introduced as an extreme variety of animation: large black text alternately appeared on the left and right halves of a cyan rectangle every 150 msec. Flashing text banners cycled asynchronously, offset from each other by 80 msec. Text for these banners was randomly generated from the target and distractor words for a given trial. Figure 2 shows the screen layout. The target and distractors were arranged in three columns of eight rows. Banners spanned two columns each, removing four distractors per trial. Banners emerged in random rows, analogous to the unpredictable placement of ads on the Web, and targets appeared in all 24 positions across the experiment.

After finishing the timed search trials, participants completed Task Load Index (TLX) surveys. Developed by NASA [Hart and Staveland 1988], TLX measures workload, defined by the following factors: mental, physical, and temporal demand; effort; frustration; and impression of performance. For each banner type, participants rated these factors from 1 to 100 and then indicated for all possible pairs which factor contributed more to the overall workload. A combination of these values would reveal the relative importance of each factor, providing a metric with which participants' subjective responses could be compared.

Experimental software was written with Macromedia Director, a platform chosen for its facility with both animated and static graphics, and for its similarity to Macromedia Flash, a popular authoring tool of modern ad developers.

Design. Trials consisted of two stages: precue and search. During the precue stage, four randomly-ordered words appeared, one of which was the target. After studying the four words, the participant would click a box, hiding the precue and initiating the search stage. Pilot tests indicated single-word precues yielded very fast searches, so we used four words; this is somewhat analogous to how multiple terms (i.e.,

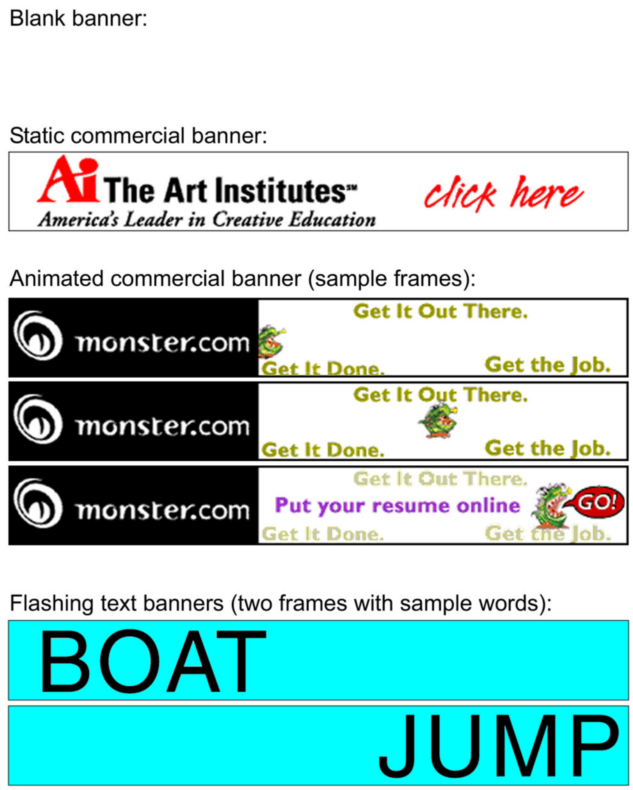

Fig. 1. The four banner types tested.

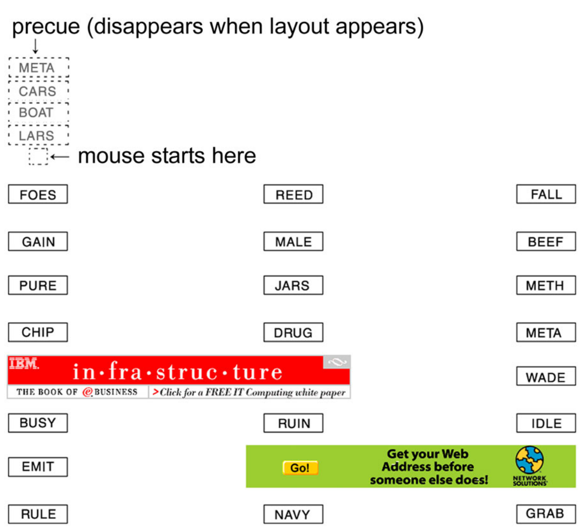

Fig. 2. Layout of target and distractors in Experiment 1.

"careers" and "employment") might satisfy a Web searcher's goal ("find a job"). The participant would look for a word matching one of the precue items and select it.

Search and selection times were separated using the Point Completion Deadline (PCD) [Hornof 2001]. In short, participants were instructed not to use the mouse until they visually located the target, at which point they should quickly click on it. Upon moving the mouse more than five pixels from its original location, participants had a limited amount of time, scaled by Fitts' law, to reach the target. If they failed to select an object in time, a buzzer sounded and the trial was recorded as an error.

Fast and generally accurate responses were motivated with incentives and penalties. By performing well, participants could increase their baseline pay of $10 to $15. Each trial had a potential bonus of seven cents with one cent deducted per second. Clicking the target earned the bonus and a 150 msec chime, but clicking anything else or exceeding the PCD warranted a 350 msec buzzer and a five-cent penalty. PCD expirations also triggered an alert box with a reminder not to move the mouse early. Cumulative earnings were displayed at the end of each block.

Each participant completed 96 trials, one for each unique combination of target location (24) and banner type (4). The trials were divided into four blocks and randomized, with blocks counterbalanced across participants through a Latin square. Banner-target combinations in error trials were repeated, shuffled into the remainder of the block. Thus, participants correctly completed all trials for all combinations. Two percent of the combinations were omitted due to subtleties in the experimental software. To gather baseline pointing times, participants completed an additional target-only block, in which the target appeared in each of the 24 positions three times, and no distractors were present; the target-only block occurred after the first two blocks for all participants, and the results will not be discussed here.

