

INTERNET ADVERTISING: IS ANYBODY

WATCHING?

XAVIER DRÈZE is an assistant professor, Wharton School, The University of Pennsylvania, Philadelphia, Pennsylvania; e-mail: xdreze@wharton.upenn.edu

FRANÇOIS-XAVIER HUSSHERR is director of Internet and new media activities at Mediametrie, Paris, France; e-mail: fxhusherr@mediametrie.fr

The authors would like to thank Voilà and France Telecom for funding this research and Philippe Taupin for his help with the Eye-Tracker device.

Xavier Drèze
François-Xavier Husherr



ABSTRACT

Click-through rates are still the de facto measure of Internet advertising effectiveness. Unfortunately, click-through rates have plummeted. This decline prompts two critical questions: (1) Why do banner ads seem to be ineffective and (2) what can advertisers do to improve their effectiveness? To address these questions, we utilized an eye-tracking device to investigate online surfers' attention to online advertising. Then we conducted a large-scale survey of Internet users' recall, recognition, and awareness of banner advertising. Our research suggests that the reason why click-through rates are low is that surfers actually avoid looking at banner ads during their online activities. This implies that the larger part of a surfer's processing of banners will probably be done at the pre-attentive level. If such is the case, click-through rate is an ineffective measure of banner ad performance. Our research also shows that banner ads do have an impact on traditional memory-

© 2003 Wiley Periodicals, Inc. and
Direct Marketing Educational Foundation, Inc.



JOURNAL OF INTERACTIVE MARKETING
VOLUME 17 / NUMBER 4 / AUTUMN 2003
Published online in Wiley InterScience (www.interscience.wiley.com).
DOI: 10.1002/dir.10063

based measure of effectiveness. Thus, we claim that advertisers should rely more on traditional brand equity measures such as brand awareness and advertising recall. Using such measures, we show that repetition affects unaided advertising recall, brand recognition, and brand awareness and that a banner's message influences both aided advertising recall and brand recognition.

INTRODUCTION

Advertisers were one of the early proponents of the Internet, embracing its dual promise of global reach and one-to-one targeting. This should not come as a surprise as advertisers have long used every conceivable vehicle to display their messages in front of the gazing eyes of potential customers, be it magazines, television, or racecars. However, after a promising start, the burst of the Internet bubble has forced many companies to curb their plans for further expansion of their online advertising spending. Industry experts nevertheless believe that this setback will be short-lived and project a resurgence of online advertising in 2004 (Forrester, 2001; Internet Advertising Bureau, 2002; The Wall Street Journal Online, 2003) with a full recovery by 2007 (Forrester, 2002).

A possible obstacle to this resurgence of online advertising is its apparent lack of effectiveness. Indeed, the most widely used measure of online advertising effectiveness is the percentage of the total number of ad exposures that induce a surfer to actually click on a banner in response to an advertised message (Forrester, 2001, 2002). This measure is known as the click-through rate (Novak & Hoffman, 1997). Click-through rates started in 1996 at around 7%. However, they have declined to around 0.7% in 2002 (DoubleClick, 2003). This is problematic because advertisers typically do not knowingly allocate budgets to media that are not effective.

Past research on advertising effectiveness cautions us regarding click-through as a valid effectiveness measure. Briggs and Hollis (1997) have shown using Milward Brown's Brand Dynam-

ics™ system (Dyson, Farr, & Hollis, 1996), that banner ads can have an impact on consumers' attitudes toward a brand independent of click-through.

Given the doubts one may have concerning online advertising effectiveness, we believe that the following two questions must be answered before marketers can pour money back into the online realm: (1) Why do banner ads seem to be ineffective and (2) what can advertisers do to improve their effectiveness? The purpose of this paper is to answer these two questions. We first show that banner ads will most likely operate at the pre-attentive processing level (Shapiro, MacInnis, & Heckler 1997). Thus, traditional effectiveness measures are more appropriate than click-through rates. We then use memory-based measures to study some of the factors that might impact banner ad effectiveness.

The remainder of the paper is organized as follows. The first section discusses the results of a study that utilized an eye-tracking device to determine whether Web surfers see banner ads and which factors increase or decrease the probability that a banner ad is seen. We use the results from this first study to generate hypotheses about the characteristics of banner ads that might increase or decrease viewers' attention. The second section relates the results of the follow-up study that tested the hypotheses generated in the first study on a broader sample of Web surfers (807 respondents). The study explored the effects of Internet advertising on recall, recognition, and awareness. We then consider the results of both studies and discuss their managerial relevance. We close with concluding remarks, a discussion of the limitations of our methodology and results, and directions for future research.

STUDY I: EYE-TRACKING

The Internet differs from traditional media in at least one significant way. When an advertiser uses television or radio to deliver his messages, he preempts the program being broadcasted (e.g., a sitcom or song) and uses all the bandwidth of the medium to transmit his message (see Drèze & Zufryden, 2000). This means that,

by default, the viewer or listener is paying attention to the advertisers, and the message is only interrupted if the listener zaps away. Zapping, however, is quite infrequent. Siddarth (2002) reports zapping rates for commercials of less than 3%. By contrast, online banner ads share their bandwidth with other elements of the pages in which they are being displayed. A banner ad typically occupies less than 10% of the area of a Web page on a standard VGA computer screen (640×480 pixels). Therefore, the attention of the Web surfer is generally focused on other elements of the page. The task of the banner ad is thus first to grab a surfer's attention and second to induce the surfer to click on the ad. If surfers never look at a banner, they cannot click on it!

Shared bandwidth might explain why click-through rates are low, but not why they are declining. There is some evidence that some online surfers dislike banner ads (Bass, 1999). One can thus hypothesize that, as surfers gain more familiarity with the medium, they learn to differentiate informational content from advertising. Ultimately, this would give them the ability to disregard banner ads. Janiszewski (1998) shows that peripheral vision allows subjects to recognize objects that are located outside their focal point of attention (as measured by the eye-tracking device). This ability, coupled with the fact that most banner ads have the same shape (468×60 pixels) provides Web surfers with the ability to train themselves into recognizing banner ads for what they are without having to actually focus on them. Both of these explanations assume that surfers learn over time and develop strategies to avoid devoting attention to advertising.

Given this possible learning and avoidance behavior, we start our investigation by measuring the extent to which surfers pay attention to banner ads. We begin by formulating the following two hypotheses:

H₁: Internet users avoid looking at banner ads.

H₂: The more time users have spent on the Internet; the less they pay attention to banner ads.

