Automated Computational Aesthetics for Pervasive Graphical Advertising

Thomas Phan*

Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
thomas.phan@acm.org

Abstract. Graphical advertisements can be seen on digital signage, mobile phones, and online webpages, and fomenting user interest in the ads is important for both publishers and advertisers. This work addresses the issue of making Web graphical ads more attractive and compelling, with the end result hopefully being a positive emotional response from the user that produces a click. We developed a system for displaying advertising images in an attractive, pleasing manner using automated heuristics based on fundamental colour theory whereby our system automatically analyses an image, extracts important colour characteristics, and dynamically sets key colours in a surrounding frame.

1 Introduction

Graphical advertisements are an increasingly pervasive and important component of advertising campaigns, with such ads being seen in many electronic contexts, including digital signage in public spaces [5], targeted messaging in mobile phones [6], and online static banners or Adobe Flash objects on a webpage. However, drawing the user's attention to the ads in a positive manner continues to be difficult. In this paper, we directly address this problem in the context of online Web advertisements, although our approach can be applied to other pervasive advertising mediums that display graphical ads.

In Web advertising, advertisers own an ad, and webpage publishers show ad impressions (the rendered instances of an ad). Graphical online advertising is a multi-billion dollar industry that generates revenue for the publishers and product interest for the advertisers [2]. Graphical ads can be part of either cost-per-click or cost-per-mille advertising campaigns, and in both cases, fomenting user interest and generating clicks on graphical ads are important for both the publishers and the advertisers. A high click-through rate (the ratio of the number of clicks to the number of impressions) is thus very desirable.

However, attaining clicks on a given ad impression can be difficult because users may simply not be interested in the advertisement or may be experiencing advertising fatigue. Publishers attempt to fight this indifference by offering

* This work was conducted while the author was with Yahoo! Inc.

more relevant advertisements based on user demographics or the content of the page [1]. Nonetheless, graphical ads in particular perform poorly, with click-through rates above 1% typically being considered successful.

This work addresses the issue of *making the graphical advertisements themselves more attractive and compelling*, with the end result hopefully being a positive emotional response from the user that produces a click. In Section 2 we discuss a design and implementation that automatically chooses aesthetically-pleasing advertisement colour schemes from advertisers' image catalogs, and in Section 3 we discuss our vision of advertising in the future.

2 A system for automated aesthetics

We developed a prototype for displaying graphical ad images in an attractive, pleasing manner using automated heuristics based on fundamental colour theory. Although aesthetics are inherently subjective, we note that in informal demonstrations, viewers of our prototype were very pleased by our results, some of which are shown in Figure 1.

This work addresses a use case where an advertiser wants to run a Web-based graphical advertising campaign that rotates images from a catalog of their goods. Since the number of images may be large (for example, it may be a collection of clothing images on the order of hundreds), crafting a pleasing colour scheme for all the images may be too large a task, so the inclination might be to place the image inside a generalised graphical frame with some default colour. While this approach is sound, it also produces a rather visually-dull set of ads.

In our work we developed a system that automatically analyses each image, extracts important colour characteristics, and dynamically sets key colours in a surrounding Adobe Flash frame that we can co-develop with the advertiser. For images smaller than 1000×1000 pixels, our approach is generally very fast, taking approximately a second to process an image using an off-the-shelf Core2Duo desktop running Red Hat Linux. Our system was written in C++/STL using the open-source ImageMagick library.

The processing proceeds in a series of steps. First, given an image, we crop and retain its center 25% region by area: given an image of width W and height H, we retain the rectangular region with diagonal corner coordinates $(\frac{1}{4}W, \frac{1}{4}H)$ and $(\frac{3}{4}W, \frac{3}{4}H)$. This cropping allows us to concentrate on the image's subject matter so that important colours can be later extracted. We note that the subject of a photo may not be in the center, and in fact in general photography, it is very often recommended to place the subject on one of the $\frac{1}{3}$ axes of the image according to the Rule of Thirds [4]; however, in commercial product photography the subject is almost always directly in the center.

Second, we analyse the cropped image pixel-by-pixel, quantise each pixel's colour, and create a histogram of bucketised RGB values, typically with 32,768 buckets and sorted by brightness. From this histogram we can determine the image's characteristics with regard to the colour values in its highlights, midtones, and shadows as well as the colours' frequency of occurrence.