Procedure. After reading instructions about the precue, target, bonuses, and PCD, participants positioned themselves 56 cm from the screen with the precue at eye level. Eye-to-screen distance was reestablished before each block. Participants were allowed an unlimited number of practice trials from the first block type to become accustomed to the PCD. When they were ready, the software was reset and the data collection began. Five additional practice trials initiated every block.

For each trial, participants studied the precue words as long as necessary. Then they clicked the box to dismiss the precue and display the layout. After visually locating the target, they selected it quickly; its colors would invert briefly and either the chime or buzzer would sound.

After the timed trials, participants reported their experiences with the four banner types through a TLX survey. They completed a short "reminder" block with a single banner type and then provided TLX weights for the workload factors of that condition, repeating the process for each banner type. Blocks were again counterbalanced by a randomized Latin square.

After completing the TLX evaluations, participants were interviewed and asked to describe their overall impressions and search strategies.

3.2 Results

Search Times and Error Rates. Table I shows the mean search time for each banner type. Error and practice trials are excluded from analysis. Results from a repeated measures ANOVA fail to show a significant difference in search time. A position effect was observed as expected: participants found targets in the upper-left positions much faster than those to the lower-right, $F(23,72) = 5.08$, $p < .0001$. One significant effect was that when a target was sandwiched between two flashing text

Table I. Mean Search Times for Each Banner Type in Experiment 1

Banner Type	Search time (ms)	Standard deviation (ms)
Blank	5,831	1,675
Flashing Text	5,234	1,116
Animated Commercial	4,795	1,010
Static Commercial	5,155	1,238

Note. Times are averaged across target positions and participants.

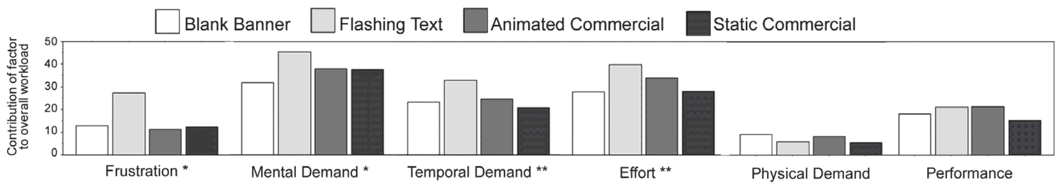

Fig. 3. Workload factors per banner type. Each factor is measured in relative units where a greater height indicates greater workload for all factors except performance, in which the inverse is true.

Mental demand and effort were the greatest contributors for all banner types.
* significantly higher for flashing text banner (p<0.05)
** marginally higher for the flashing text banner (p< 0.1)

banners it took an average of 75% (4.3 seconds) longer to find than if the banners were located elsewhere, $F(1, 298) = 7.0$, $p < .01$. This effect was not seen with the other banner types, nor did it hold if the target was adjacent to just one of the two banners.

Participants committed two kinds of errors: *misses*, in which they clicked something other than the target, and *timeouts*, in which the point completion deadline expired. The percentage of misses and timeouts were not significantly different across banner types.

Overall Workload. Participants reported tasks with flashing text banners to have the greatest workload. A repeated measures ANOVA revealed a significant difference between banner types, $F(3, 36) = 6.52$, $p < .001$, and a Fisher's PLSD *post hoc* test showed the flashing text banners were more workload-intensive than the others (*Mean* = 160.71). Animated banners also had greater overall workload, though non-significantly ($M = 121.44$). The static and blank banners ranked approximately the same ($M = 107.83$ and 110.54, respectively).

Workload Factors. Participants found the flashing banners more frustrating and mentally demanding than the other banner types, $F(3,32) = 3.50$ and 1.62, respectively, $p < .05$. Marginally greater temporal demand and effort were reported, as well, $p < 0.1$. Figure 3 shows how each of the six TLX factors contributed to the overall workload by banner type.

3.3 Discussion

Search Times and Error Rates. Participants found targets just as quickly in the presence of animated banners as in all other conditions. This suggests that banner animation does not necessarily capture attention in a relatively simple visual search task. However, like the goal-driven participants in Pagendarm and Schaumburg's [2001] experiment, our participants also had an incentive *not* to look at the banners—the bonus pay decreased every second. Like Diaper and Waelend [2001] and Zhang [2001], we propose that animation's power to distract is dependent on the nature of the task. Thus, in Experiment 2, we introduced a more challenging and ecologically realistic task, hoping to elicit more conclusive search-time results.

Workload. Participants found that searching in the presence of the fastest animation, the flashing text banners, was the most frustrating and mentally demanding. However, the same results were not observed for the commercial animated banners, suggesting several possibilities. Comparing the two banner types, participants may have felt the animation of the commercial banners to be more subdued, and thus no more workload-intensive than the static banners. Or, the words in the flashing banners may have been similar enough to the targets to impose a cognitive burden, as predicted by the contingent involuntary orienting hypothesis [Folk et al. 1992]. Perhaps participants believed themselves capable of ignoring typical advertisements, whether animated or static, consistent with the findings of Benway and Lane [1998] and Bayles [2002]. Animated commercial banners were rated slightly higher than static banners for all workload factors, suggesting that for another task or with a different set of commercial banners, a significant effect might be observed.

Exit Interview. Participants discussed their search strategies and overall impressions in the post-experiment interview. In general, they reported being able to "tune out" the banners, although some found the flashing text and brightly-colored ones difficult to ignore. One participant even admitted to intentionally clicking the wrong word if she had not found the target within a few seconds just to make the banners disappear. Her error rate, however, was similar to that of other participants.