To test these hypotheses, we asked a group of subjects to look at various Web pages while hooked up to an eye-tracking device that records their eye movements and fixations. The use of eye-tracking devices in marketing studies is not new. Russo and Leclerc (1994) studied in-store brand choice, Fischer et al. (1989) studied warning labels on tobacco ads, Janiszewski (1998) looked at exploratory search behavior with catalogs, Kroeber-Riel (1979) investigated the effect of arousal on advertising copy processing, and Lohse (1997) studied Yellow Page advertising. Most of the eye-tracking studies, including ours, use Pupil Center/Corneal Reflection (PCCR) monitoring devices to track the eye movement of their subjects.

Study Design

Our study was conducted using information portals as a background. The cover study was an ergonomic research on the design for one of the largest French portals: *Voilà* (www.voila.fr). The subjects were asked to perform five searches using three portals: *Voilà*, an alternate layout for *Voilà* (henceforth called *Voilà Bis*), and *Voilà*'s largest competitor. Three of the searches related to general topics (e.g., find information about "Le Louvre"), the other two related to individuals (e.g., find the phone number of "Jean Dupont"). Each of the three general-topic searches was made using a different portal. The two other searches were made with *Voilà* and *Voilà Bis*.

To accomplish the task they were assigned, the subjects would click on a link or enter a search string. For the three general-topic searches, this would prompt the display of an answer page (see Figure 1 for the answer page to *Voilà*) containing the information they were asked to look for. They were then asked to indicate with the mouse where the information they were looking for was located. The answer pages were designed to match the look and feel of the question page. In other words, the *Voilà* search page led to a *Voilà* answer page of similar design; the *Voilà Bis* search page led to a *Voilà Bis* answer page, and likewise for the competing portal.

The order in which the pages were shown was

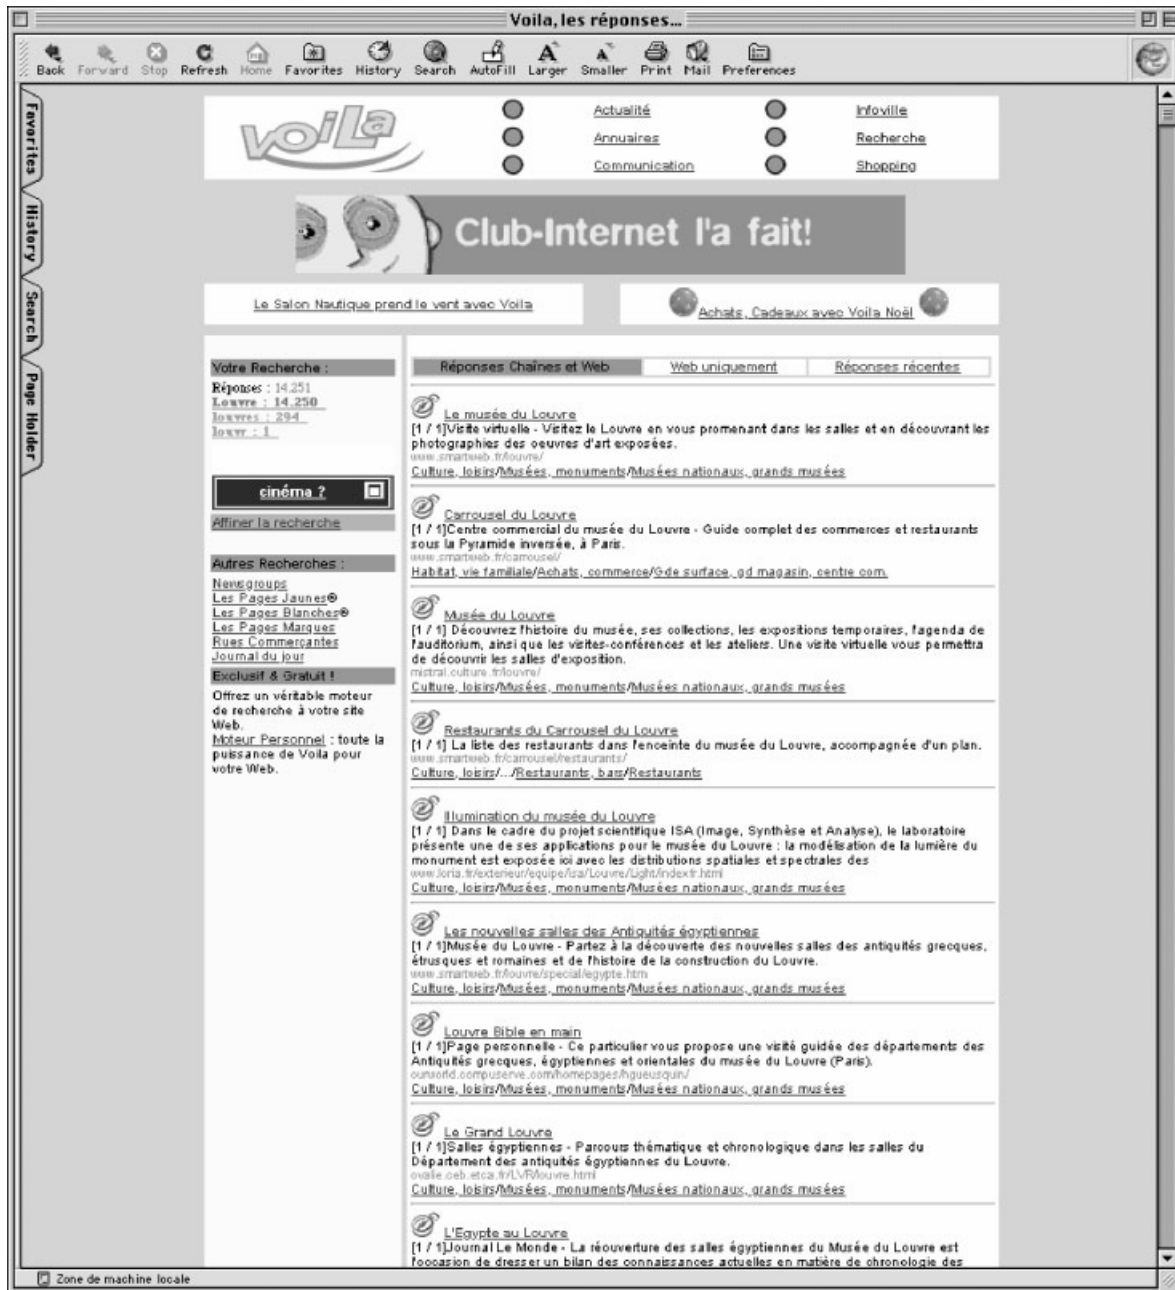


FIGURE 1
Voilà Answer Page

rotated across subjects with the only restriction being that the three search/answer page pairs were kept unbroken. Eight banner ads were integrated within the design of the eight Web pages (see Figure 1 for the “Club Internet” banner of *Voilà Answer*). At no time before or during the search task on these eight pages was any reference to banners ads made.

