The Effect of Animated Banner Advertisements on a Visual Search Task

Moira Burke and Anthony J. Hornof
Computer and Information Science Department
University of Oregon, Eugene, OR 97403
{ moira, hornof } @ cs.uoregon.edu

ABSTRACT
Though animated banners are the predominant form of advertising on the Web, much controversy surrounds their efficacy. Users and experts complain, often bitterly, about being distracted by animation and subjectively believe that their performance suffers when flashing objects are on the screen. A visual search experiment was designed to measure both subjective impression of workload and objective task performance in the presence of animated and static banners. Participants reported greater workload, stating that animation increased the frustration and mental demand of the search task. However, participants were not significantly slower or more error-prone in the presence of animated banners.

Keywords
Animation, visual search, banner advertisements, flashing

INTRODUCTION

As processor and Internet connection speeds increase, more designers are choosing to use animation in software and on the Web. Animation provides additional visual information in a screen with limited real estate, and can instruct or assist a user performing a task for the first time. For example, Microsoft Office's animated assistant, "Clippit," occasionally blinks and dances between answering a user's questions (see Figure 1). Giving the traditional help system an animated, lively demeanor could make the program more attractive to users. However, empirical evidence suggests that both the persona and even the mere presence of an animated assistant play a large role in effectiveness and user anxiety [16]. If the animated agent is too intrusive, users may avoid or even resent it.

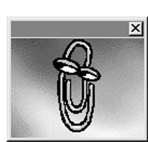

Figure 1. Clippit, the MS Office assistant. Clippit's animated personality is disliked by some users.

Department of CIS Technical Report 02-03

Yet animation can provide meaningful clues in an interface, especially one with a traditionally steep learning curve, such as sophisticated painting software. For example, the selection tool is a common digital painting instrument used to delineate a specific area in an image. What is the best way to represent the selection tool in a toolbar, where each icon must fit in a small area? Photoshop and other popular painting programs use marquee and lasso metaphors for two versions of the selection tool, but these might not be immediately recognizable to a beginner (see Figure 2). Several studies show that animating an icon clarifies its purpose and invites a user to try it [2, 4, 18], though others warn that animation should be haltable by expert users [1]. While animation could assist a novice, experts may be annoyed by gratuitous movement.

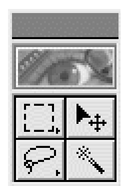

Figure 2. Part of Photoshop's toolbar, including the marquee (upper left) and lasso (lower left) selection tools. Animation might help to explain these tools' purpose.

The problems of digital animation carry over to the analog world, as well. The Federal Highway Administration [5] released a report investigating the safety and aesthetic effects of flashing billboards. In particular, the type of animation, cycle length, and rate of change all have a high impact on traffic safety. The report indicated that aspects of animated road signs could potentially cause driver error, particularly in less-than-optimal conditions, and called for careful consideration in the development of these signs.

Animation in Banner Advertisements

The most prevalent form of advertisement on the Web today is the banner: a 468 x 60 pixel image, often at the top of a webpage. Advertisers value banners for their customizability and automatic log of clickthrough statistics. The Internet Advertising Bureau (IAB), founded in 1996, recommends standard sizes adopted by most businesses to maintain a professional appearance. The IAB lauds the effectiveness of banner ads [15], although a few years after their inception in the mid-nineties, banner clickthrough rates had fallen, and Benway and Lane [3] announced the "banner blindness" phenomenon, in which people missed information prominently displayed in colorful, banner-shaped images.

Thus, advertisers have turned to more drastic methods to attract consumers, and found animation to be a successful option. ZDNet cited in [17] found that animation increased clickthrough rates by at least 15%, and as much as 40%. With the new prevalence of high bandwidth Internet connections, more advanced animations created with Flash, DHTML, and other rich media allow advertisements to interact with users. People can now purchase a book through a Barnes and Noble advertisement, typing their credit card number into a form field embedded in the ad, all without leaving the current site they are exploring.