When asked to rank banners from most distracting to least, nearly every participant rated the flashing text banners as the most distracting. From there, participants varied in rankings, often citing color and type of animation as important factors. Many participants explained that the layouts with blank banners were the easiest to search, but others preferred the presence of banners because they helped to divide the screen into smaller search regions, and thus helped to structure the search. According to one participant, blank banners led to "too much of the same" visually. On average, blank banner layouts were indeed the slowest, but again, the difference was not significant.

The participants' comments and the uniform search times across banner conditions support the notion of banner blindness. The increase in perceived workload for the flashing banners and the marginal increase for commercial animated banners reflect the opinions of many Web users that despite the supposed blindness some price is paid. Therefore, under different task conditions the workload increase might translate to an increase in search time. Experiment 2 tests this hypothesis and adds eye tracking to explore the conditions under which people look at the banners.

4. EXPERIMENT 2

Experiment 1 suggested that animated banners might distract people performing visual search tasks, but might not necessarily slow them down. Previous studies [Zhang 2000, Diaper and Waelend 2002] indicated that the nature of the search task influences the strength of the distraction, and so we introduced a different task in the second experiment—searching through news headlines that appeared as links on a web page. To ensure ecological validity, we modeled the layout and format of links after several news sites, including CNN, Google News, and Yahoo! News.

Pilot tests for Experiment 2 revealed that many participants had difficulty with the point-completion deadline (PCD) on this task. Evidently, several participants wanted to use the mouse as a visual placeholder as they read, but the PCD would not allow it. Timeout errors were frequent. Researchers (such as Sears and Shneiderman [1994]) posit that people use the mouse as a visual placeholder when searching menus. Brumby and Howes [2004] observed this behavior, as well. Though participants do require a number of trials up front to grow accustomed to not moving the mouse before finding the target, the PCD seems to work fine, and participants have low timeout errors, in tasks that are closer to a laboratory visual search task, such as finding single words in Experiment 1. However, in the current experiment, participants did not grow accustomed to the PCD, and timeout errors remained high. This suggests that, in a real-world setting, moving the mouse may be an integral part of the visual search. More empirical work is needed to explore this possibility. We removed the PCD, letting participants move the mouse freely, as they would on the Web. Response times were actually much faster in the exact-match block in Experiment 2 than in Experiment 1.

To determine whether participants actually looked at the banners, we added eye tracking to collect fixation data. Tracking participants' gazes should also reveal common search patterns.

4.1 Method

Participants. Twenty-four undergraduate students from Lewis & Clark College (sixteen female) with a mean age of 21 participated in the experiment for compensation. All participants were experienced with graphical user interfaces and had normal or corrected-to-normal vision.

Materials. Three hundred twenty-nine headlines were gathered from humorous news sites like CNN's Offbeat News (http://www.cnn.com/offbeat) from April to September 2003. All were displayed in underlined, blue, 12-point Arial, a common format for Web links. One hundred twenty-four of the headlines were used as targets; the remaining headlines were distractors. Target headlines appeared once per participant; distractor headlines appeared up to three times.

The animated and static commercial banner types from Experiment 1 were reused in Experiment 2. The flashing text banners were removed for their lack of ecological validity. Gray rectangles were used in place of blank banners because of participants' previous suggestions that the banners helped to partition the search space.

Equipment. Eye movements were recorded using the LC Technologies' Eyegaze system. During data collection, participants used a chin rest to keep their heads relatively still. A small, unobtrusive camera was mounted below the computer monitor. Two separate computers were used in the data collection, one to collect the gaze position (represented by a small yellow plus sign on the screen) and one to run the experimental software. Both computers were connected to the same monitor with a two-way switch. Participants viewed the output from the Eyegaze computer during a short calibration. When they performed the experimental tasks, the monitor was switched to display the output from the computer running the experimental software. The software used to display the stimuli was directly derived from that used in Experiment 1.

The video signals from both computers were sent to a digital video mixer (Videonics model MX-1) where the plus sign showing the participant's gaze point was superimposed over the screen from the computer running the experiment using a chroma key effect. This composite image was then recorded to digital video and later transferred to Quicktime format for data analysis.

Design. To manipulate the mental workload of the search task, two precue conditions were used. In the *exact precue* condition, the precue contained the text of the target headline, word for word. For example, both the precue and target headline might be "Drop-outs doing just fine, thanks." In the *semantic precue* condition, a sentence or two from the beginning of the news article was used. No key content words in the headline appeared in the semantic precue; when necessary, synonyms found in the article were substituted. For example, the semantic precue for "Drop-outs doing just fine, thanks" was as follows:

> New research debunks the common belief that leaving school before completing year 12 diminishes a teenager's chance of a successful career.

In the semantic precue condition, participants could not merely look for a keyword in the target headline. Instead they had to read the headlines and compare them to the precue to find sufficient overlap in meaning to make the match. All precues were written in black, non-underlined, 14-point Arial so that they would have a slightly different overall appearance than the target and distractor headlines. This difference in font size and color would prevent participants in the exact precue condition from simply matching letter shapes.

Figure 4 shows the screen layout. Target and distractor headlines were arranged in two columns of six rows each. Each trial contained two banners. One always appeared at the top of the screen, directly above the area where the headlines were displayed. This location was selected to ensure that a participant's gaze would pass over a banner on every trial, and to mimic a common position of banner ads on the Web. The second banner was randomly placed in one of the six rows of the headline search area, spanning both columns. For each trial, both banners were of the same type (static, animated, or gray). Participants never saw both the animated and static versions of the same commercial banner.

