Journal of Computers Vol. 30 No. 5, 2019, pp. 279-289
doi:10.3966/199115992019103005022

Is It Worth to Deliver Display Ads on Content Farm Websites?

Cheng-Jye Luh*, An-Chen Wu

Department of Information Management, Yuan Ze University, Taoyuan 320, Taiwan, ROC
imcjluh@saturn.yzu.edu.tw, j5720363@gmail.com

Received 15 March 2019; Revised 30 June 2019; Accepted 4 September 2019

Abstract. The study evaluated the performance of display ads on content-farmed pages as compared to those on non-content farm sites. Content farms generate a significant number of low-quality articles on many different topics and then use keywords to get their pages better positions on search engine results pages that enable web users to browse and click on ads. We compared the performance of display ads placed on content-farmed pages and regular ones. The results revealed that display ads in remarketing campaigns performed better than those in standard display campaigns on all type of websites. On the remarketing campaigns, display ads placed on the non-content-farmed pages have better overall and purchase conversion rates but lower click-through rate than those on content-farmed pages.

Keywords: advertising performance evaluation, content farm pages, display ads, Internet advertising

1 Introduction

Online advertising is a great way to acquire customers and to get engaged with customers in the digital era. Generally speaking, there are two kinds of online advertising available: one is search ads appearing on search engine results pages, and the other is display ads appearing on some websites. Google has been the most prominent search advertising platform on the market [1]; the top two multimedia display platforms are Facebook and Google [2]. Notably, display advertising spending exceeded search advertising spending since 2016.

Facebook delivered display ads on their platforms, so the quality is more comfortable to control, but Google delivered most of the display ads on the websites of the Google Display Network (GDN), except YouTube. Although Google has some mechanism to review the websites that apply to join GDN, we found from Google ads performance reports that about 8 percent of websites that ads appeared were content farm websites. Content farms are websites whose primary goal is to earn online advertising traffic. They usually use a variety of unethical or illegal methods to produce a large number of poor quality articles on popular keywords and then use some shaking article titles to attract users to click [3-5]. Interestingly, on the one hand, Google adjusts its search ranking algorithm to block content farms from appearing in search results, but on the other hand, it delivers display ads on this kind of websites, allowing them to earn revenue to maintain operations.

This study aimed to assess whether the performance of the display ads on content farms was significantly different from that on non-content farm sites and to provide recommendations to advertisers. The data we used were the performance data of display ads shown on the Google Display Network.

2 Literature Review

2.1 Low-quality Websites

* Corresponding Author

According to two Google guidelines [6-7], low-quality sites have three main characteristics: (1) automatically generating content; (2) participation in link manipulation plots -- either with a large number of outgoing links to other websites, or lots of incoming links from other websites; (3) having only a small amount of or no original content. Their common purpose is to manipulate the ranking of the spamming website in search engines. Many techniques are available for web spam detection [8], but to effectively detect spam pages is just like an “arms race” between search engines and website operators. For example, “keyword stuffing,” which adds some hot, but not relevant keywords to the content of webpages, is the most commonly used web spam technique for years. One study [9] proposed several heuristic features, including the number of words on a webpage, the proportion of pages containing popular keywords, to train a spam page classifier. The approach has a major flaw in that as long as a suspicious website uses different words, the detection performance would become significantly worse.

Becchetti et al. [10] built a link spam classifier using the link architecture obtained from the statistical features in the vicinity of training pages. Urvoy et al. [11] developed an HTML style similarity measure to identify suspicious spam pages. Dong and Zhou [12] found that the content of spam sites often spread across several topics and the variation among topics was relatively more substantial than that of regular websites. Thus they used a website theme diversity value as the spamming probability to detect web spam; however, this method was inclined to make false positives on themes like the sale of health care products and online mall.

