

The Impact of Visual Appearance on User Response in Online Display Advertising

Javad Azimi^{* ◊}, Ruofei Zhang[†], Yang Zhou[†]
 Vidhya Navalpakkam[†], Jianchang Mao[†], Xiaoli Fern[◊]
[◊]Oregon State University
[†]Yahoo! Labs, Silicon Valley
 {azimi,xfern}@eecs.oregonstate.edu
 {rzhang,yangzhou,nvidhya,jmao}@yahoo-inc.com

ABSTRACT

Display advertising has been a significant source of revenue for publishers and ad networks in the online advertising ecosystem. One of the main goals in display advertising is to maximize user response rate for advertising campaigns, such as click through rates (CTR) or conversion rates. Although the visual appearance of ads (creatives) matters for propensity of user response, there is no published work so far to address this topic via a systematic data-driven approach. In this paper we quantitatively study the relationship between the visual appearance and performance of creatives using large scale data in the world's largest display ads exchange system, RightMedia. We designed a set of 43 visual features, some of which are novel and some are inspired by related work. We extracted these features from real creatives served on RightMedia. Then, we present recommendations of visual features that have the most important impact on CTR to the professional designers in order to optimize their creative design. We believe that the findings presented in this paper will be very useful for the online advertising industry in designing high-performance creatives. We have also designed and conducted an experiment to evaluate the effectiveness of visual features by themselves for CTR prediction.

Categories and Subject Descriptors

I.4.7 [Computing Methodologies]: Image Processing and Computer Vision—*Feature Measurements*

General Terms

Algorithms, Design, Performance

Keywords

Online Advertising, Visual Features, Creative Recommendation

1. INTRODUCTION

The problem of predicting the user response rate for online ads, especially CTR, has been studied by several researchers in the last few years. One major research focus has been in predicting clicks by studying the relationship between CTR and user, webpage, and ads features. For example in [2], the authors considered the ad's relevancy to the content of the webpage in predicting CTR. Although it is generally believed that visually appealing display ads

*This work was done while the first author was an Intern at Yahoo! Labs.

can perform better in attracting online users, as a result of which advertisers always care about the creative designs, there is no, to the best of our knowledge, published work so far to quantitatively study the effect of visual appearance of creatives on advertising campaign performance in online display advertising. This motivates us to investigate the correlation between the visual features of the creative and CTR, independent of other factors such as (page, user, ads) relevance. This can help us to provide a set of practical and actionable recommendations to designers to help design better creatives that are more likely to be clicked by users.

Our proposed approach consists of three main steps, 1) feature extraction, 2) correlation investigation and 3) click prediction. We first extract some informative visual features from the creatives. We introduce 43 visual features classified into three categories, 1) *global features* which characterize the overall properties of a given creative, 2) *local features* representing the properties of specific parts within a given creative, and 3) *advanced features* which are a group of features developed based on more complicated algorithms such as the number of human faces and number of characters in a creative. We then evaluate the correlation between proposed features and CTR to propose a set of recommendations to ad designers. Finally, we develop a regression method to predict the CTR based on all developed features. The study is conducted using real creatives and their performance data from the world's largest display ads exchange system, RightMedia.

The benefit of this work is three-fold. First, our findings on the visual features and their relationship to CTR can provide useful recommendations to designers on what features to consider while designing creatives, and/or can help in automated creative generation. Second, the visual features and the regression methods employed here can augment the traditionally investigated ad factors (such as ad relevancy, ad position, etc.) to improve CTR prediction in online ads selection. Third, it provides the research community with the first ever data set, initial insights into the effect of visual appearance on user response propensity (e.g., CTR), and evaluation benchmarks for further study.

2. DEVELOPED FEATURES

We developed a set of 43 different visual features for this study. We categorize the developed features into three different sets, 1) global features, 2) local features and 3) advanced features. Below we list the developed features in each category.

Global Features. Global features are a set of features which represent the overall properties of the whole image. We developed 19 different global features as follows: f_1 : gray level contrast, f_2 : number of dominant bins in gray level histogram, f_3 : standard deviation of gray level images, f_4 : number of dominant bins in

RGB histogram, f_5 : size of the dominant bin in RGB histogram, f_6 : number of dominant bins in HSV histogram, f_7 : size of the dominant bin in HSV histogram, f_8 : deviation from the best color harmony model, f_9 : average deviation from the best two color harmony models, f_{10} : number of connected coherent components (CCC), f_{11} : size of the largest CCC, f_{12} : size of the second largest CCC, f_{13} : color size rank of the largest CCC, f_{14} : color size rank of the second largest CCC, f_{15} : number of dominant hues, f_{16} : contrast of dominant hues, f_{17} : standard deviation of hues, f_{18} : average lightness, and f_{19} : standard deviation of lightness.