Two blocks of trials were presented in a counterbalanced order. One block was the exact match condition; the other was semantic match. Each block consisted of 5 practice trials followed by 36 data collection trials, 12 trials each containing animated, static, or gray banner ads. The target headline appeared in a different position for each of the 12 trials. The type of banner presented was randomized across trials and within blocks. Banner and target combinations in error trials were repeated, shuffled into the remainder of the block.

Procedure. Search trials proceeded in the same manner as in Experiment 1. Participants studied the precue for an unlimited period and when ready, clicked a box to make the precue disappear and the search area appear. Unlike in Experiment 1, participants were allowed to move their mouse while searching. The cursor changed from an arrow to a hand over the headlines, as it would for links on the Web. After locating the target, participants selected it; it would briefly turn magenta and either the reward chime or penalty buzzer would sound. The payoff matrix was the same as in Experiment 1, except that exact search trials started with a potential bonus of 9 cents bonuses and semantic search trials with a bonus of 14 cents. One cent was deducted per second, and errors earned a five-cent penalty.

After the visual search tasks, participants were given a short break and then asked to view and identify banners that were shown in the study. This was the first

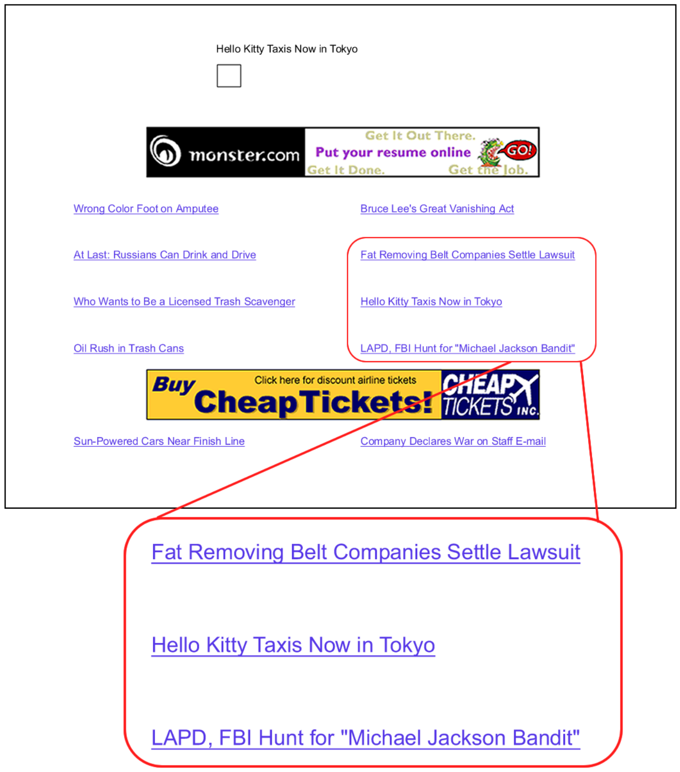

Fig. 4. Screen layout for a literal precue trial with a zoomed-out view of three headlines. The precue at the top disappeared when the layout appeared.

mention of the banners to the participants. It was explained that they would see some banners that had been in the study and others that were not. The banners were displayed on the screen one at a time, and participants responded by clicking a "yes" or "no" button at the bottom of the screen. Each click triggered the presentation of the next banner. A total of 60 banners were presented (30 animated and 30 static). Of these, 40 banners had appeared during the visual search tasks and 20 had not. Participants were not given feedback on accuracy for this memory task and speed was not recorded or emphasized.

4.2 Results

Search time. The type of precue (exact vs. semantic) produced the strongest effect in the experiment. Search times for the exact precue condition (M = 2134 msec., *Standard Deviation* = 299 msec.) were much faster and less variable than for the semantic precue condition (M = 6129 msec, SD = 1567 msec.), $F(1, 23) = 231$, $p < .0001$. Due to overwhelming differences and the unequal variance in the search times for these precue conditions, the remaining search time analyses are broken down by precue condition.

Table II shows the mean search time for each banner type. For the exact precue condition, a repeated measures ANOVA revealed a significant difference among the banner types $F(2,46) = 5.5$, $p < .007$. Post hoc paired t-tests showed that both the static

Table II. Mean Search Times for Each Banner Type in Experiment 2

Banner type	Search time (ms)	Standard deviation (ms)
Exact precue		
Gray	2,040	289
Static	2,169	300
Animated	2,193	297
Semantic precue		
Gray	6,065	1,614
Static	6,210	1,736
Animated	6,110	1,397

Note. Times are averaged across target positions and participants.

and animated banners resulted in slower search times compared to the gray banners ($p < .005$ for both comparisons), but equivalent search times when compared to each other ($p = .65$).

Significant differences do not persist in the semantic precue condition. A repeated measures ANOVA revealed no significant differences in search times for the different banner types $F(2,36) = 0.18$, $p = .83$. While the general pattern is that banner ads appear to have slowed the search process even in the semantic precue condition, the high between-subject variability overwhelms the statistical significance of the difference.

The error rates in this experiment were uniformly low in both precue conditions (4.6%). There was no significant correlation between speed and accuracy, $r = -.103$, $p = .63$.

Memory. Participants responded "yes" or "no" to each ad during the recognition memory test. A "yes" to a banner that did appear in the experiment is a *hit*. A "yes" to a banner that did not appear is a *false alarm*. The number of "no" responses to banners that did appear (misses) and that did not (*correct rejections*) can be derived from the number of hits and false alarms.