Google has been promoting the importance of high-quality websites since 2011 [13] and applying some search filter (e.g., Panda Update) to penalize low-quality websites in search ranking. According to one Google patent [14], Google’s search ranking algorithm checks whether a ranked website is low-quality when it produces the draft search results for a user query. If the number of low-quality websites exceeds a threshold, it will automatically perform alternative queries for better quality sites to replace those low-quality sites. According to Moz’s comparison report on winners and losers of Panda Update [15], the winners usually had a smaller number of advertisements on their websites.

2.2 Content Farms

A content farm is a website that generates a large number of low-quality articles on many different topics. The underlying business model is to generate revenue by placing advertisements on the pages or selling contextual links within the content [15-16]. The idea is that the more network traffic a web page receives, the more revenue the page may generate. To this end, content farms usually produce content pages that not only rank well in search engines but also attract user clicks with exaggerated titles and lovely images.

Some issues arise after we examined several Google display ads performance reports for advertising campaigns running from 2015 to 2016. We found that about 8 percent of the web pages on which display ads appeared were content farms, and some content farms had continual clicks and conversions for two years. Several research questions, for example, why these content farms continue to attract users’ attention, and whether the clicks lead to valuable purchase conversions or others are all worth studying.

2.3 Advertising Performance Metrics

Running online advertising with Google Ads usually needs to set up campaigns and select a campaign type (e.g., Search, Display, or Video) for each campaign, then create ad groups under each campaign as well as ads within each group. Generally, each search ad need to create an ad text and a landing page, and choose several relevant keywords with a matching option and a bid for each keyword. When a search query matches an ad keyword, the corresponding ad may appear on the search engine results page. Display ads require creating ads as images in different sizes or HTML5, targeting options and set a default bid at the ad group level or a bid for an individual targeting option. Ads can appear on websites and mobile apps that are part of the Google Display Network.

Commonly used advertising performance metrics are click-through rate, conversion rate, average cost per click, and average cost per acquisition. Click-through rate (CTR) is the ratio of clicks to impressions: CTR = clicks/impressions; conversion rate (CVR) is the ratio of conversions to clicks: CVR = conversions/clicks, where a conversion can be signup, subscribing newsletter, placing an order, or others; average cost per click (Avg. CPC) is the average cost per click advertisers pay: Avg. CPC = cost/clicks; average cost per acquisition (Avg. CPA) is the amount advertisers pay to attain a conversion, i.e., Avg. CPA = cost/conversions.

Some studies suggested that using the above metrics alone does not lead to accurate performance evaluations because of high conversion rates for ads on expensive products, hard-to-trace offline conversions, long lead times, and attribution complexity [17-18]. For brand websites, the number of website visits is a more reliable performance indicator rather than the clicks [17]. Kireyev et al. [18] found that display ads helped increase search conversion rates as well as the number of clicks and the total cost. Barford et al. [19] found that more than 80 percent of the websites their crawler visited used some targeting mechanisms in their advertising setting; that is, the ads appearing on the web pages highly correlated with the simulated user profile (e.g., browsing interests, and cookies). Currently, Google Ads offers a remarketing display ads, which show ads only to some past visitors to the advertiser's website. Therefore, the remarketing display ads are likely to perform far better than general display ads.

3 Research Objective and Methods

This study intended to answer this research question: Is there a difference in mean ads performance for differing the type of campaign and the category of placement? The ads performance variables of interest are click-through rate (CTR), overall conversion rate (CVR), and purchase conversion rate (purchase CVR). To this end, we went through two main steps: pre-processing of the dataset and conducting two-way analyses of variance on specified performance variables with two factors: campaign and placement.

3.1 Data Pre-processing

This study obtained ten placement performance reports (https://developers.google.com/adwords/api/docs/appendix/reports/placement-performance-report) of display ads run in 2016 by a local online advertising agency. Each of these reports comes from an individual advertiser's account that has tracked conversions happening on their destination web pages. On an individual report, each row shows all statistics of performance data like impressions, clicks, CPC, among other metrics aggregated for individual URL (called a placement) where the advertiser's display ads have appeared. According to Google Ads definition (https://support.google.com/google-ads/answer/2470108?hl=en), a placement can be a website, a specific page on a site, a Mobile App, or a piece of video content.