However, like animated software assistants, animated banner ads also risk irritating people. Daily complaints appear in online newsgroups, questioning the effectiveness of Web ads and criticizing their tactics [12, 19]. One e-journalist called banner ads the Web's "preferred means of exchanging ideas through hyperventilation, screeching and hooting" [12]. People might try to ignore banner ads they dislike. This experiment tests whether people can do this successfully for a particular task.

Animation and Attention Capture

Traditional attention capture research may help to explain how animation grabs attention, if it does at all. Two opposing theories regarding animation have been proposed: (1) automatic attentional capture versus (2) potential capture contingent on other factors. Proponents of the first theory assert that humans involuntarily orient their attention toward moving objects regardless of the objects' importance [8, 10], while those of the second maintain that animation only attracts a user when the animation is relevant to the person's task [6, 14]. Both hypotheses concern exogenous attention orientation, or "shifts of spatial attention that are elicited by salient external events and that occur involuntarily (i.e., without conscious intent)" [6]. Neither hypothesis involves endogenous attention, visual attraction caused by the semantic content of an animated distractor, such as the presumed appeal of a banner with the word "FREE."

Hence, advocates of the automatic attentional capture theory would claim that animated banners automatically capture people's attention. Assuming limited attentional capacity, banner ads would make any concurrent visual task more difficult. Several researchers have investigated the attentional capture of animated icons in a simple visual search task, finding that abrupt visual onset (such as flashing) leads to longer search times, while static discontinuities (such as simple feature changes) do not [20, 11].

Contrarily, the contingent capture hypothesis imposes a strong restriction on animation's purported power: It asserts that animation is only distracting when the motion is relevant to a person's primary task. The contingent involuntary orienting hypothesis states that "involuntary orienting of attention will occur if the event shares the critical property and will not occur if it does not" [6]. This theory is bolstered by a study [14] in which participants actually performed faster in the presence of irrelevant, animated (shimmying) distractors.

Based on the two conflicting hypotheses, one cannot predict the effect of animated banner advertisements. Are people involuntarily drawn to the motion, or does exogenous capture occur only if the banner has relevant material on it? In a related study [13], participants exposed to a ticker that popped up intermittent headlines at the bottom of the screen were distracted from their primary editing task. Would an animated advertisement produce a similar effect?

RESEARCH QUESTION

How do animated banner advertisements affect a simple visual search task? Do users feel a greater sense of workload in the presence of animated banners, and can they perform a task as quickly and accurately as in the presence of similar static banners or no banners? With the ubiquity of flashing advertisements, swirling logos, and blinking icons on the Web, site designers need to know how this motion affects their visitors, both in subjective factors such as frustration and mental demand, and in task performance.

EXPERIMENTAL OVERVIEW

A study was designed to test how animation affected a simple task, analogous to what users might encounter on the Web. The study investigated participants' subjective rating of a visual search task using a well-known workload rating system, NASA Task Load Index (TLX). Search times and errors were also recorded and examined.

In the task, participants were asked to find a particular word hidden amongst distractor words. Meanwhile, two banner advertisements appeared on the screen, some of which were animated. This scenario is common on the Web: People look through lists of links while ads blink around the page.

In a four-person pilot study, the banners were centered just above and below the search area. As links are sometimes arranged in columns and sometimes scattered on a webpage, the pilot study investigated how layout (columnar vs. scattered) affected search outcome. No statistically significant differences were found between the different layouts in the pilot study, so this factor was removed to simplify the main experiment, leaving only the columnar layout.

Pilot study participants found their targets very quickly. Thus, in the main experiment, participants were given four words for which to search, only one of which would appear on the screen. This is much like a real-world web task, in which multiple words could fit someone's search criteria. It was expected that having to maintain four words in memory would lengthen participants' search times, increasing their exposure to the banners, and thus magnifying any distracting effects of the ads.

Since Web users cannot always predict where banners will appear on a page, the banners in the main experiment appeared in different regions of the search area in different trials, instead of being centered above and below the search area as in the pilot study. This allowed the experiment to examine (1) the effect of animated banner ads when they are randomly distributed, from trial to trial, within the same physical space as the targets and distractors, and thus (2) whether interspersing animated ads with relevant content would contribute to the attentional capture of the ads.