The hit and false alarm rates are used to assess the participants' memory for the banners. Overall, memory for the banner ads was poor, with a hit rate of only 20.1% and a corresponding false alarm rate of 20.2%. Perfect performance would have been 100% and 0%, respectively. The hit rates did not differ by precue condition (exact = 20.0%, semantic = 20.2%), [$X^2(1, N = 24) = .008$, $p = .94$]. Though the increased mental workload of the semantic precue search task greatly increased visual search time, it did not affect memory for the banner ads.

Though recognition memory for the banners was low overall, Signal Detection Theory can be used to show that recognition memory was better for the static banners. The hit and false alarm rates were transformed into a single measure of memory strength known as d' [Green and Swets 1966]. A positive, non-zero d' value is an indication of memory strength, controlling for guessing behavior and decision strategies that participants might adopt. One group t-tests showed that the d' value for the static banners ($M = 667$) is significantly higher than zero, $t(23) = 2.66$, $p = .01$, while the d' for the animated banners ($M = -.07$) is not, $t(23) = -0.3$, $p = .77$. A paired t-test further revealed that the d' score for the static banners is significantly higher than for the animated banners, $t(23) = 2.14$, $p = .04$. This shows that when we correct for their guessing strategies, the participants have significantly worse memory for the animated banners than for the static.

One final round of analyses on the memory for banner ads concerns their screen location. For each search trial, one ad was placed in the same location at the top of the screen, between the precue and the topmost headlines. The second ad was randomly placed in one of the six rows of the headline search area, spanning both columns. By

Table III. Correlation Between Number of Items Viewed and Search Time

Banner Type	Correlation
Exact precue	
Gray	.749**
Static	.536*
Animated	.790**
Semantic precue	
Gray	.702**
Static	.549*
Animated	.372

**p* < .05 ** *p* < .001.

combining the hit rate data across all 24 subjects we are able to determine that screen location affected the recognition memory for these banners. Overall, there was a trend for the top banner to be better remembered [$X^2(1, N = 24) = 2.83, p = .09$]. Breaking the data down further by banner types indicated that the top ad was remembered significantly better than the randomly-placed ad for the static banners ($p = .01$), but not for the animated banners ($p = .82$). Note that the *d′* analysis is not used here because too many participants had hit and false alarm rates of 0 in this combination of factors, thus making *d′* undefined.

Eye tracking. A digital video composite was created by superimposing the screen output from the two computers used in the experiment: one that collected gaze position (represented by a small plus sign) and another that presented the experimental software. Fixations were encoded by watching the video. For each trial, it was noted whether the first saccade (rapid, ballistic eye movement) from the precue occurred before or after the appearance of the banners and headlines (hereafter collectively called "items"). Each time the gaze landed on a new item, the item number was recorded. Multiple saccades within a single item, such as those necessary to read a headline, were not counted. Revisits to an item were counted, as long as the gaze moved to another item and then returned.

Participants looked at nearly twice as many items in the semantic condition ($M = 8.35$) as in the exact condition ($M = 4.92$), $F(1, 16) = 120.86, p < .0001$. Across blocks, the number of items viewed was not affected by banner type, $F(2,32) = 1.11, p = .34$, although in the exact precue block, a marginally significant difference was observed, $F(2,34) = 2.96, p = .07$. Gaze time per item was also greater in the semantic condition ($M = 786.73$) than the exact condition ($M = 452.14$).

There was a strong positive correlation between the number of items viewed and the search time, $r = .845, p < .0001$. Table III shows the correlations for each combination of banner type and precue type. One exception was noted: there was no correlation in the semantic precue condition when animated banners were present.

Participants looked at banners in 11.7% of the trials, regardless of banner type, $F(2,32) = 1.28, p = .29$. They looked at gray banners ($M = 8.8\%$) almost as often as the static ($M = 13.0\%$) and animated ($M = 13.2\%$) ones. Of the trials in which participants looked directly at banners, 70% of the banner gazes occurred during the participant's first eye movement. The precue type (exact versus semantic) did not affect whether participants looked at banners, $F(1,16) = .137, p = .72$. Of the 164 banners that participants correctly remembered in the memory test (hits), only 10 received direct gazes during the search portion of the experiment. Thus, participants did not directly look at 93.9% of the banners they "remembered."

Table IV. Number of Trials with Direct Fixations on Top and Inside Banners

	Top banner		Inside banner	
Precue type	Total	On first fixation	Total	On first fixation
Exact	10	6	58	47
Semantic	36	25	38	21

Table IV shows how precue type and banner position affected whether participants looked at banners. Banners that were positioned in a random row within the search area ("inside" banners) were viewed more often than banners that appeared in a fixed location between the precue and topmost headlines ("top" banners), $F(1,16) = 6.73$, $p = .02$. Participants rarely looked at the top banner in the exact condition but did look at inside banners. In the semantic condition, gazes at top and inside banners were approximately equal.

Participants frequently made anticipatory fixations: eye movements from the precue to the search area before the headlines and banners appeared. Anticipatory fixations occurred in approximately 40% of the trials, irrespective of precue type (M = 40.4% for semantic and 42% for exact), $p = .76$. Of the trials in which participants looked directly at banners on their first eye movement, 54% were anticipatory gazes, in which case the banners appeared *after* participants had moved their eyes to the location.

Regardless of whether the eye movement away from the precue occurred before or after the onset of the layout, the gaze was almost always near the top of the screen when the layout appeared. In 73% of all trials, at the moment that the layout (and thus the top banner) appeared, the gaze was either on the precue, on the top banner position, or in position for the top row headlines.