The descriptive statistics of the data revealed that the distributions of clicks, CTR and CVR were all extremely positive skewed, indicating a majority of the data are under their individual mean and on the right-hand side there are some abnormal data samples such as very high conversion rates from low volume clicks. We were intended to highlight the placements where a proper volume of clicks and some conversions have appeared. Thus we conducted pre-processing of the data to make them ready to the test against the normality assumption and the transformation to normal distribution if needed.

The data pre-processing steps for each report are as follows:

(1) Keep rows that contribute the 80 percent of the total number of clicks by Pareto analysis.

(2) Remove rows without any conversions.

(3) Remove rows which are extreme outliers in CTR or CVR.

As an example, Fig. 1 shows the histograms of CVR data before and after pre-processing.

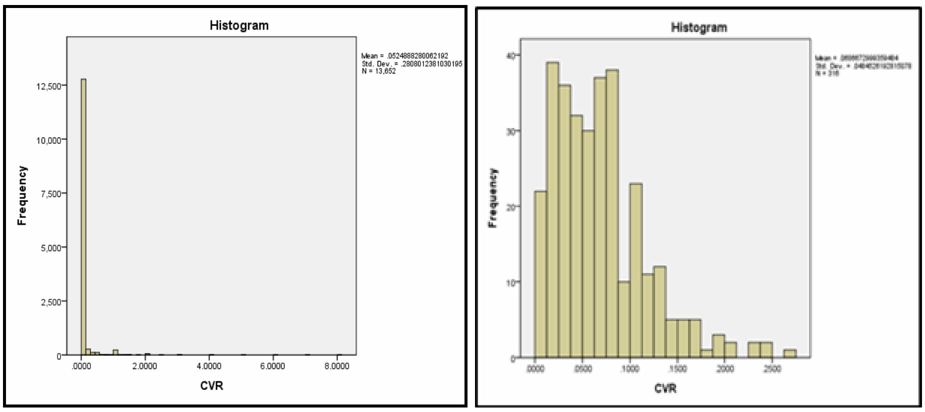

Fig. 1. Before and after pre-processing of CVR data

Then we classify the placements into content-farm (CF), non-content-farm (NonCF), and anonymous (Unknown) using the content farm classifier we developed in a previous study [20]. The Unknown placements include anonymous sites (e.g., ***.anonmous.google) and mobile apps (e.g., mobileapp::1-*********), which are not the focus of this study. The classification proceeds as follows:

(1) Mark anonymous sites and mobile apps with the Unknown label
(2) Conduct feature extraction for each of the remaining placements
(3) Classify each placement as CF or NonCF based on its feature values
(4) Manually double check the predicted CF placements which are new to our content farm blacklist
(5) Append confirmed CF placements to our content farm blacklist

3.2 Two-way Analysis of Variance on Ads Performance

This study conducted two-way ANOVA comparing ads performance for different combinations of campaign and placement. The factor, campaign, has two levels: RMKT (short for Remarketing) and GDN (short for Google Display Network) based on the campaign's criteria type. RMKT is a specific type of display campaigns which show ads only to the previous visitors while GDN show ads to general visitors of the advertiser's website. The other factor, placement, has three levels: CF, NonCF, and Unknown, as described above.

To conduct cross-account analysis, we transformed the dependent variables, CTR and CVR, of each report to the standard normal distribution with a mean of zero and a standard deviation of one. The inverse normal transformation using SPSS [21-22] proceeds as follows:

(1) Convert the values of a variable into fractional ranks by selecting Rank Cases from the Transform option.

(2) Transform the fractional ranks into normally distributed values with a mean of zero and a standard deviation of one by using the Idf.Normal formula from the Inverse DF function group in the Compute Variable form under the Transform option.