We collected our data in two steps. First, we used an eye-tracking device to collect eye movements and fixations during the experiment. Second, we asked our subjects to fill out a short survey after completing their assigned task on the eight Web pages. The survey asked questions about their Internet savvy, the experimental process (e.g., did they encounter any

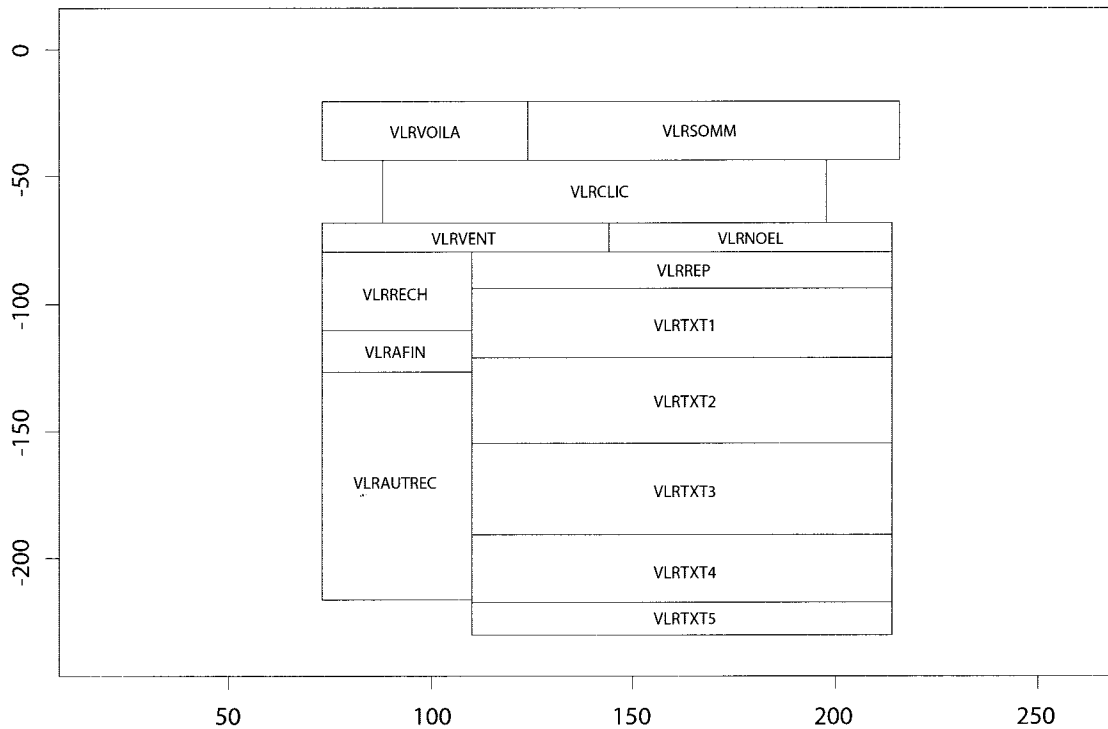


FIGURE 2
Voilà Answer Zone Definition

stress?), their preferences regarding the various pages to which they were exposed, and a series of questions regarding the banner ads they saw (e.g., do you remember seeing any banner ads?).

To simplify the analysis of the eye-tracking data, each page was dissected in a series of mutually exclusive rectangular zones. One zone was created for each paragraph of text, banner ad, or graphical element of the page (see Figure 2 for the zone definition of *Voilà Answer*). The eye-fixation data were then coded by zone. Hence, for every subject we have a list of each zone that they fixated on during the experiment and for how long. Of the 60 subjects that were recruited for the experiment, 11 had to be eliminated because they suffered from heavy nystagmus or because the calibration of the eye-tracking device could not be performed satisfactorily on them. This left us with 49 usable subjects. The subjects were selected through a street intercept in the center of Paris and paid €15 to participate in the study.

Analysis

The primary goal of this first experiment was to measure the extent to which surfers actually look at the banner ads that are embedded within the Web pages. Each of our 49 subjects was exposed to eight banner ads (one per page). Looking at the zones that were focused on by our subjects, we find that every subject looked at one or more banner ads (i.e., nobody managed to avoid every ad). On average they looked at 3.96 banners during the experiment, which yields a probability of being seen of 0.49 for each individual banner. This probability is low relative to other media such as television (> 90%, per Siddarth, 2002) or Yellow Page ads (89% for small display ads, 93% for large display ads, per Lohse, 1997).

To test Hypotheses 1 and 2, we created a data file that lists all of the zones that a subject might focus on (i.e., 111 zones × 49 subjects). Each zone is associated with variables describing their location, shape, and content (see Appendix A for a description of the zone description vari-

TABLE 1
Logit Regression (Zone Attractiveness)

<i>Variable</i>	<i>Coefficient</i>	<i>Wald χ^2</i>	<i>Pr > χ^2</i>
Intercept	-3.1947	30.22	0.0001
Test			
AD	-0.1888	5.10	0.0240
Expertise	-0.0633	2.50	0.1136
AD \times Expertise	-0.0554	0.54	0.4639
Location			
X	0.0203	8.24	0.0041
Y	0.0211	32.73	0.0001
X ²	-0.0001	14.52	0.0001
Y ²	-0.0001	59.24	0.0001
Distance	-0.0380	326.27	0.0001
Shape			
Area	0.0005	348.96	0.0001
Zone Orientation	0.0482	0.37	0.5437
Page Orientation	0.0597	1.00	0.3190
Orientation Int.	0.1135	3.68	0.0550
Content			
IMAGE	1.2067	91.52	0.0001
PICTURE	-3.9477	236.82	0.0001
CONTRAST	0.2813	12.06	0.0005
TEXT	1.8351	150.20	0.0001
INPUT	0.1751	2.23	0.1353
Control			
# of Zones/Page	0.0083	0.15	0.7029
Age	0.0595	2.89	0.0892
Gender	0.0531	1.60	0.2013
Regression		1234.03	0.0001

ables). We then ran a logit regression using as the dependent variable an indicator that was set to 1 if the subject fixated on the zone, and 0 otherwise. As test variables we used an *Ad* dummy (1 if the zone is a banner ad, 0 otherwise), an *Expert* dummy (1 if the subjects have been on the Internet at least 25 times, 0 otherwise), and an *interaction* term. To control for page layout as well as possible differences across gender or age, we use the variables described in Appendix A as control variables along with two demographic control variables (Gender and Age—over or under 40 years old).