4.3 Discussion

Search time. Consistent with the findings of Zhang [2000] and Diaper and Waelend [2000], the nature of the task influenced the effect of the distractors. Participants searching for the simple exact match were more adversely affected by banners than were participants searching for the harder semantic match. This can be seen in the significant increase in search times in the exact match condition when commercial banners were present, and the non-significant difference in search times in the semantic match condition. One possible explanation is that simple searches require less attention, and thus other stimuli can be perceived and processed, slightly delaying the primary task. It may be that difficult searches necessitate greater focus, leaving less processing available for irrelevant objects, which are thus ignored. The high between-subject variability in the semantic condition suggests an alternative explanation: by their very nature difficult tasks result in a wider range of response times, making the banners' effect harder to detect.

Contrary to the prevailing logic of designers, animated commercial banners affected performance no differently than static commercial banners. Search times were slower for both commercial banner types than for the gray banners, but were equivalent when compared to each other. Given that the animated banners had graphics roughly identical to their static versions, attention capture could be attributed to the images, colors, and text within the banners, rather than their motion. Even in the semantic condition, where participants were exposed to an average of six seconds of animation per trial, the animation did not affect search times. In fact, though the difference was not significant, searches tended to be slightly faster with animated banners than static in the semantic condition, loosely analogous to the findings of Pashler [2001].

Participants looked at twice as many items per trial when searching for a semantic match than when searching for an exact headline. This result is consistent with the nature of the task. In the semantic condition, the best match could rarely be determined in isolation; a potential target had to be compared to the remaining headlines. Return visits to a headline were common. Additionally, gaze time per item was greater in the semantic condition, as participants considered the content of each headline.

The strong positive correlation between the number of items viewed and search time is expected, but the lack of a correlation in the semantic condition when animated banners were present is surprising. Some other factor must influence these variables. Perhaps participants consider headlines for a much wider range of times. This is speculation, and one of the many possible explanations, but perhaps the animation increases this variance.

Fixations on banners. Graphics and animation in the commercial banners did not attract participants' gazes: Direct fixations occurred on the gray banners as often as the commercial ones. Though the static and animated banners did increase search time, the increase cannot be attributed to participants looking directly at the banners and thus processing their detailed content. Instead, the delay might be caused by graphics and animation viewed peripherally.

Participants' initial eye movements further suggest conscious efforts to avoid banners. The majority (70%) of fixations on banners occurred in the first eye movement. Half of these (54%) happened when the eyes moved to a location before the banner appeared there. A maximally efficient search strategy would avoid banners altogether, and it is possible that at the time of the first eye movement, participants had not yet encoded the locations of the banners, and thus, any banner fixations were accidental. The low number of top banner fixations in the exact condition further supports the idea of banner avoidance. Participants knew the top banner would always appear in the same place, so it was easier to avoid, at least in the exact match block. Just as the nature of the search task influenced the degree to which different banner types delayed search, it also impacted at which banner positions participants looked. A simpler search allowed for the formation of a strategy to avoid the known banner location. In the semantic condition, top banner fixations occurred as often as inside banner fixations. Perhaps the semantic processing precluded the formation of a highly efficient visual search strategy. Participants may have been so busy interpreting the precue that their eyes simply drifted toward the top banner location.

Memory. The simplest explanation for participants remembering certain banners would be that they looked directly at those banners. However, the eye-tracking data do not support this explanation: 94% of the banners correctly identified in the recognition test had not been directly looked at by the participant. Perhaps graphical elements in the banners were salient enough peripherally to aid memory, but more likely participants' recognition scores were determined by luck. The high false alarm rates suggest as much. The low memory rate is consistent with Bayles [2002] and Pagendarm and Schaumburg [2001]: People simply do not remember banners that are irrelevant to their goals. Recognition also did not vary by length of exposure. Banners in the semantic condition appeared three times longer than those in the exact condition, but participants did not remember them any better.

Static banners fared somewhat better than animated banners in the recognition test. Neither Bayles [2002] nor Pagendarm and Schaumburg [2001] had compared the memorability of static to animated banners, but advertisers employ animation so often that we were surprised to find it to be an ineffective memory aid. Though neither type of banner was very memorable, a small percentage of static banners was correctly identified, after we accounted for participants' guessing strategies. Animated banners were recognized no better than chance. In part, this may be due to the changing appearances of these banners: a slogan or logo that disappeared or moved may not have been visible or obvious during the brief period in which participants processed the banners. This is not to say that participants were not exposed to a full cycle of animation in each trial, simply that they did not pay attention to the entire cycle. They may have observed too few details to recognize the banners later. Static banners, in which the message appeared continuously throughout the trial, were more memorable.

Banners positioned at the top of the screen were remembered more often than the inside banners. This is surprising, given that participants looked directly at the top banners less frequently. However, as was mentioned, the gaze was usually near the top of the screen when the layout and thus the top banner appeared. Perhaps attention was slightly drawn to the onset of the top banner and, since that banner was relatively close to the gaze at that time, a small amount of visual information was processed.

Like the goal-directed searchers in Pagendarm and Schaumburg [2001], the participants in the present experiment had little incentive to look at the banners. Nonetheless, the commercial banners slowed visual search responses. In trials where the task was easy enough to allow participants to formulate a search strategy, they intentionally avoided looking in places where banners were known to appear. On Web sites where banner ads are unrelated to page content and viewers' goals, the same results may be expected: ads will increase visual search times and people will attempt to avoid looking at them.