(3) Use the Kolmogorov-Smirnov (K-S) test to confirm whether the transformed variables achieve statistical normality or not.

We then prepared the cross-account treatment groups as the six combinations of two campaign levels and three placement levels. The data of the treatment groups were normalized as well as described above, but this time we used the original mean and standard deviation of a variable for each treatment group as the desired mean and standard deviation at the second step. That is, the mean and standard deviation of each variable for each treatment group remains the same after transforming to statistical normality.

We proceeded to perform several two-way analyses of variance on these treatment groups. Each two-way analysis of variance was to determine whether there is a significant interaction between campaign and placement on one single ad performance variable (CTR, CVR, or purchase CVR). This study set the significance threshold at .05. If there is a statistically significant interaction between campaign and placement on some variable, we then conduct the simple main analysis to compare the group means of that variable. Simple main analysis is equivalent to a one-way analysis of variance (ANOVA) that determines the mean difference on a variable between campaigns at each level of placement, as well as between placements for each type of campaign. If any of the ANOVA results yielded a significant difference in the group means, we then conduct post hoc tests to make multiple comparisons of the observed means. Alternately, if no statistically significant interaction exists, we then examine the main effect for each factor on a specific performance variable. For example, we would compare the marginal means of CTR between the two campaign levels to see if there are differences in CTR due to differing the campaign level.

4 Experimental Results

4.1 Content Farm Detection

One of our previous studies framed the content farm detection as a classification problem - predicting a discrete class label (CF or NonCF) for a given URL [20]. The Weka SMO based classifier used data from the principal features of domain authority, the maximal number of ads on any page, URL match degree between pages, and the use of shopping cart to make the decision. A ten-fold cross-validation test on a training set of 1000 websites (including 250 content farms) yielded 93.6%, 93.8%, and 93.7% on recall, precision, and F1 measure, respectively.

The classifier made 954 positive predictions of content farms from 86,385 unique hostnames extracted from the URLs in the dataset. Here a hostname refers to the domain name part (including subdomain) of a URL address, for example in a typical URL "http://www.example.com/path_to_file," the hostname is "www.example.com." Among the 954 predicted content farms, 234 hostnames were already in our CF blacklist, and 540 hostnames out of the remaining 720 positive predictions were true positives after manual verification. In short, the classifier reached a precision of 81 percent (=774/954) in content farm detection. We expanded our CF blacklist with the newly found content farms to a total of 790 hostnames.

4.2 Exploratory Statistical Analysis of Placements

This section examined some basic statistical measurements of the placements present in the dataset. Table 1 shows the percentage shares of hostnames, impressions, clicks and conversions segmented by placement category.

Table 1. Statistics of placements segmented by category

Placement	Hostnames	Impressions	Clicks	Conversions
CF	7.48%	16.78%	33.28%	26.58%
NonCF	77.85%	49.01%	60.23%	68.18%
Unknown	14.67%	34.21%	6.49%	5.24%

We can see that hostnames in the CF category account for only 7.48 percent of the total hostnames, but their accumulated impressions and clicks comprise 16.78 percent and 33.28 percent of the overall impressions and clicks, respectively. The reason seems to that content farms generally allocate more slots for display ads on their pages, and the more impressions they have, the more clicks on the ads they may generate. We further refined the data into Table 2 and found that content farms yielded 0.37 percent higher CTR than non-content farm websites; however, they produced 0.77 percent lower CVR than the latter. Obviously, display ads shown on content farm websites had higher chances to get clicks but had fewer opportunities leading clicks to conversions than on regular websites. Notably, Unknown placements had spared clicks, resulting in very low CTR, but the clicks together made up a well-performed CVR.

