Banner Ads Hinder Visual Search and Are Forgotten

Moira Burke[1], Nicholas Gorman[2], Erik Nilsen[2], and Anthony Hornof[1]

[1]Computer and Information Science Department
University of Oregon, Eugene, OR 97403
{moira, hornof} @cs.uoregon.edu

[2]Psychology Department
Lewis & Clark College, Portland, OR 97219
{gorman, nilsen} @lclark.edu

ABSTRACT
Banner ads persist on the Internet in spite of evidence against their efficacy. Many ads include animation in an attempt to increase their attentional capture. An experiment was conducted to examine how various banner ads affect the visual search of news headlines on the Web, and whether participants could recall the ads they saw. The results both support and contradict the notion of "banner blindness," the idea that people ignore banner ads. Participants could not recall the ads that they saw, but those ads did distract the users and significantly increased search times. The most surprising result is that recall was especially bad for animated banners. This paper examines issues of attentional capture in an applied domain and provides guidance for web designers.

Author Keywords
Animation, visual search, banner ads, memory, World Wide Web.

ACM Classification Keywords
H.5.2. User Interfaces: GUI, Screen Design

INTRODUCTION
One of the most prevalent forms of advertisement on the Web today is the banner. While they are increasingly varied in appearance and size, the prototypical banner ad is a 468 x 60 pixel image, often found at the top of a web page. This study examines two aspects of banner ads—animation and location—to quantify their impact on visual search performance for news headlines and recognition memory for the advertisements.

Banner ads became prevalent in the mid-nineties, and their usage was quickly challenged. Benway [2] coined the term "banner blindness" in 1998 with his observation that only 20% of users could recall even the presence of brightly colored banners. One e-journalist called banner ads the Web's "preferred means of exchanging ideas through hyperventilation, screeching, and hooting" [6]. Yet in 2004, banner ads are still ubiquitous, and have increased in size on some sites, such as NYTimes.com.

In an attempt to improve the effectiveness of banner ads, designers have incorporated animation into the advertisements to capture users' attention. A frequently cited ZDNet report [8] claims that the use of animation improves click-through rates up to 40%.

Scientific research on animation and attentional capture is more equivocal. While there is clear support for the idea that moving objects sometimes capture our visual attention [4], there is still a healthy debate concerning whether this is an automatic process or if it is dependent on the relevance of the animation to the person's task [7].

A study of animated icons found that abrupt visual onset (such as flashing) leads to longer search times, while static discontinuities (such as simple feature changes) do not [5]. On the other hand, a few research studies have found that irrelevant animations that do not share critical properties of the search targets sometimes lead to equivalent or even faster performance [7].

There have been several HCI studies related to the effects of animation in the recent past. Zhang [9] had users count letter strings in a display that had simple animated letter strings and images appearing at the top, side, or bottom of the screen. Both types of animation interfered with performance. She also found that the negative effect of the animation was greatest with easier counting tasks. Greater mental workload reduced, but did not eliminate the negative impact of animation. Bayles [1] studied memory for animated and static banner ads. Participants completed four information search tasks on a web page that included one animated and one static banner ad that were not relevant to the search tasks. Only 40% of subjects correctly recalled at least one of the banner ads. There was no difference in the recall or recognition of static versus animated banners. One shortcoming of the Bayles study was that there were only two ads so it is hard to generalize her findings. A shortcoming of the Zhang study is that the task and the animations are unlike those on the Web.

The study presented here examines both (a) the impact of banner ads on task performance and (b) memory for the ads

CHI 2004, April 24–29, 2004, Vienna, Austria.
ACM 1-58113-703-6/04/0004.

in a visual search task that is common on the Web: people looking through lists of links while ads appear around the page. Participants search through lists of real news headlines on screens that contain pairs of banner ads selected from a set of 100 advertisements obtained from commercial websites. Our study manipulates the type of banner ads used (static, animated, or blank), the screen location of the banners (top or embedded within the headline list), and the mental workload of the search task (literal or semantic precue). These manipulations allow us to address a number of questions of interest to both researchers and web designers: Do animated banners affect a visual search task more than static banners or no banners? Are animated banners more memorable than static banners? Does the screen placement of the banners matter for task performance or recognition memory? How are the answers to these questions influenced by the mental workload of the user when performing the visual search?

METHOD

Participants

Twenty-four adults (sixteen female) ranging from 19 to 22 years of age participated in the experiment for compensation. Every participant was paid a minimum of $10, but each had the opportunity to earn an additional bonus of up to $8 based on speed and accuracy. All participants were experienced web users and had normal or corrected-to normal vision.

Equipment and Materials

The visual stimuli were presented on a Planar 17" LCD monitor with 0.28 mm pitch, controlled by a 350 MHz Pentium II processor running Windows 98. The participants responded using a new optical Logitech Wheel Mouse.

The participants' eye movements were recorded using 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. The analysis of the eye movements is ongoing and this data will not be reported in this paper.

Two computer programs were developed for the study. One presented the visual search task and collected reaction times and error rates. The other assessed participants' memory for the banner ads. The software was written in Lingo, Macromedia Director's scripting language.