5. GENERAL DISCUSSION

The popular notion of "banner blindness" suggests that people just ignore banner ads. Nonetheless, may web users still dislike them. Motivated by the seeming contradiction between "banner blindness" and distaste of ads, we discovered results consistent with both schools of thought. Typical commercial banner ads hinder searches through lists of links, even if the searchers never look directly at the banners. Given the prevalence of animation on the Web, it was surprising to find little disparity between the effects of animated and static versions of the same ads. We had hypothesized that search times and workload for animated commercial banners would be greater than those for static banners. However, people do not look at animated banners more often than static ones, and they can search equally quickly under both. Extreme forms of animation, such as the flashing cyan banners from Experiment 1, do increase frustration and perceived mental demand, but mainstream animated banners performed no worse than their static equivalents, at least when compared head-to-head with extreme flashing banners. The one surprising difference between animated and static banners was that animation makes ads less memorable. Contrary to widespread practice, animation may not provide a benefit to advertisers; in fact, it may even reduce an ad's success.

In this task, there is evidence of both goal-directed (top-down) attentional control and stimulus-driven (bottom-up) attentional capture. Though people rarely looked at the banner ads, when they did it was independent of banner type. This suggests that participants adopted search strategies that enabled them to ignore irrelevant distractors, consistent with the notion put forth by Folk et al. [1992] and others that people adopt task-strategies that prevent involuntary shifts of attention to at least some stimulus properties.

Nonetheless, there is some evidence for stimulus-driven capture. In Experiment 2, participants searched longer in the presence of commercial banners than in the presence of gray rectangles. This suggests that some graphical elements may have captured attention regardless of the participant's strategies. Interestingly, it is unlikely that the stimulus-driven attentional capture was due to animation, as animated banners did not have greater search times than static banners.

Unlike the previous studies of animated distractors on the Web [Bayles 2002, Diaper and Waelend 2000, Pagendarm and Schaumburg 2001, Zhang 2000], the present experiments include eye tracking data that reveal the underlying behavior of people searching in the presence of the distractors. People rarely look directly at banners, and adding graphics does not appear to matter. Gray banner "placeholders" were fixated as frequently as commercially-designed banners. In fact, most banner fixations in Experiment 2 occurred on participants' first eye movements, perhaps before they encoded the banner locations so as to avoid them. The infrequent fixations on top banners in Experiment 2 and the "sandwich effect" from Experiment 1 (in which targets sandwiched between two rows of flashing banners took longer to find) further support the idea that people intentionally avoid looking in locations where they expect banners, at least once the search has begun.

Clearly, the nature of the primary task strongly interacts with the attention-getting capacity of the banners. Across the two experiments, three kinds of simple visual search tasks with only subtle differences between them were tested and they resulted in markedly different outcomes. The negative effects of banner ads are subtle and not always easy to directly measure. Experiment 1 may not have been ecologically realistic enough to elicit a search-time disparity. The semantic condition in Experiment 2 led to high between-subjects variability, again potentially concealing a time effect. Diaper and Waelend [2000] present yet another study in which the effect could not be detected. Even in a highly controlled environment, the interaction between banner factors, such as color and semantic appeal, and task factors, such as participants' reading abilities, make quantifiable results difficult to detect.

The negative effects discussed here apply directly to experienced Web searchers who know exactly—or nearly—for what they are looking. To them, banners are irrelevant. Many Web surfers fall into Pagendarm and Schaumburg's [2001] "aimless browser" demographic. The effect of banner ads on this population is still unclear.

Practical implications. Web designers and site owners should post ads closely related to page content if they hope to attract their viewers' attention. Participants in the present studies had an overriding incentive not to look at banners, and no amount of banner manipulation increased their pull. Longer exposure time, animation, and the presence of images did not make the task-irrelevant ads more conspicuous. Connecting advertising to viewers' goals may make ads more successful; Yahoo! received positive feedback when it deployed ads related to page content (see Rohrer and Boyd [2004] for a discussion of user experience and advertising).

Banners positioned at the top of the screen may be more memorable, although this effect could be due to the specific screen layout of the present experiments. Top banners were favorably situated between the precue and the content to be searched. Browser address bars and standard site navigation areas appear roughly in the same screen region as our precue; perhaps Web users' eyes might follow paths similar to those observed experimentally. Designers should be wary, nevertheless, of habituating viewers to predictable banner locations: People avoid looking in areas where they expect to find ads. People's success in avoiding banners may be dependent on the cognitive complexity of their tasks; top banner avoidance was only clearly observed in the exact precue condition of Experiment 2, but the Web presents numerous tasks of varying complexity. Usability guidebooks (i.e., Nielsen and Tahir [2002]) warn against placing site navigation above banner ads to lessen the spread of banner blindness to critical page elements. The "sandwich effect" from Experiment 1 supports this advice.

Further investigation is needed into all aspects of visual search on the Web. Traditional attention-capture studies, such as those discussed in Section 1.1, explain some search behaviors, but the myriad of interacting Web factors should be explored in a more ecologically valid context. Additional studies of animated distractors are needed for more involved tasks, such as multi-page surfing and form-filling. Though the results are too premature to report here, we noticed dramatically different gaze paths over the same headline layouts in Experiment 2, depending on whether the precue had been exact or semantic. That people may scan identical screens differently based on mental load has implications for cognitive modelers, especially those seeking to predict eye movements. Future empirical studies into other tasks may lead to a more comprehensive understanding of visual search on the Web. In the present experiments, the implications are clear: banner ads degrade visual search and are quickly forgotten.

REFERENCES

ABRAMS, R.A., MEYER, D.E., AND KORNBLUM, S. 1990. Eye-hand coordination: oculomotor control in rapid aimed limb movements. *Journal of Experimental Psychology: Human Perception and Performance 16*, 248-267.

BACON, W.F., AND EGETH, H.E. 1994. Overriding stimulus-driven attentional capture. *Perception & Psychophysics* 55, 485-496.

BAYLES, M. 2002. Designing online banner advertisements: should we animate? In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Minneapolis, MN, April 2002, ACM Press, New York, 363-368.