METHOD

Subjects

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

Equipment

The visual stimuli were presented on a Sun 18" 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 Microsoft Wheel Mouse.

Experiment software was programmed with Macromedia Director, a multimedia authoring tool commonly used to make interactive CD-ROMs and Shockwave movies for the Web. The software was written in Lingo, Director's scripting language. Director was chosen for its adeptness with both animated and static graphics, and for its similarity to Macromedia Flash, a less robust application with which many modern banner advertisements are made.

DESIGN

Basic Task

For each experimental trial, the participant was instructed to find a particular target and click on it with the mouse. Each trial had two parts: the precue stage and the search stage. In the precue stage, a precue consisting of four randomly-ordered items, one of which was a copy of the target, appeared in a column on the screen. When ready, the participant would click on a button beneath the precue, making the four words disappear, and initiating the search stage. Just below the precue position, a two-dimensional layout of objects would appear. Only one of the objects would match one of the precue words; the other visible objects were distractor words and banners, designed to make the search harder. The participant would search for the target item and click on it.

Target and Distractors

Each experimental trial contained one target object and 19 similar distractor objects. Targets and distractors were four-letter words written in 18-point Helvetica type enclosed in a 2.4 x 0.7 cm rectangle. Six hundred seventy-two four-letter words were used, gathered from an automated, exhaustive search of the first two levels of text on Wired.com. Words were filtered for appropriateness and limited to one or two syllables to facilitate silent vocal repetition during the search process.

Layout

The target and distractors were arranged in three columns of eight rows (see Figure 3). Two banners were randomly placed so that each covered two columns of one row, leaving nineteen spots for the target and distractor words. Banners were placed in random rows so that targets appeared in all 24 positions across the experiment.

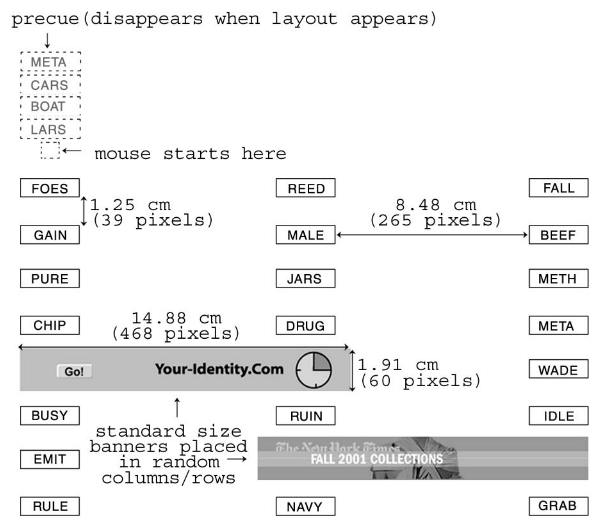

Figure 3. Layout of target and distractors, drawn to scale. The precue is also shown.

Banners

Figure 4 shows the four banner types used in the experiment: (a) blank banners, (b) animated commercial banners, (c) static commercial banners, and (d) flashing text banners. In the case of blank banners, distractor words that would be "covered" by the banners were simply removed.

A selection of one hundred ten animated commercial banners from popular news websites, search engines, and portals, including the New York Times website, AltaVista, Yahoo!, AOL, and CNN.com was chosen. Static commercial banners consisted of a representative frame from each animated commercial banner. Finally, to incorporate an extreme case of animation, flashing text banners were created: each had a bright cyan background with large black text that alternated flashing on the left and right halves of the banner every 150 msec. The cycles for the two banners were offset by 80 msec, making them change asynchronously. Text for these banners rotated through the target and distractor words for any given trial. Thus, their content was somewhat relevant to the task, potentially increasing their ability to capture attention.