Three banner types were used: (a) blank (gray) banners, (b) animated commercial banners, and (c) static commercial banners. A selection of one hundred animated commercial banners was chosen from popular news websites, search engines, and portals, including the New York Times website, AltaVista, and AOL. Static commercial banners were created by extracting a representative frame from each animated commercial banner.

Design and Procedure

For each experimental trial, the participant was instructed to find a particular news headline from a hyperlinked list and click on it with the mouse. Each trial had two parts: the *precue stage* and the *search stage*. Two different types of precues were used to manipulate the mental workload of the search task. In the easier *literal precue* condition, the precue contained the exact headline, word for word. For example, both the literal precue and the target headline might be "Drop-outs doing just fine, thanks."

In the more challenging *semantic precue*, a sentence or two from the beginning of the news article was used. Special care was taken to make sure that none of major words in the headline appeared in the semantic precue. 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 this 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 with the precue to find sufficient overlap in meaning to make the match.

In each precue condition, participants were given as much time as they wanted to read and memorize the precue. When ready, the participant clicked on a button beneath the precue. This made the precue disappear and the layout appear, thus initiating the search stage (see Figure 1 for screen layout).

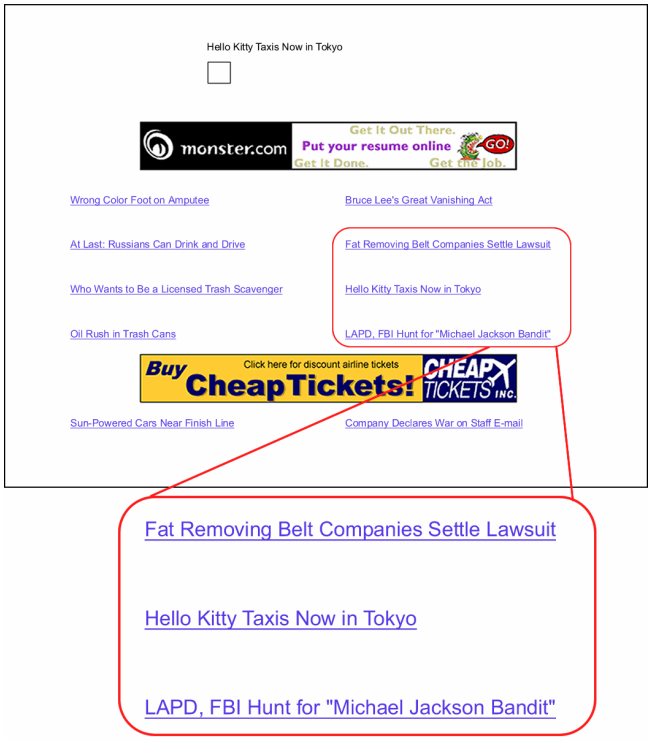

Figure 1. Screen layout for a literal precue visual search trial with a zoomed-out view of three headlines. The precue at the very top disappeared when everything else appeared.

In the search stage of an experimental trial, two columns of text filled the screen area (see Figure 1). Each column had six rows, five of which contained headlines. The remaining row contained a banner spanning both columns. Headlines were 12-point, blue, underlined Arial text, a common format for hyperlinks on the Web. Only one of the headlines was the target item; the remainder acted as distractors. Each target headline was used in the experiment just once, and each distractor headline appeared no more than three times. The headlines and precues were selected from offbeat articles published from April – September, 2003 in actual online news sites (e.g. CNN.com/offbeat). Participants located and clicked on the target headline as quickly as possible. Auditory feedback provided an assessment of accuracy, and then the precue for the next trial appeared at the top of the screen. Error trials were repeated later in the experiment.

Two blocks of trials were run in a counterbalanced order. Each block used a single type of precue (literal or semantic) and consisted of 5 practice trials followed by 36 data collection trials (12 trials each containing animated, static or blank banner ads). The target headline appeared in a different position for each of the 12 trials. Each trial included two banner ads. One banner was always located at the top of the screen, directly above the area where the headlines were displayed. This location was selected to ensure that the participants' gazes 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, blank). The type of banner presented was randomized across trials and within blocks.

Immediately following the visual search tasks, subjects were given a short break and then asked to view and identify banners that were shown in the study. At this point they were told for the first time that they would be asked to identify banners from the study. It was explained that they would see some banners that were in the study and others that were not. The banners were shown on the screen one at a time and participants responded by clicking on a "yes" or "no" button at the bottom of the screen to indicate whether they had seen each ad earlier in the study. Each click brought up the next banner. A total of 60 banners were presented (30 animated and 30 static). Of these, 40 banners were actually shown during the visual search tasks and 20 were not. Participants were not given feedback on accuracy for this memory task and speed was not recorded or emphasized.

RESULTS

Search Time and Error Rates

The mental workload manipulation caused by the type of precue (literal vs. semantic) produced the strongest effect in the experiment. The literal precue condition (Mean=2,134 msec., Standard Deviation= 299 msec.) was much faster and less variable than the semantic precue condition (M=6,129 msec., SD= 1,567 msec.), $F(1, 23) = 231$, $p<.0001$. Due to overwhelming differences and the unequal variance in these precue conditions, the remaining search time analyses are broken down by precue condition.