BENWAY, J.P., AND LANE, D.M. 1998. Banner blindness: Web searchers often miss “obvious” links. *Internetworking 1*.

BRUMBY, D.P., AND HOWES, A. 2004. Good enough but I’ll just check: Web-page search as attentional refocusing. In *Proceedings of the Sixth International Conference on Cognitive Modeling*, Pittsburgh, PA, August 2004, 46-51.

DIAPER, D. AND WAELEND, P. 2000. World Wide Web working whilst ignoring graphics: good news for Web page designers. *Interacting with Computers 13*, 163-181.

ENNS, J.T., AUSTEN, E.L., DI LOLLO, V., RAUSCHENBERGER, R., AND YANTIS, S. 2001. New objects dominate luminance transients in setting attentional priority. *Journal of Experimental Psychology: Human Perception and Performance 27*, 1287-1302.

FOGG, B.J., MARSHALL, J., LARAKI O., OSIPOVICH, A., VARMA, C., FANG, N., PAUL, J., RANGNEKAR, A., SHON, J., SWANI, P., AND TREINEN, M. 2001. What makes Web sites credible? A report on a large quantitative study. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Seattle, WA, April 2001, ACM Press, New York, 61-68.

FOLK, C.L., REMINGTON, R.W., AND JOHNSTON, J.C. 1992. Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance 18*, 1030-1044.

FRANCONERI, S.L., AND SIMONS, D., J. 2003. Moving and looming stimuli capture attention. *Perception & Psychophysics* 65, 999-1010.

GREEN, D.M., AND SWETS, J.A. 1966. *Signal Detection Theory and Psychophysics*. Wiley, New York.

HART, S., AND STAVELAND, L.E. 1988. Development of NASA-TLX (task load index): results of empirical and theoretical research. In *Human Mental Workload*, P.A. HANCOCK AND N. MESHKATI, Eds. Elsevier Science Publishers B.V., Amsterdam, 139-183.

HILLSTROM, A.P., AND YANTIS, S. 1994. Visual motion and attentional capture. *Perception & Psychophysics 55*, 399-411.

HORNOF, A.J. 2001. Visual search and mouse-pointing in labeled versus unlabeled two-dimensional visual hierarchies. *ACM Transactions on Computer-Human Interaction, 8,* 171-197.

INTERACTIVE ADVERTISING BUREAU. 2001. IAB/Dynamic Logic ad unit effectiveness study. Available from http://www.iab.net/resources/dynamic_lbrandrl.asp

JONIDES, J. AND YANTIS, S. 1988. Uniqueness of abrupt visual onset as an attention-capturing property. *Perception & Psychophysics 43*, 346-354.

NIELSEN, J. AND TAHIR, M. 2001. *Homepage Usability: 50 Websites Deconstructed*. New Riders, Indianapolis, IN.

PAGENDARM, M. AND SCHAUMBURG, H. 2001. Why are users banner-blind? The impact of navigation style on the perception of Web banners. *Journal of Digital Information 2*.

PASHLER, H. 2001. Involuntary orienting to flashing distractors in delayed search? In *Attraction, Distraction, and Action: Multiple Perspectives on Attentional Capture*, C. FOLK AND B. GIBSON, Eds. Elsevier Science B.V., Amsterdam, 77-92.

POST, R.B., WELCH, R.B., AND BRIDGEMAN, B. 2003. Perception and action: Two modes of processing visual information. In *Visual perception: The influence of H. W. Leibowitz,* J. ANDRE AND D.A. OWENS, Eds. American Psychological Association, Washington, DC, 143-154.

REMINGTON, R.W., JOHNSTON, J.C., AND YANTIS, S. 1992. Involuntary attentional capture by abrupt onsets. *Perception & Psychophysics 51*, 279-290.

ROHRER, C. AND BOYD, J. 2004. The rise of intrusive online advertising and the response of user experience research at Yahoo!. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Vienna, April 2004, ACM Press, New York, 1085-1086.

ROWE, N.C., COFFMAN, J., DEGIRMENCI, Y., HALL, S., LEE, S., AND WILLIAMS, C. Automatic removal of advertising from Web-page display. In *Proceedings of the Second ACM/IEEE-CS Joint Conference on Digital Libraries*, July 2002, Portland, OR, 406.

SCHUMACHER, E.H., LAUBER, E.J., GLASS, J.M., ZURBRIGGEN, E.L., GMEINDL, L., KIERAS, D.E., MEYER, D.E. 1999. Concurrent response-selection processes in dual-task performance: Evidence for adaptive executive control of task scheduling. *Journal of Experimental Psychology: Human Perception & Performance 25*, 791-814.

SEARS, A. AND SHNEIDERMAN, B. 1994. Split menus: effectively using selection frequency to organize menus. *ACM Transactions on Computer-Human Interaction 1*, 27-51.

THEEUWES, J. 1991. Cross-dimensional perceptual selectivity. *Perception & Psychophysics 50*, 184-193.

YANTIS, S. 2000. Goal-directed and stimulus-driven determinants of attentional control. *Attention and Performance 18*, 73-103.

YANTIS, S. AND EGETH, H.E. 1999. On the distinction between visual salience and stimulus-driven attentional capture. *Journal of Experimental Psychology: Human Perception and Performance 25*, 661-676.

YANTIS, S. AND JONIDES, J. 1990. Abrupt visual onsets and selective attention: voluntary versus automatic allocation. *Journal of Experimental Psychology: Human Perception and Performance 16*, 121-134.

ZHANG, P. 2000. The effects of animation on information seeking performance on the World Wide Web: securing attention or interfering with primary tasks?. *Journal of the Association for Information Systems 1*, 1-28.