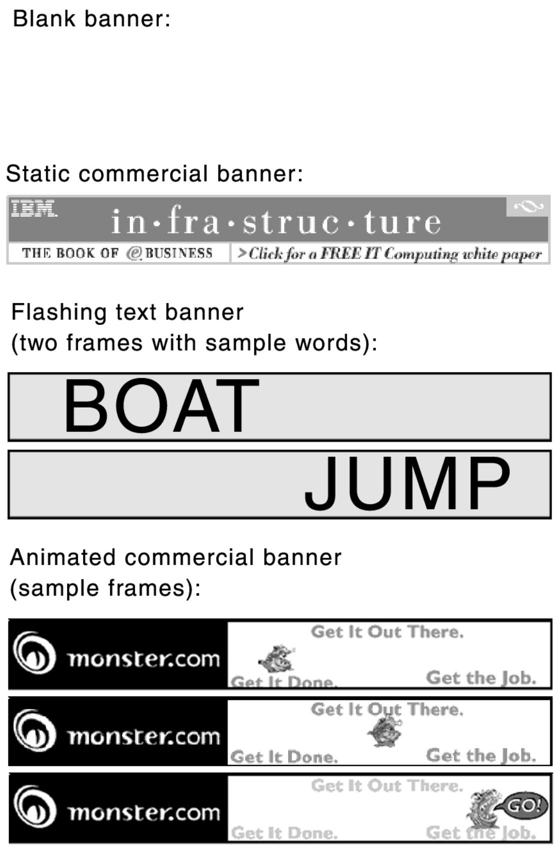

Figure 4. Samples of the four banner types used in the experiment. Sample frames from the animated banners are shown. All banners were in color, and the flashing text banner had black text on a cyan background.

Point Completion Deadline

To isolate search time and selection time, the point completion deadline (PCD) designed by Hornof [9] was used in this experiment. In short, participants were instructed to hold the mouse stationary until they found the target. Once they moved the mouse more than five pixels in any direction, they had a small amount of time to click on the target. If participants took substantially longer than Fitts' Law would predict to be necessary for the mouse point, a buzzer sounded and the trial was recorded as an error. Thus, both the search time and selection time were separately recorded. Participants were given as long as they liked to practice with the PCD before the experiment so that they could grow accustomed to not moving the mouse until they found the target, and thus error rates were minimized during the reported trials.

Incentives and Penalties

Incentives and penalties were weighted to encourage fast but mostly accurate responses. Every participant was paid a minimum of $10, but each had the opportunity to earn an additional bonus of up to $5 based on speed and accuracy in the trials. Each trial had a potential bonus of seven cents, with one cent deducted for every second that elapsed on the clock until the participant clicked on the target. When participants hit the target, they would hear a 150 msec chime and earn the trial bonus. If participants clicked on anything other than the target, or if the PCD expired, a 350 msec buzzer sounded and a five-cent penalty was deducted from their earnings. In the case of a PCD expiration, an alert box would also appear reminding participants not to move the mouse until they had seen the target. The software displayed cumulative earnings at the end of each block of trials.

Blocking

Each banner type and target position combination (a total of 4x24=96) was presented to each participant once. The set of all combinations was divided into four equal-length blocks to give participants a chance to rest between blocks. Trials were randomized in each block, and blocks were counterbalanced across participants with a randomized Latin square.

Whenever a participant made an error, an additional trial with the same banner type and target position combination was added to the end of the block, and the remaining trials for that block were reshuffled. This way participants completed a correct trial for every combination of banner type and target position. Two percent of all combinations were not presented to every participant due to subtleties in the experimental software.

To gather baseline pointing times, an additional target-only block was added. In this block, the target appeared with no distractors. Each of the 24 target positions appeared three times to give a mean search and selection time for each position. Participants began target-only trials with two cents, rather than seven cents, since the task was much easier and there were three times as many trials as in other blocks. The target-only block was presented after the first two blocks for all participants. These results will not be reported here.

PROCEDURE

Search Task

Participants began the experiment by filling out a preliminary questionnaire regarding their usage of computers and the Web. Next, they read instructions about the precue, target, point completion deadline, bonuses, and block structure. The lights were dimmed and participants positioned 56 cm from the screen with the precue at eye level. This eye-to-screen distance was remeasured before each block. Approximately half of the participants chose to use a mousepad, although the optical mouse did not require one.