Banner Type	Mean Search Time (in msec.)	St. Dev. (in msec.)
Literal Precue Condition		
Blank	2040	289
Static	2169	300
Animated	2193	297
Semantic Precue Condition		
Blank	6065	1614
Static	6210	1736
Animated	6110	1397

Table 1. Search times for each banner type separated by precue condition, averaged across target position and participant. The standard deviations are also shown.

The top of Table 1 shows the mean search time for each banner type for the literal 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 and animated banners resulted in slower search times compared to the blank banners ($p<.005$ for both comparisons), but equivalent search times when compared to each other ($p=.65$).

The bottom half of Table 1 shows the mean search time for each banner type for the semantic precue condition. A repeated measures ANOVA revealed no significant differences among the banner types $F(2, 46) = 0.18$, $p=.83$. 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 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, $r= -.103$, $p=.63$.

Recognition Memory for Banners

Participants responded "yes" or "no" to each ad during the recognition memory test. If participants respond "yes" to a banner ad that was in the experiment, this is termed a *hit*. Answering "yes" to a banner that was not shown in the study is termed a *false alarm*. The number of "no" responses to banners that actually appeared in the study (a *miss*) and to those that did not (a *correct rejection*) can be derived directly from the number of hits and false alarms. Thus, we just report the hit and false alarm rates here to assess the participants' memory for the banners.

Overall, memory for the banner ads was quite 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 reliably by precue condition (literal precue = 20.0%, semantic precue = 20.2%), [X^2 (1, N=24) = .008, p=0.94]. Though the increased mental workload of the semantic precue search task greatly increased visual search time, it did not affect the memory for the banner ads at all.

Though recognition memory for the banners was very poor overall, recognition memory was better for the static banners, as shown using Signal Detection Theory to transform the hit and false alarm rates into a single measure of memory strength known as d' [3]. 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 guessing strategies of the participants, they 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 banner ad was placed in the same location at the top of the screen, in between the precue and the news headlines. The second banner ad was randomly placed in one of the six rows of the headline search area, spanning both columns. By 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=0.09]. Breaking this data down further by banner types indicated that the top banner was remembered significantly better than the randomly-placed banner for the static banners (p=.01), but not for the animated banners (p=.82).

CONCLUSION

The study presented here contradicts the notion of "banner blindness," that people just ignore and are effectively "blind" to banner ads [2]. The current study demonstrates that, yes, it is true that people do not remember the ads that they see, but it also demonstrates that people are not blind to the distracting effects of the advertisements. The ads do interfere with a user's primary visual search task.

The study also reveals some surprising effects of animation in banner ads. Though online marketers continue to add interactivity and obtrusiveness with Flash, this experiment shows that animation actually hinders ad recall. It appears as if web designers and site hosts should be cautious of burdening their customers with such stimuli. Other metrics of ad success, such as click-through rates, are not evaluated here, but given the high number of ad impressions required to generate click-through revenue, the accumulated impediment those impressions cause to typical viewers is substantial.

Further analysis of the eye tracking data will reveal whether participants even looked at the banners or if they intentionally adopted search strategies to avoid them. Future studies may also indicate whether some factors, such as color or imagery, increase memorability.

ACKNOWLEDGMENTS

This research is supported in part by the National Science Foundation through Grant IIS-0308244 to the University of Oregon with Anthony Hornof as the principal investigator.

REFERENCES

1. Bayles, M.E.. (2002). Designing online banner advertisements: should we animate?, *Proceedings of CHI '02: Conference on Human Factors in Computing Systems*, 365-368.
2. Benway, J. P. (1998). Banner blindness: the irony of attention grabbing on the World Wide Web. *Proceedings of the Human Factors and Ergonomics Society 42nd Annual Meeting, USA, 1,* 463-467
3. Green, D. M. & Swets, J. A. (1966). Signal Detection Theory and Psychophysics. New York: Wiley.
4. Hillstrom, A. P., & Yantis, S. (1994). Visual motion and attentional capture. *Perception & Psychophysics*, *55*(4), 399-411.
5. Jonides, J., & Yantis, S. (1988). Uniqueness of abrupt visual onset in capturing attention. *Perception & Psychophysics*, *43*(3), 346-354.
6. La Vache Q. R.. (2001). Jocko homo. *Suck.com*. [Online]. Available: http://www.suck.com/daily/2001/04/24/1.html [2001, May 4].
7. Pashler, H. (2001). Involuntary orienting to flashing distractors In delayed search? In Charles L. Folk and Bradley Gibson (Eds.), Attraction, Distraction, and Action: Multiple Perspectives on Attentional Capture. *Advances in psychology*. Elsevier Science.
8. Research sources and findings [Banner ad placement study]. (1998, April 13). *Webreference.com*. [Online]. Available http://www.webreference.com/dev/banners/research.html [2001, March 23].
9. 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 AIS*. *1*(1), 1 – 28.