The results shown in Table 1 are very revealing. As one could have expected, a zone's loca-

tion and size are important. The positive coefficient on *Area* shows that the bigger a zone is, the more likely it is to capture subjects' attention. The significant interaction between the page and the zone interaction shows that the page layout is important. Similarly, the zone's content is important as evidenced by the significant content dummies.

The negative coefficient on the *Ad* dummy provides support for Hypothesis 1. It indicates that viewers avoid looking at ads. It also indicates that they are able to recognize that an item is an ad without having to look at it directly. Although the *Expert* dummy is marginally significant ($p = 0.11$), the interaction term be-

TABLE 2
Differences Among Groups of Surfers

	<i>Number of Fixes</i>		<i>Number of Zones</i>		<i>Fix Time</i>	
	<i>Estimate</i>	<i>p-Value</i>	<i>Estimate</i>	<i>p-Value</i>	<i>Estimate</i>	<i>p-Value</i>
Intercept	26.59	0.0001	6.99	0.0001	12.01	0.0001
Expert	-6.03	0.0006	-0.67	0.0162	-2.52	0.0002
Old	5.36	0.0018	0.21	0.4536	2.65	0.0001
Female	-2.32	0.2570	-0.59	0.0753	-0.86	0.2822
Voilà	-0.91	0.7742	0.28	0.5851	-0.23	0.8507
Voilà Answer	4.73	0.1363	1.59	0.0019	2.22	0.0742
Yahoo!	4.62	0.1413	-0.49	0.3309	0.63	0.6078
Yahoo! Answer	12.18	0.0001	-2.27	0.0001	4.39	0.0004
Voilà Bis	4.13	0.1912	0.43	0.3872	1.65	0.1806
Voilà Bis Answer	10.20	0.0014	2.26	0.0001	4.20	0.0008
Voilà Annuaire	6.48	0.0436	2.78	0.0001	0.95	0.4481
<i>N</i>	378		378		378	
<i>R</i> ²	0.11		0.28		0.13	

tween expertise and banner is not significant ($p = 0.46$). Hence, we do not find support for Hypothesis 2.

Although we did not find support for Hypothesis 2, it is still interesting to contrast the behavior of our various demographic groups. During the experiment, as well as throughout our analysis, we found significant differences in behavior between novices and experts as well as between young and older surfers. To illustrate these differences, we ran a series of regressions on the number of fixes, number of zones looked at, and time spent during fixes across these groups. For each page looked at by each respondent, we regressed the three dependent variables against an expertise dummy, an age dummy, and a gender dummy, as well as seven control variables to account for the differences across pages. As the results in Table 2 show, experts tend to process each page by making fewer fixes, looking at fewer zones, and spending less time than novices. An illustration of this can be seen in Figure 3, where we show an example of eye movements for both an expert and a novice. The expert clearly makes fewer fixations than the novices. Further, the eye movement seems more systematic with less back and forth movement. Older people look at the same number of zones as

young people, but it takes them longer and they fixate on a larger number of positions. Finally, males and females seem to behave similarly.

As part of the debriefing questionnaire, we asked our subjects if they remembered seeing any banner ads. Only 46.9% of the subjects indicated they did. After asking them if they remembered seeing any ad, we showed our subjects four banners and asked them if they recalled seeing the ads during the test. Two of the ads were fake ads that had not been part of the test; the other two were real. We did not find significant differences in recognition level between the fake and the real ads ($m = 23.5\%$ vs. 18.4, respectively, $p = 0.38$). The number of false positives we encountered is similar to those reported in other studies. Janiszewski (1990a) reported 21% of false positives (119 subjects, 5 different ads). The French division of the Internet Advertising Bureau and SOFRES (IAB France/SOFRES, 1999) reports false positive levels of 17% (6,872 subjects, 14 ads).

Take away from Study 1

Our first study is very revealing. It provides us with an answer to the first question motivating this study (why are banner ads not effective?)

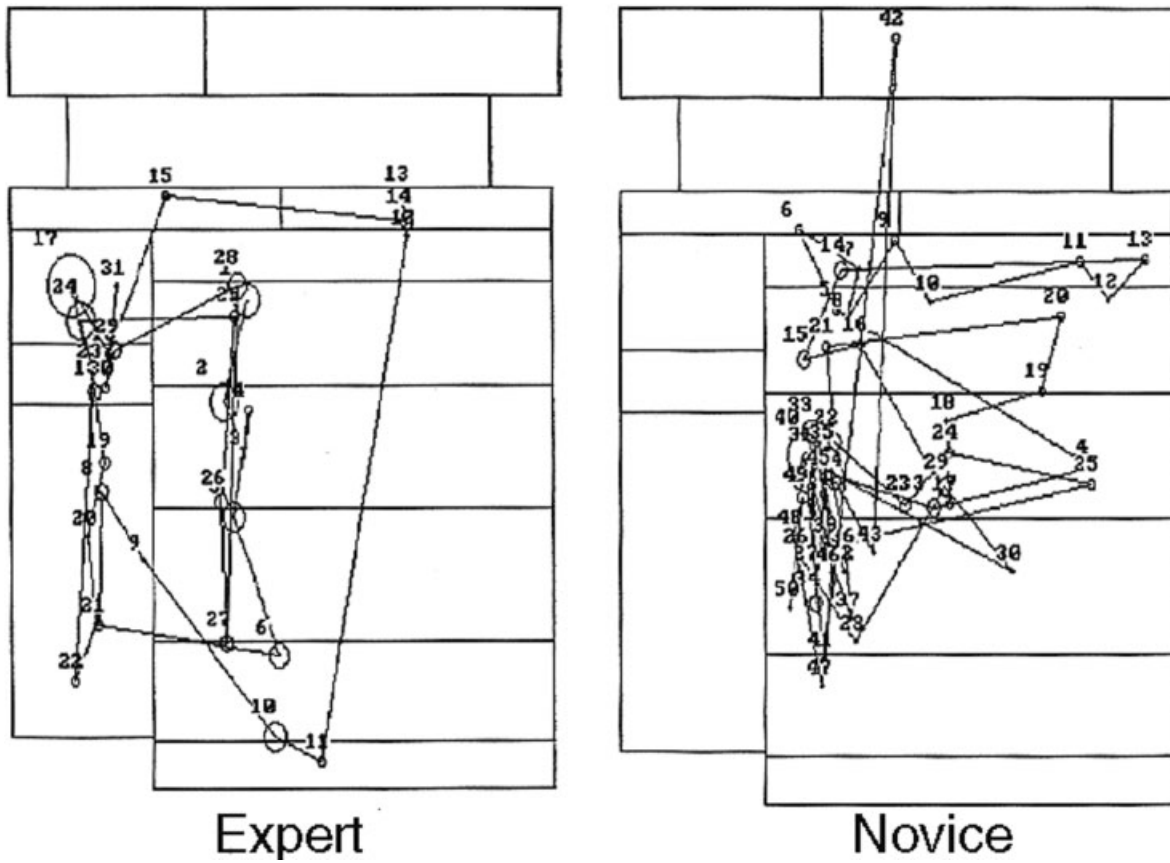


FIGURE 3
Experts vs. Novices

and a hint of an answer to the second question (what can advertisers do to improve banner effectiveness?). The study shows that one of the problems hindering banner ad effectiveness is that *half of the banner exposures are not attended to*. The problem is not only that surfers do not look at the banners, but they also seem to *purposefully avoid looking at them* (Hypothesis 1).