Before the timed portion of the experiment, participants were given unlimited practice trials from the first block to grow accustomed to the point completion deadline. When they were ready, the software was restarted and data collection began. At the start of each block, participants were given five additional practice trials. During these trials, the words "Practice trial" appeared above the precue.

For each trial, the four precues appeared alone on the screen. Participants took as long as they needed to study the words. When ready, participants clicked the box beneath the precue. The precue disappeared, and the target, distractor words, and banners appeared. When the participant clicked on a target or distractor, the item's colors inverted briefly, and either a chime or buzzer sounded, as previously described. At the end of each block, a summary screen showed how much money the participant had earned so far.

NASA Task Load Index (TLX) Ratings

After the timed experiment, participants were asked to subjectively evaluate their experience with the different conditions. The Task Load Index (TLX) system developed by NASA [7] was used to allow comparison of subjective ratings across participants. TLX measures workload, defined by the following factors: mental demand, physical demand, temporal demand, impression of performance, effort, and frustration. Participants rated each factor on a visual scale between one and 100 and then performed a pairwise weighting comparison, indicating for each pair which factor contributed more to the sense of workload. A combination of these values reveals the relative importance of each factor, providing a reasonable metric with which many participants' responses can be compared.

Before being asked to record their subjective ratings for each banner type, participants were shown blocks with each banner type again, in randomized Latin square order. All four banner types were rated along with the target-only condition. Participants ran a few trials in a given block to refresh their memory, and then gave TLX ratings and weights for the condition in that block.

In addition to the TLX evaluations, participants were also given an exit interview, conducted verbally with the experimenter. Participants described the level of difficulty of different parts of the experiment, ranked the different banner types for distraction, and described any strategies they used to perform well.

RESULTS

TLX Workload Scores

Overall workload

As indicated by the TLX workload scores, participants reported a much greater sense of overall workload in the presence of the flashing text banners (see Figure 5). A repeated measures ANOVA revealed a significant difference in overall workload between the banner types, $F(3, 36) = 6.52$, $p<0.001$. A Fisher's PLSD post hoc test showed that flashing text banners were reportedly more workload-intensive than the other banner types. Though not significantly greater, the animated commercial banners had a somewhat greater overall workload than the static and blank banners. The static and blank banners ranked approximately the same in terms of overall workload.

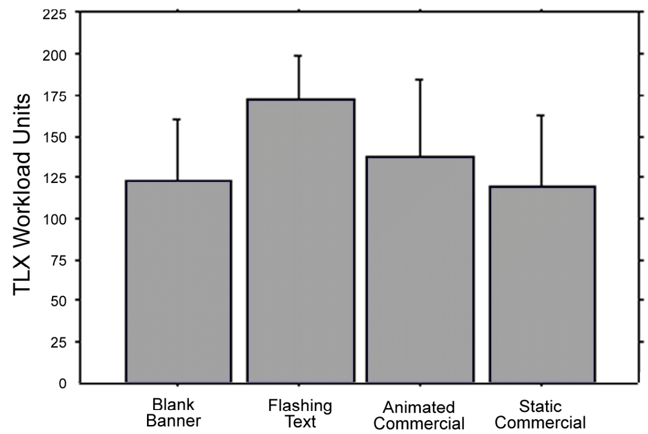

Figure 5. Reported overall workload for the four banner types, averaged across participants and target locations. Bars indicate one standard deviation from the mean of the twelve participant means.

Factors that contributed to workload

Participants found the flashing banners more frustrating and mentally demanding than the other banner types, $F(3,32) = 3.50$ and 1.62, respectively, $p<0.05$. Participants reported marginally ($p<0.1$) greater amounts of temporal demand and effort required to process the flashing banners, as well. Figure 6 shows how much each of the six TLX factors contributed to the overall workload for each banner type.

Search Time and Error Rates

Table 1 shows the mean search time for each banner type. Error and practice trials are excluded from analyses unless specifically noted. Results from a repeated measures fail to show a significant difference in search time caused by banner type. Statistical power was low, however, (approximately 0.45), so only a large effect could have been detected.