Local Features. Local features represent a set of features extracted from specific parts of the image rather than the whole image. We apply the normalized cut segmentation method [4] to partition the image into 5 smaller segments. We then develop the following features, f_{20-32} , based on the segmentation result: f_{20} : size of the Largest Segments (LS), f_{21} : segments size contrast, f_{22} : number of image dominant hues in the LS, f_{23} : number of dominant hues in the LS, f_{24} : largest number of dominant hues in one segment, f_{25} : contrast of dominant hues among segments, f_{26} : contrast of hues in the LS, f_{27} : standard deviation of hues contrast among segments, f_{28} : deviation from the best color harmony model for LS, f_{29} : average deviation from the best two color harmony models for LS, f_{30} : average lightness in LS, f_{31} : standard deviation of average lightness among the segments, f_{32} : contrast of average lightness among the segments.

Advanced features. Advanced features are based on more complicated algorithms such as object (face/character) detection. Most of them are based on the saliency map (SM) [3] of the image which determines the visually salient areas in the image that are more likely to be noticed by the humans. Features f_{33-41} are then developed as follows: f_{33} : background size in SM, f_{34} : number of connected components in SM, f_{35} : size of the largest connected components in SM, f_{36} : average saliency weight of largest connected component (CC), f_{37} : number of CCs in image background SM, f_{38} : size of the largest CC of background in SM, f_{39} : distance between CCs in SM, f_{40} : distance from rule of third points in SM, f_{41} : distance from center of image in SM. We also extract two additional features related to the number of characters, f_{42} , and number of faces, f_{43} , in an image. The complete description of the generated features can be found in [1].

3. EXPERIMENTAL RESULTS

In this section we evaluate the relationship between visual features and the performance of creatives in online display advertising.

Data Set. We extracted creatives of advertising campaigns from the world’s largest online advertising exchange system, RightMedia. We filtered out animated creatives because our features are designed for static images. Then we calculated the average CTR of these creatives from online serving history log during a two-month period. To remove the impact on performance introduced by ad *position* and *size*, we created two data sets, each of which consists of creatives with the same size and same position. The first data set, ID2, consists of 6272 creatives with size 250×300 pixels and ads position LREC, and the second data set, ID6, includes 3888 creatives with 90×730 pixels and ads position SKY¹. All of the creatives have a minimum of 100K impressions guaranteeing that their CTRs have converged to their true values.

Creative Design Recommendation. We first calculate the Linear Correlation and Mutual Information between all features and CTR in each data set. Note that mutual information can provide us with the information of non-linear correlation between features. We also used forward feature selection to select the most informa-

tive features in predicting CTR. Based on our results, we provide the following set of recommendations to designers for optimizing creative performance:

DOs

- Creatives with higher gray level contrast achieve higher CTR.
- Small number of salient components, with all components close to the center of the creative and the major component consistent with the rule of third, achieves higher CTR.
- Creatives with good color harmony (those with small deviation from color harmony models) achieve higher CTR.
- Average lightness across whole image and the largest segment of the image has a positive correlation with CTR.

DONTs

- Cluttered creatives (those with large number of connected components) are unlikely to achieve high CTR.
- Creatives with large number of characters cause textual clutter and are unlikely to achieve high CTR.
- Too many different hues, in both the whole image and the largest component in the image, is not desirable.

CTR Prediction. In addition to providing recommendations to designers for optimizing creative design, we further studied the correlation between visual features with ads performance by testing the ability of visual features to predict CTR, as a result we use visual features only regardless of page and user relevance factors here. We used two different regression algorithms to predict CTR, 1) Linear Regression (LR) and 2) Support Vector Regression with RBF kernel (SVR). To evaluate the CTR prediction accuracy of the algorithms, we run each algorithm for 200 independent runs; for each run, 80% of the data set is selected randomly for training and 20% for testing. The accuracy evaluation results are reported over the prediction of the test data. Mean Squared Error (MSE) is used to measure the prediction accuracy for each algorithm. To meaningfully interpret the MSE value, we introduce two baseline approaches, 1) *Weighted Sampling* policy which samples from the CTR distribution of the training data to predict the CTR of each testing creative, and 2) *Uniform Voting* (UV) policy that assigns average CTR of the training creatives to testing creatives. Table 1 shows the result such that each entry is the MSE value of the *Weighted Sampling* policy divided by MSE value of each algorithm. Results show that we can perform up to 3.27 times better than *Weighted Sampling* policy in predicting CTR from visual features only. This result demonstrates the non-trivial impact of visual appearance of the creative on its advertising performance. In future work, we will use the visual features developed here along with the existing, non-visual features (user, page, ads) to further improve CTR prediction of creatives.

Table 1: Performance of each method.

Data set	Samples	UV	LR	SVR
ID2	6272	1.71	2.28	3.27
ID6	3888	1.75	2.27	2.77

4. REFERENCES

- [1] J. Azimi, R. Zhang, Y. Zhou, V. Navalpakkam, J. Mao, and X. Z. Fern. The impact of visual appearance on user response in online display advertising. *CoRR*, abs/1202.2158, 2012.
- [2] D. Chakrabarti, D. Agarwal, and V. Josifovski. Contextual advertising by combining relevance with click feedback. In *WWW*, pages 417–426, 2008.
- [3] X. Hou and L. Zhang. Saliency detection: A spectral residual approach. *IEEE CVPR*, pages 1–8, 2007.
- [4] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE PAMI*, pages 888–905, 2000.

¹LREC and SKY are two pre-defined ad positions on web pages.