That only half the banner ads are looked at is highly detrimental to click-through rates. One cannot click on something one does not look at! It does not mean, however, that half of the banner exposures are wasted. Research has shown that consumers do not need to fully process a message in order to be influenced by it. Janiszewski (1990a, 1990b, 1993) has researched the topic extensively. Among other things, his research shows that incidental exposure to advertising can enhance a consumer's liking for the ad and brand advertised despite

the consumer's inability to recognize having previously seen the ad and brand (a situation similar to ours). Other researchers (e.g., Shapiro, MacInnis, & Heckler 1997) have reached similar conclusions. This means that a large part, if not most, of a consumer's processing of banner ads will be done at a pre-attentive level rather than at a full attention level. Further, it implies that click-through rates will not capture the full extent of an ad's effectiveness since pre-attentive processing does not lead to immediate action.

Our study also reveals that experts are more efficient at processing Web pages than novices and that young surfers are more efficient than older ones. This does not, however, translate into fewer banners seen by experts or young surfers.

As to what factors might help improve banner effectiveness, we found that location, size,

and zone content are important factors when trying to predict whether a zone is attended. We further investigate these factors in Study 2.

STUDY 2: IMPACT OF ONLINE ADVERTISING

Study 1 provides us with somewhat discouraging results about Internet advertising. Although everybody fixated on at least one ad, only half of the subjects remembered seeing a banner ad during the experiment. Further, the reported recognition rate for ads they have been exposed to is similar to the recognition rate for ads they have never seen before.

The purpose of Study 2 was to further explore these findings using a different methodology. Indeed, eye-tracking studies have two major drawbacks. First, they are artificial environments; second, their cost and infrastructure prevent their use on large samples. Hence, we sought to validate the results of the first study in a much broader experiment (~1,000 respondents). We build on Study 1 and use traditional memory-based measures of advertising effectiveness to establish how size, shape, and content affect advertising effectiveness. The study uses the same benchmarks for online advertising as those that have been used for advertising in traditional media: unaided advertising recall, brand recognition, aided advertising recall, and brand awareness (Tellis, 1998). These four measures are sorted by increasing level of advertising effectiveness. Brand awareness is the most desired effect; aided advertising recall is the least preferred.

Previous studies (e.g., Chatterjee, Hoffman, & Novak, 1998) have shown that repetition has a negative impact on click-through rates. Click-through rates are maximum after the first exposure. This is at odds with studies of television advertising (Pechman & Stewart, 1989) that shows how a low level of repetition is beneficial. The conflict in results might be due to the different metrics used to measure effectiveness: click-through (short-term action) versus recall (long-term memory). If one believes that some of the banner ad processing is made at a pre-

attentive level, one would expect to see the same beneficial impact of repetition online that one sees offline. Hence, to resolve the issue of the positive or negative impact of repetition we phrase hypothesis three as:

H₃: Advertising effectiveness increases with frequency of exposure.

Our logit analysis performed in Study 1 gives us some insights as to what attributes of a banner ad will affect its ability to attract attention. The factors included in the logit analysis were size, shape, content, and location. We will not study the impact of location, as it is an attribute over which advertisers have little control. As we have mentioned before, site designers typically use the prime locations to display the informational content of their site, and relegate banner ads to the top of their pages.

As far as size is concerned, our analysis shows that bigger is definitely better, although the S-shaped nature of the logit link function implies that at some level one will see decreasing marginal returns from increases in size. This leads to our next hypothesis:

H_{4a}: Larger banner ads will be more effective than smaller banner ads.

Study 1 also revealed that orientation matters. A vertical zone is more likely to be attended to than a horizontal one. Consequently, we will have the following corollary to Hypothesis 4a:

H_{4b}: Banner ads that are laid out vertically will be more effective than banner ads that are laid out horizontally.

Besides location and shape, Study 1 showed that a zone's content impacts its ability to attract attention. The next series of hypotheses will deal with this issue. The size of the *Contrast* coefficient in Table 1 shows that it is an important factor. Hence, we will formulate our next hypothesis as:

H_{5a}: Banner ads that contrast with their environment will be more effective than ads that do not.

One characteristic that sets banner ads apart from other elements on a Web page is animation. There is some anecdotal evidence that users are put off by blinking banner ads (Hamilton, 1998). However, animation provides advertisers with the ability to put more content in the same space by spacing the content temporally rather than spatially. Animation also allows advertisers to build dramatic effects by delivering the message sequentially. This increased flexibility should translate in more effective ads. Thus, we hypothesize that:

H_{5b}: Animated banners will be more effective than static ones.

As they have gained experience with the online medium, advertisers also have become more and more creative with their use of the limited space available on a banner. They have gone from simple images to fully interactive games, from static words to audible messages. The tools they have at their disposal are now richer. This should, hopefully, translate into more effective ads. This presupposes, of course, that advertisers use these tools effectively and that execution matters.

While it is hard to evaluate the ability of advertisers, it is relatively easy to figure out whether their ability matters by testing the impact of different executions of the same concept. We will thus formulate Hypothesis 5c as:

H_{5c}: The effectiveness of a banner ad depends on its message.

In summary, we have generated six hypotheses that fall into three different types. H₃ relates to the effect of frequency on online advertising effectiveness. H_{4a} and H_{4b} relate to the physical characteristics of the banner box. Finally, H_{5a}–H_{5c} relate to the content of the banner box.

Study Design

The second study consisted of a two part self-administrated Web-based questionnaire. Respondents were recruited through a mass e-mail sent to 6,000 individuals. The e-mail asked them to participate in a study about *Voilà*, the search engine of France Telecom. The stated purpose

of the study was to better understand how people use the Internet. The e-mail contained a hyperlink that lead the respondent directly to the Web pages containing the survey. As an incentive, two digital cameras and 25 CD-ROMs were to be awarded to participants through a random drawing.