Table 1. Search times for each banner type, averaged across target position and participant. The standard deviation of the twelve participant mean times for each banner type is also shown.

Banner Type	Mean Search Time (in msec)	Std. Deviation (in msec)
Blank	5831.24	1674.61
Flashing Text	5234.39	1115.65
Animated Commercial	4795.10	1010.12
Static Commercial	5154.68	1237.51

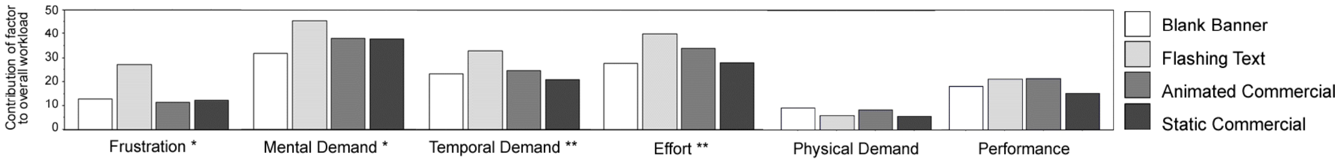

Figure 6. TLX Workload separated by contributing factors and banner types. Mental demand and effort were the greatest contributors for all banner types. Participants ranked the flashing text banners higher for four of the six factors.

* indicates a factor that was significantly higher for the flashing text banner

** indicates a factor that was marginally significantly higher for the flashing text banner

Two significant effects were found in the search time data. First, a position effect is evident: Participants found targets in the upper-left positions much more quickly than those to the lower-right, $F(23,72)=5.08$, $p<0.0001$ (see Figure 7).

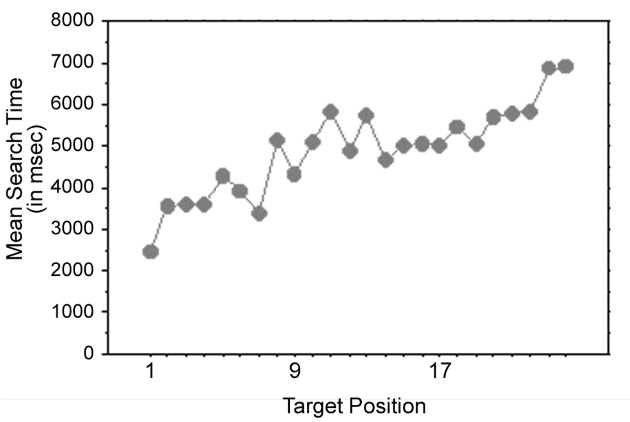

Figure 7. Mean search times split by target position show a strong position effect. Target positions are numbered 1 to 24, starting at the upper left and counting down each column. Thus, the numbers 1, 9, and 17 (shown on the x-axis) are the indices for the topmost item in each column.

The second significant effect occurred when a target was sandwiched between two flashing text banners. In this case, the target took an average of 75% (4.3 seconds) longer to find than if the flashing text banners were placed elsewhere, $F(1, 298)=7.0$, $p<0.01$. 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.

Error rates are shown in Table 2. Participants committed two kinds of errors: *misses*, in which they clicked on something other than the target, and *timeouts*, in which the point completion deadline expired. The percentage of misses was not significantly different for the different banner types; the percentage of timeouts was also approximately the same for the different banner types.

Table 2. Error rate for each banner type, averaged across target position and participant.

Banner Type	Misses	Timeouts
Blank	0.06	0.05
Flashing Text	0.06	0.10
Animated Commercial	0.07	0.10
Static Commercial	0.06	0.09

DISCUSSION

TLX Workload Scores

Concurrent with popular opinion, participants did find the most extreme type of animation, the flashing text banners, more frustrating and mentally demanding than the other types of banners. Additionally, participants felt they had to expend slightly more effort, and that there was greater time pressure when the flashing text banners were on the screen.