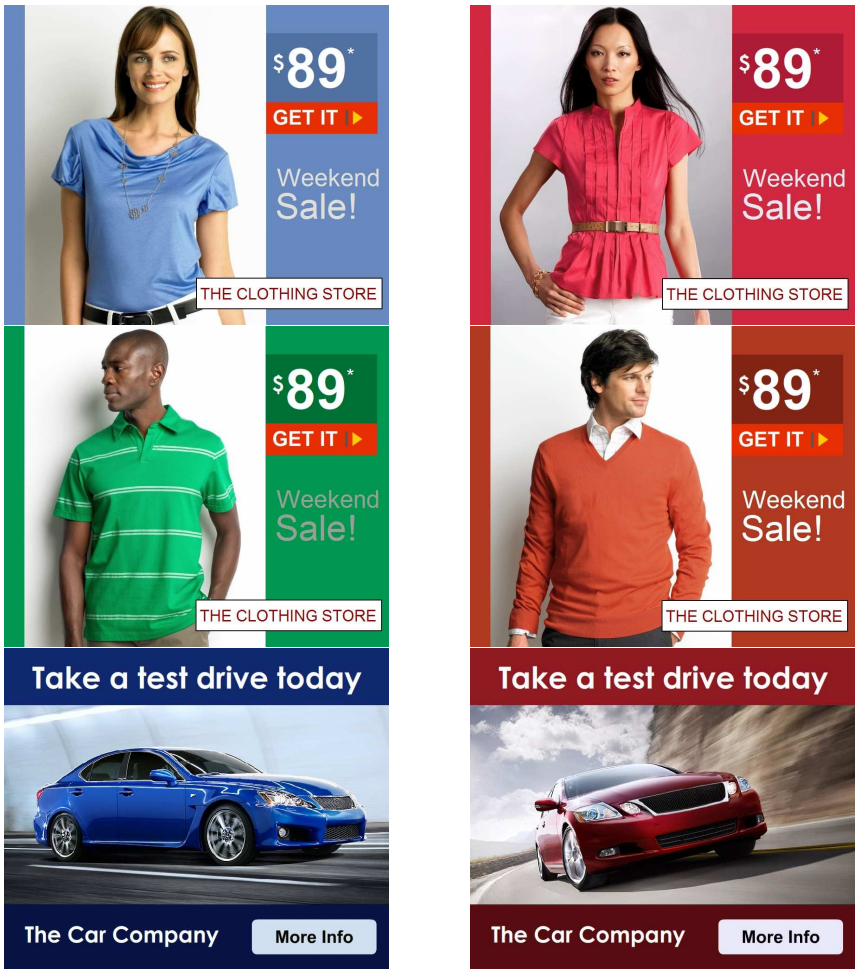

Fig. 1. Sample ads showing the surrounding colours being based on an image's highlights, midtones, and shadows. The images are copyrighted by their respective owners.

Third, following a set of business rule heuristics, we pick colours from the histogram from which we can set the surrounding colours in the frame. In the images shown in Figure 1, we picked up to three colours from the image which are then stated in an XML file that the Flash object reads and uses at display-time to set its colours (using ActionScript). In the clothing ads, the colours for the surrounding frame and the block around the price were set with the image's prevalent midtone and shadow colours, respectively. A similar approach was taken in the car ads, with the colour of the button using the prevalent highlight.

The business rules used in last phase are obviously the most crucial and must be done in accordance to the advertisers' tastes. In other trials of our prototype, we also used frame colours that were chosen as the complement of colours found in the image (where we use the term "complement" in the strict

sense of the colour that is 180 degrees diametrically opposite another colour on an RGB-CMY colour wheel [3]). A variety of other rules may be used, such as constraining the colours for users in particular demographics (such as pastels for women and vibrant colours for users under 25). Further, in a pervasive computing environment, these rules might leverage such context as audience eye movement, the weather and season, or ambient light. Once such heuristics are laid out, images can be driven through the system to produce the automated renderings.

3 Conclusion

Over the next 25 years, webpages, mobile phones, and digital signage will host an ongoing war between the advertisers eager for user interest and interaction and the users irritated by increasingly-pervasive advertisements. We envision a future where ads will be completely personalised and delivered with surgical precision. Even when the anonymity of a user is preserved, a tremendous amount of information can today already be drawn and inferred from basic Web-surfing behaviour that can reveal, for example, the user's gender, age, income, political leanings, and personal interests. Given this information, an advertiser (with the help of the publisher's platform) is able to deliver a specifically-tailored ad. On the other hand, users have already developed a keen sense of dullness in responding to online advertising. Additionally, ad-blocking systems deployed at either the user's browser or in proxies make the advertiser's job more difficult.

Our opinion is that digital pervasive advertising 25 years from now must take the form of self-standing, aesthetically-compelling works with which users will *want* to interact. Such advertising must be able to provoke a positive emotional response using the digital medium in a way similar to how TV commercials during the Super Bowl are considered a significant form of entertainment. Our work in developing an automated system to produce visually-pleasing graphical advertisements is based on basic colour theory, and we expect future successful pervasive advertising to leverage an intersection of the innovativeness of computer science research and engineering, the aesthetics of graphic design and photography, and the business acumen of advertising agencies.

References

1. A. Broder, M. Fontoura, V. Josifovski, and L. Riedel. "A Semantic Approach to Contextual Advertising," In *Proceedings of SIGIR*, 2008.
2. S. Clifford. "A Web Shift in the Way Advertisers Seek Clicks," The New York Times, April 21, 2008.
3. R. Hirsch. Exploring Color Photography. Mcraw-Hill, 2004.
4. B. London, J. Upton, J. Stone, K. Kobre, and B. Brill. Photography, 8th ed. Prentice Hall, 2004.
5. K. O'Hara, M. Perry, and E. Churchill (eds.). *Public and Situated Displays: Social and Interactional Aspects of Shared Display Technologies*. Kluwer, 2003.
6. M. Reardon. "Advertising seeps into the cell phone," CNET News, September 14, 2006. http://news.cnet.com