The questionnaire began with three questions addressing a respondent's current Internet usage and brand awareness. Once these questions were answered, respondents were led through nine Web pages in which they had a specific task to accomplish or a question to answer (e.g., "Enter the word 'Voyage' in the SEARCH string input box"). Seven of the nine pages contained a banner ad. The ads related to five different products. Four of the products were online services offered by France Telecom: *@près l'école* (a subscription based homework help Web site), *Le Mél* (a free e-mail service), *Nouba* (a listing of the time and places of films, plays, and concerts), and *Tout En Ville* (a city guide). The fifth advertised product was *Cortal*, a large French financial institution. Each respondent saw one banner ad for each of the five services plus a second ad for both *Cortal* and *@près l'école*, for a total of seven exposures.

Each of the banner ads was executed with different creative specification purposefully designed to test the hypotheses derived in the previous section. For instance, there were five ads for *Le Mél*, each with a different colored background (blue, yellow, orange, red, or green). *Cortal* was executed in two different sizes: 230 × 33 pixels and 468 × 60 pixels. A total of 18 different banners were used. The details of each execution will be discussed along with the experimental results in the next section. The order in which the banners were presented as well as which execution of each banner was used was counterbalanced across respondents.

Twenty-four hours after they filled out the first survey and went through the nine test Web pages, the respondents who completed the first task were sent a second e-mail asking them to fill-out a follow-up survey. This survey asked them several questions about their advertising recall, and to rate the banner ads to which they



FIGURE 4
Sample Censored Banner

were exposed (see Appendix B for a list of the questions). Of the 6,000 e-mails that were originally sent, 211 came back as undeliverable; 2,297 of the remaining 5,789 prospects successfully completed the first part of the experiment. After being contacted the next day, 807 respondents filled out the second survey, for an overall response rate of 13.5 percent.

Analysis

Using our pre- and post-stimuli questionnaire, we computed the four effectiveness measures in the following way. Changes in brand awareness were calculated by taking the difference in brand awareness between the pre and the post questionnaire (see Question 1 in Appendix B). We find an increase in brand awareness across all five advertised brands (2.8% average increase) while the nonadvertised brands exhibit a small decrease (-1.5% on average). The increase in awareness for the advertised brands is statistically significant ($p < 0.0001$).

Unaided advertising recall was measured using Question 2. Out of 807 respondents, 460 (57%) recalled having seen an ad during their task the previous day. The five brands with the highest recall (11.4% average recall) are those that were advertised during the test. This unaided recall is significantly higher than the 'recall' of brands that were not advertised (2.47%, $p < 0.001$).

Aided advertising recall and brand recognition were measured by showing respondents copies of the five banner ads where the name of the announcer had been masked (see Figure 4 and Question 3 of the questionnaire). We then asked them whether they recalled seeing the ads during the experiment and, if so, could they name the brands that are advertised. We find

that, across all five banners, 30.1% of respondents recalled seeing the banners. We can compare this recall level to the level obtained for phony ads in Study 1 (23.45%) or to the one obtained by the Internet Advertising Bureau (IAB) France/SOFRES (17%). We find that our recall measure is greater than either benchmark ($p < 0.001$). Further, 61.8% of the respondents who stated that they had seen the ad the previous day were capable of correctly naming the brand advertised. This translates to 18.5% aided brand recognition.

At the outset, our results (summarized in Table 3) show that banner advertising has a long-term impact beyond the immediate click-through. This replicates the findings of Briggs and Hollis (1997) using a more traditional measure. It also means that we can use these measures to test Hypotheses 3-5.

The effect of repetition was tested across all five advertisers. For aided and unaided recall, as well as for brand recognition, we used a simple logit model using a dummy variable indicating whether one or two exposures had been used. For brand awareness, we looked at all 10 brands

TABLE 3
Overall Banner Exposure Effects

<i>Advertising Construct</i>	<i>Effect</i>
Exposure	100.0%
Aided Advertising Recall	30.1%
Aided Brand Recognition	18.5%
Unaided Advertising Recall	11.4%
Increased Brand Awareness	2.8%
Click-through*	0.7%

* From DoubleClick (2003).

TABLE 4
Effect of Repetition

Construct	<i>H₃: Repetition</i>	
	<i>Pr</i> > χ^2	N
Aided Advertising Recal	0.0047*	3597
Aided Brand Recognition	0.0021	1078
Unaided Advertising Recall	0.0088	6440
Brand Awareness	0.0001	16140
Overall	Effect	

* Although significant, this effect is of the wrong sign.

measured in Question 1. To accommodate for baseline awareness levels, we included brand intercepts and used both pre- and post-measurement, and used a dummy variable to estimate the effect of the test (negative coefficient, but not significant) and a continuous variable for the effect of exposure. We report in Table 4 the significance level of the exposure parameter for each of the four measures. As one can see, we have significant effects for all four measures. However, the effect on aided recall is the opposite of the other three measures. That is, repetition improves unaided advertising recall, brand recall, and brand awareness, but does not improve aided advertising recall. These data suggest support for Hypothesis 3 that repeated exposures increase online advertising effectiveness.

The procedure to test the remaining five hypotheses was the same for each hypothesis. In order to be as concise as possible, we will describe it only once, leaving it to the reader to make the necessary adjustments to fit each case. We test each hypothesis by fitting logit regression models on our four advertising effectiveness constructs: unaided advertising recall, aided advertising recall, brand recognition, and brand awareness. In each case, we set up a series of dummy variables that reflect the experimental design used, and we test whether these dummies are significantly different from zero. The regression only used the data gathered for the specific banners used in each test (e.g., when testing for background effect, only the banners

related to *Le Mèl* were analyzed). Rather than report detailed numbers about each of the regression, we report only the significance level of the dummies as a whole.

When interpreting these tests, we consider that an ad is more memorable than another if, at the minimum, it leads to higher *brand recognition*. A banner that leads to higher aided advertising recall, but does not improve brand recognition, unaided advertising recall, or brand awareness will not be considered effective as surfers can remember the banner, but not the brand that was advertised.

H_{4a} and H_{4b} relate to the size and orientation of banners. We tested three different banner sizes. We used the traditional 468×60 pixel banner as our null case, and pitted it against both a quarter (230×33^1) and a double (468×120) size banner. To test vertical versus horizontal orientations, we pitted a 144×240 pixel banner against the traditional 468×60 . The significance levels for the logit regressions are shown in Table 5.

Overall, we do not find any support for H_{4a} . Small banner ads perform just as well as large ones. This contradicts the findings of Study 1, which showed that bigger is better. The lack of size effect might be due to the decreasing returns from larger size. The small banner ads might be big enough to attract attention and the benefits gained from the larger ads might be too small to be measured accurately given our sample size. We find very weak support for H_{4b} . It seems that banner orientation affects aided advertising recall, but not unaided recall or brand recall. It might also have a small effect on brand awareness, although the effect is only marginally significant at best ($p = 0.17$). The fact that only aided advertising recall is significant might indicate that our subjects remember the shape of the banner more than its content.