Interestingly, the same magnitude of results was not found for the animated banners culled from commercial websites; in most cases, the animated commercial banners were only slightly (nonsignificantly) more workload-intensive. These results suggest that the current crop of professional animated banners on the Web are somewhat difficult to process mentally, but have the potential to be extremely workload-intensive (and subsequently unpleasant) if the animation were more extreme, like the flashing text banners. More ostentatious advertisements are found on the Web, but the banners used in this experiment came from popular professional sites, and thus comprise a representative sample of ads accepted by influential members of the commercial Web community.

For all of the TLX workload factors, the animated commercial banner ranked equally or higher than its static counterpart. Though the difference was not statistically significant, it still shows a strong trend. Since the content of the animated and static commercial banners was virtually the same, the animation must have contributed to the increased workload. Thus, advertisers should be wary of using animation, since in this case it added to people's feelings of workload without producing a noticeable increase in search time. People were annoyed by the animation even without spending extra time looking at it.

Since the flashing text banners contained four-letter words from the list of targets and distractors, the contingent involuntary orienting hypothesis [6] may have led to the increased workload scores. Participants might have noticed the flashing words and recognized some of them from the precue list or distractor set, and thus needed to use more mental capacity to filter them out. Perhaps if the flashing text on the banners had been the same size and/or color as the target and distractor set, search times might have been even greater.

Search Times and Error Rates

That participants were able to find a target just as quickly in the presence of animated banners as with static or blank banners suggests that animation does not necessarily capture attention to the point of impeding a simple visual search task.

This experiment only delves into one particular visual search task: people looking for a particular word they know will appear. The experiment does not consider the effect of peripheral animation on a more demanding search task or cognitive task such as reading. Quite frequently, websites list their internal links at the top of the page, just below an animated advertisement, which presents a task similar to that performed in the experiment, but different results could occur on a reading-intensive site. When people visually scan the results of a Google search query, for example, they must interpret longer phrases taken out of context. In that circumstance, perhaps animated banners would be more distracting.

Also, this experiment intentionally ignores the appeal of banner content; participants had no incentive to look at the banners, and they were encouraged to complete their search task quickly. Web surfers are often lured by the content of ads rather than the medium, especially when ads are carefully aimed at a target audience.

Interestingly, one experimental result calls into question previous advertising tips suggested by WebWeek, cited in [17]. In 1996, the online magazine recommended that site designers "bracket important content between two banners." However, this experiment shows that content sandwiched between two flashing text banners takes significantly longer to find. Perhaps WebWeek's suggestion was prudent in 1996, but five years of animated advertisements have since dulled consumers' attention.

Post-Experiment Interview

Participants discussed their search strategies and impressions in the post-experiment interview. In general, participants reported being able to "tune out" the banners, although some found the flashing text and brightly colored banners difficult to ignore. One participant even admitted to intentionally clicking on 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 ranked 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 found the blank banners easiest, but several others preferred having static commercial banners to break up the visual field. According to one participant, blank banners led to "too much of the same," making the search more difficult. On average, participants did have the longest search time when the (invisible) blank banners were on the screen, though, again, the differences were not significant.

The interview also revealed the robustness of the PCD and the various strategies that participants used to perform well. When asked if they were able to move the mouse before finding the target word, participants reported being buzzed nearly every time. The only exception occurred for two participants when they began moving the mouse to a word they thought was the target, and then realized that an adjacent word was actually the correct target, which they were able to reach in time.

Most participants reported searching the columnar layout from left to right, scanning either down or up each column, although some performed a preliminary overall inspection around the center of the screen first.

CONCLUSION

Though many people consider advertising a necessary evil to keep the Web cost-free, much thought should be given to the implementation of banner ads. This research shows that for a simple visual search task, animation does not catch peoples' attention to the extent that they were slowed in their task. While further investigation is warranted, this study reveals that advertisers should not automatically equate animation with attentional capture, especially considering the increased workload reported by participants. Advertisers using animated banners should be wary of irritating their audience with a potentially ineffective technique.

Though animation did not affect search time in this experiment, it might play a larger role concurrent with a more difficult task. Reading is a common Web task, and it requires much greater mental processing than simple visual scanning. Even filling out an online form entails more thought; investigating the effect of animated banner ads on either of these two tasks would be worthwhile to the Web community.