H_{5a} to H_{5c} relate to the graphical content of banner ads. To test H_{5a} we used five banner ads that differ only in their background color (red, blue, yellow, orange, or green). Given the back-

¹ 230×33 is slightly larger than a quarter of 468×60 , but it is a widely accepted banner size.

TABLE 5
Size and Orientation Significance Level

Construct	<i>H_{4a}</i> : Size		<i>H_{4b}</i> : Orientation	
	<i>Pr</i> > χ^2	N	<i>Pr</i> > χ^2	N
Aided Advertising Recall	0.9483	3926	0.0003	918
Aided Brand Recognition	0.6397	1018	0.1716	240
Unaided Advertising Recall	0.3056	2360	0.9314	552
Brand Awareness	0.8785	4146	0.1714	962
Overall	No effect		Weak effect	

ground color of the pages and the color of the elements surrounding the ads, the yellow and orange ads were coded as not contrasting, while the other ads were contrasting. To test *H_{5b}*, we used a still and an animated banner ad, both displaying the same message. Finally, to test *H_{5c}*, we compared four different ads for the same product. Our base case ad is a traditional 468 × 60 pixel banner ad with the brand name of the sponsor. The second ad is an offer to participate in a game associated with the sponsor. The third execution has a dropdown menu embedded in the banner ad. A sound file that enunciates the name of the brand being advertised as the banner ad is displayed accompanies the last one.

As reported in Table 6, we find a weak effect for *Contrast* (*H_{5a}*). However, the effect goes opposite to the direction predicted. High contrast ads seem to perform worse than noncontrasting one. We do not find any effect for *Animation*

(*H_{5b}*), but we find support for *H_{5c}*. The fact that brand recognition is significant for *H_{5c}* is important because it shows that the different messages used affect not only the recognition of the visual component of the ads as is the case for *H_{5a}* and *H_{5b}* but also the comprehension of the message. Hence, ad content is important. Further, *the message is more important than how the message is sent*.

CONCLUSION

We started with an eye-tracking study conducted on 49 subjects. This first study revealed that a banner ad typically has a probability of about 50% of being seen by a surfer looking at a page in which the banner is embedded. This number is dramatically lower than the 97% reported by Siddarth (2002) for television ads or 93% reported by Lohse (1997) for Yellow Pages ads, indicating that pre-attentive ad processing

TABLE 6
Contrast, Execution, and Animation Significance Level

Construct	<i>H_{5a}</i> : Contrast		<i>H_{5b}</i> : Animation		<i>H_{5c}</i> : Message	
	<i>Pr</i> > χ^2	N	<i>Pr</i> > χ^2	N	<i>Pr</i> > χ^2	N
Aided Advertising Recall	0.0012	1500	0.0989	900	0.0001	1182
Aided Brand Recognition	0.5407	540	0.6084	216	0.0022	418
Unaided Advertising Recall	0.1194	902	0.2726	540	0.9297	716
Brand Awareness	0.2719	1582	0.6447	944	0.6114	1266
Overall	Weak effect		No effect		Effect	

might be much more important on the Internet than in more traditional media. We also found that different levels of expertise, age, or gender lead to differences in page processing behavior. However, these differences did not translate into different probabilities of seeing banner ads. A logit analysis of the eye-tracking data revealed that some of the factors that would most likely influence the probability of content being looked at on a Web page are a zone's size, shape, content, and location.

We used the results of the eye-tracking experiment to design a large-scale experiment, where 807 subjects were exposed to nine pages containing six banner ads. The banner ads were designed to test a series of six hypotheses relating to size, shape, color, animation, message, and repetition. The study was also used to test whether traditional memory-based measures of advertising effectiveness could be applied to the Internet.

We found that traditional memory-based effectiveness measures provide valuable insight into the effects of Internet advertising. These measures outperform the immediate effects measured by the click-through rate. On average, for 100 surfers exposed to a banner ad, 11 recall seeing the ad and can mention the brand name on the ad without any aid 24 hours later (refer back to Table 3). Thirty respondents remember seeing the banner when they are shown the same banner but without brand name. Of those 30 surfers, 18.5 (62%) can name the banner's brand. In addition, three of the 100 surfers become aware of the brand. We can compare these numbers with the average click-through rate of 0.7% (DoubleClick, 2003). The effect on unaided brand awareness is four times larger than the click-through rate; the effect on unaided brand recall is 19 times larger.

In terms of what factors influence online advertising effectiveness we found that frequency is important. Repetition affects unaided advertising recall, brand recognition, and brand awareness. However, artistic execution has little effect. Banner size seems to be unimportant. Contrast, animation content, and shape of the banner influence aided advertising recall, but

no other dependent measure. Nevertheless, a banner's message influences both aided advertising recall and brand recognition. This indicates that what an ad says is more important than how it says it.

We started this research seeking to answer two basic questions regarding online advertising. Given the steady decline in click-through rate over the past 5 years and questioning the predicted resurgence of online spending, we asked: (1) Why do banner ads seem to be ineffective and (2) what can advertisers do to improve their effectiveness?

As to why banner ads seem to be ineffective, we believe that they are actually effective. Banners lead to brand awareness. Click-through rate might be low, but, in the long run, awareness is more important than click-through. The low click-through rates are primarily due to the fact that surfers fixate on less than 50% of the banners to which they are exposed. Not only do they not look at banners, they actually avoid looking at them!

Further, advertisers can increase advertising effectiveness by concentrating on the message they send. What is said and how often it is said is more important than how it is said. Bigger ads or animated ads will not compensate for ineffective content.

One should also be aware of some of the shortcomings of our research. First and foremost, all of our subjects were engaged in a goal-directed search. As Janiszewski (1998) has shown, goal-directed searches are very different than exploratory searches. Second, we examined memory measures, but not behavior measures (e.g., sales or "store visits").

APPENDIX A: ZONE DESCRIPTION VARIABLES

Building on past research (e.g., Janiszewski, 1998; Steinman & Levinson, 1990), we hypothesize that a zone's ability to attract a subject's gaze will be a function of its location on the page, its size and shape, and its pictorial or textual content. Location was captured by the *x* and *y* coordinates of the center of each zone, as

well as by the distance to the nearest zone seen by the subject. The top-left corner of the screen was defined as the origin point (0, 0); the bottom-right corner was set to (255, 255). We also created squared terms for the x and y coordinates to allow for the fact that the center of the pages might be preferred.