This study bridges traditional attention-capture and Web marketing research. The Web imposes myriad factors on banner effectiveness, each of which could be investigated separately, in traditional human factors studies, or in a realistic Web setting, as is often the case in marketing surveys. This experiment probed a few individual features of the Web; subsequent studies will add more to the animation discussion.

ACKNOWLEDGMENTS

This work was supported in part by the Office of Naval Research through Grant N00014-01-10548 to the University of Oregon, Anthony J. Hornof, principal investigator.

REFERENCES

1. Alpert, S.R. (1991). Self-describing animated icons for human-computer interaction: a research note. *Behaviour and Information Technology, 10*(2), 149-152.
2. Baeker, R., Small, I., & Mander, R. (1991). Bringing icons to life. *Proceedings of CHI '91: Conference on Human Factors in Computing Systems*, 1-6.
3. Benway, J.P., & Lane, D.M. (1998, December). Banner Blindness: Web Searchers Often Miss "Obvious" Links. *Internetworking* (1.3). Retrieved from http://www.internettg.org/newsletter/dec98/banner_blindness.html
4. Dormann, C. (1994). The design of animated signs as help. *Proceedings of CHI '94 Conference Companion on Human Factors in Computing Systems*, 137-138.
5. Federal Highway Administration. (1980). *Safety and environmental design considerations in the use of commercial electronic variable-message signage* (FHWA Rep. No. RD-80/051). Retrieved from http://www.scenicflorida.org/cevms/chfhward80051a.html
6. Folk, C.L., Remington, R.W., & Johnston, J.C. (1992). Involuntary Covert Orienting Is Contingent on Attentional Control Settings. *Journal of Experimental Psychology: Human Perception and Performance, 18(4),* 1030-1044.
7. Hart, S.G., & Staveland, L.E. (1988). Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. *Human Mental Workload*, 139-183.
8. Hillstrom, A.P., & Yantis, S. (1994). Visual motion and attentional capture. *Perception & Psychophysics, 5(4)*, 399-411.
9. Hornof, A.J. (in press). Visual search and mouse pointing in labeled versus unlabeled two-dimensional visual hierarchies. *ACM Transactions on Computer-Human Interaction*.
10. Jonides, J. (1981). Voluntary versus automatic control over the mind's eye's movement. *Attention and Performance*, 9, 197-203.
11. Jonides, J. & Yantis, S. (1988). Uniqueness of abrupt visual onset in capturing attention. *Perception & Psychophysics, 43*(3), 346-354.
12. La Vache Qui Rit. (2001). *Jocko homo*. Retrieved May 4, 2001, from http://www.suck.com/daily/2001/04/24/1.html
13. Maglio, P.P., & Campbell, C. S. (2000). Tradeoffs in displaying peripheral information. *Proceedings of CHI 2000: Conference on Human Factors in Computing Systems*, 241-248.
14. Pashler, H. (under review). Involuntary orienting to flashing distractors in delayed search? *Journal of Experimental Psychology: Human Perception and Performance.*
15. PricewaterhouseCoopers. (1999). *IAB Internet Advertising Report [Executive Summary, 1999 Third-Quarter Results]*. Retrieved from http://www.iab.net/adrevenue/PwC_IAB99Q3.PDF
16. Rickenberg, R., & Reeves, B. (2000). The Effects of Animated Characters on Anxiety, Task Performance, and Evaluations of User Interfaces. *CHI 2000 Letters, 2*(1), 49-56.
17. *Research sources and findings [Banner ad placement study]*. (1988, April 13). Retrieved from http://www.webreference.com/dev/banners/research.html
18. Rist, T., André, E., & Müller, J. (1997). Adding Animated Presentation Agents to the Interface. *Proceedings of IUI '97: Conference on Intelligent User Interfaces*, 79-86.
19. St. Huck. (1999). *Add it up: does that banner yet wave?* Retrieved May 4, 2001, from http://www.suck.com/daily/1999/05/21
20. Yantis, S., & Johnson, D.N. (1990). Mechanisms of attentional priority. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 812-825.