Shape was described in terms of the zone's area as well as its orientation² (horizontal vs. vertical) and the general orientation of the page. An interaction dummy was also created to indicate whether the zone and the page follow the same orientation.

The content of each zone was coded using the following dummy variables:

- TEXT (set to 1 if the zone contains text, 0 otherwise)
- IMAGE (1 if the zone contains nonphotographic graphical elements)
- PICTURE (1 for photographs)
- INPUT (1 if the zone contains input elements such as buttons)
- CONTRAST (1 if the zone contrasts with its surrounding)
- AD (1 if the zone contains an advertisement).

Our last control variable is a variable that indicates the number of zones present on each page.

APPENDIX B: ADVERTISING RECALL QUESTIONNAIRE

The advertising recall questionnaire consisted of a brief introduction followed by three multi-part questions. The first question (see below) measured brand awareness. It was the same exact question as the brand awareness question asked in the first questionnaire that respondents completed before being exposed to the stimuli. The second question measured unaided advertising recall. The third question measured aided advertising recall and aided brand recognition.

² Remember that all zones are rectangular.

Q1—For the following brands or services, please check all those that you know or have heard of³:

- Le Mel
- Alta Vista
- Nomade
- Voila
- Tout en ville
- Lycos
- Cortal
- Nouba
- @près l'école
- Yahoo!

Q2—While completing the task yesterday, do you recall seeing one or more banner ad?

- 1—Yes
- 2—No

If yes, please enter the brand(s) advertised.
(Free flow multi-line text box)

Q3—For each of the following banner ads, do you recall seeing it during the task:

A censored version of each of the ads seen by the respondent the day before is shown here (see Figure 4 for an example of censored ad).

- 1—Yes
- 2—No

If yes, enter the name of the brand advertised.

(Text box)

REFERENCES

- Bass, S. (1999). Unclog the Net: Eliminate Web Ads, PC World Online, January. Retrieved, from <http://www.pcworld.com/resource/article/0,aid,8731,00.asp>
- Berst, J. (1999). Free PCs, ISPs and Software: Are They Worth the Price? ZDNet, February 10. Retrieved, from http://www.zdnet.com/anchordesk/story/story_3080.html
- Briggs, R., & Hollis, N. (1997). Advertising on the Web:

³ The questionnaire has been translated from French for the reader's convenience.

- Is There Response before Click-Through? *Journal of Advertising Research*, 37 (2), 33–45.
- Chatterjee, P., Hoffman, D.L., & Novak, T.P. (1998). Modeling the Clickstream: Implications for Web-Based Advertising Efforts. Project 2000. Nashville, TN: Vanderbilt University.
- DoubleClick (2003, January). DoubleClick 2002 Full-Year Ad Serving Trends. Retrieved, from http://www.doubleclick.com/us/knowledge/documents/trend_reports/dc_2002adservingtrends_0212.pdf
- Drèze, X., & Zufryden, F. (2000). Internet Advertising: The Medium is the Difference. *Consumption, Markets and Culture*, 4 (1), 23–37.
- Dyson, P., Farr, A., & Hollis, N. (1996). Understanding, Measuring, and Using Brand Equity. *Journal of Advertising Research*, 36 (6), 9–21.
- Fischer, P.M., Richards, J.W., Berman, E.J., & Krugman, D.M. (1989). Recall and Eye Tracking Study of Adolescents Viewing Tobacco Advertisements. *Journal of the American Medical Association*, 261 (1), 84–89.
- Forrester (2001). Online Advertising Retrenches. The Forrester Report, October. Retrieved, from <http://www.forrester.com/ER/Research/Report/Summary/0,1338,11005,FF.html>
- Forrester (2002). Online Advertising Picks Up Again, TechStrategy Research, May. Retrieved, from <http://www.forrester.com/ER/Research/Report/Summary/0,1338,14576,FF.html>
- Hamilton, A. (1998). Secrets of Super-High Web Ad Click-Through. ZDNet, August 19. Retrieved, from http://www.zdnet.com/anchordesk/story/story_2439.html
- Internet Advertising Bureau (IAB). (2003). IAB and PwC Release “Top-Line” Data from 2002 Internet Advertising Revenue Report. Retrieved, from http://www.iab.com/news/pr_2003_4_9.asp
- Internet Advertising Bureau (IAB). (1999, May). Paris: France/SOFRES Net.Imp@ct.
- Janiszewski, C. (1990a). The Influence of Nonattended Material on the Processing of Advertising Claims. *Journal of Marketing Research*, August, 263–278.
- Janiszewski, C. (1990b). The Influence of Print Advertisement Organization on Affect Toward a Brand Name. *Journal of Consumer Research*, 17 (June), 53–65.
- Janiszewski, C. (1993). Preattentive Mere Exposure Effects. *Journal of Consumer Research*, 20 (December), 376–392.
- Janiszewski, C. (1998). The Influence of Display Characteristics on Visual Exploratory Search Behavior. *Journal of Consumer Research*, 25 (3), 290–301.
- Kroeber-Riel, W. (1979). Activation Research: Psychological Approaches in Consumer Research. *Journal of Consumer Research*, 5 (March), 240–250.
- Lohse, G.L. (1997). Consumer Eye Movement Patterns on Yellow Pages Advertising. *Journal of Advertising*, 26 (1), 61–73.
- Novak, T.P., & Hoffman, D.L. (1997). New Metrics for New Media. Toward the Development of Web Measurement Standards. *World Wide Web Journal*, 2 (1), 213–246.
- Russo, J.E., & Leclerc, F. (1994). An Eye-Fixation Analysis of Choice Processes for Consumer Nondurables. *Journal of Consumer Research*, 21 (2), 274–290.
- Shapiro, S., MacInnis, D.J., & Heckler, S.E. (1997). The Effects of Incidental Ad Exposure on the Formation of Consideration Sets. *Journal of Consumer Research*, 24 (June), 94–104
- Siddarth, S. (2002). Describing the Dynamics of Attention to TV Commercials: A Proportional Hazards Model of the Time to Zap an Ad. USC Working Paper. Los Angeles, CA: University of Southern California.
- Steinman, R.M., & Levinson, J.Z. (1990). The Role of Eye Movement in the Detection of Contrast and Spatial Details. In E. Kowler (Ed.), *Eye Movements and Their Role in Visual and Cognitive Processes* (pp. 115–160). Amsterdam: North-Holland, Elsevier.
- Tellis, G.J. (1998). *Advertising and Sales Promotion Strategy*. Reading, MA: Addison-Wesley.
- Wall Street Journal Online. (2003). Online Ad Spending Is Expected To More Than Triple by Year 2004. Retrieved, from http://www.performersalliance.tv/archive/wsj_8_12_99.html