Table 2. Overall CTR and CVR by placement category

Placement Level	CTR	CVR
CF	0.97%	3.28%
NonCF	0.60%	4.65%
Unknown	0.09%	3.31%

4.3 Two-Way ANOVA on CTR

This section presented the two-way ANOVA results of examining the effects of campaign and placement on the click-through rate (CTR) of display ads. The interaction plot (Fig. 2) reflected that campaigns had higher CTR when ads appeared on CF and NonCF placements than on Unknown placements, and perhaps there were slight differences between ads shown on CF and NonCF placements. Additionally, RMKT campaigns had higher CTR than GDN campaigns, and perhaps the campaign difference was smaller at the Unknown level of placement.

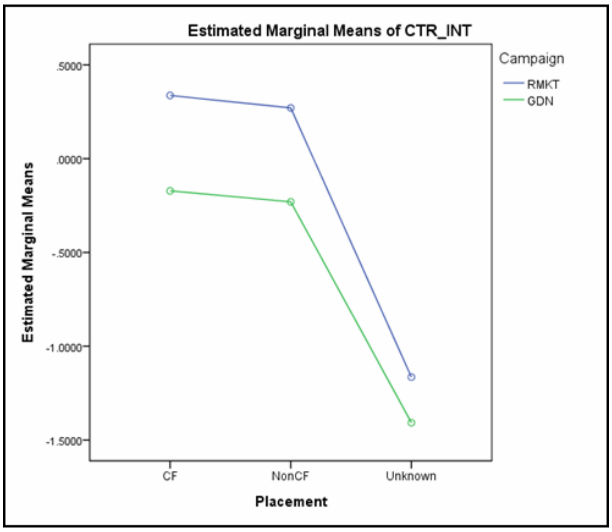

Fig. 2. Interaction plot for estimated means of CTR

The main ANOVA table (Table 3) revealed that the interaction effect between campaign and placement on CTR was significant, F(2, 5661) = 4.771, $p < .001$, indicating that the campaign difference in CTR depends on the particular placement level. Simple main effects analysis showed that RMKT campaigns yielded significantly higher CTR than GDN campaigns at all levels of placement ($p < .001$, $p < .001$, $p < .001$, respectively). Additionally, for RMKT campaigns, ads appeared on both CF and NonCF placements resulted in significantly higher CTR than those shown on Unknown placements ($p < .001$, $p < .001$ respectively), but there seemed to be no differences in CTR between ads placed on CF placements and those on NonCF placements ($p = .048$) because the p-value is very close to the significance threshold. Similarly, for GDN campaigns, ads appearing on both CF and NonCF placements produced significantly higher CTR than those shown on Unknown placements ($p < .001$, $p < .001$, respectively), but there were no differences in CTR between ads shown on CF placements and those on NonCF placements ($p = .42$). Table 4 summarizes the results of simple main effects analysis.

Table 3. Two-way ANOVA table for CTR

Dependent Variable: CTR_INT					
Source	Type III Sum of Squares	df	Mean Squares	F	Sig.
Corrected Model	1198.070	5	239.614	302.967	.000
Intercept	473.053	1	473.053	598.126	.000
Campaign	132.566	1	132.566	167.616	.000
Placement	789.202	2	394.601	498.932	.000
Campaign*Placement	7.546	2	3.773	4.771	.000
Error	4477.241	5661	.791		
Total	5675.774	5667			
Corrected Total	5675.311	5666			

Table 4. Summary of simple main effects on CTR

Factor	Levels	Post-Hoc Test
Placement	CF	RMKT > GDN
	NonCF	RMKT > GDN
	Unknown	RMKT > GDN
Campaign	RMKT	(CF ≈ NonCF) > Unknown
	GDN	(CF ≈ NonCF) > Unknown

4.4 Two-Way ANOVA on CVR

This section presented the results of a two-way ANOVA evaluating the interaction effects of campaign and placement on the overall conversion rate (CVR) of display ads. Examination of the interaction plot

(Fig. 3) suggested that campaigns had higher CVR on NonCF placements, and perhaps yielded poor CVR on CF placements; additionally, RMKT campaigns had higher CVR than GDN campaigns, and the campaign difference was smaller at the Unknown level of placement.

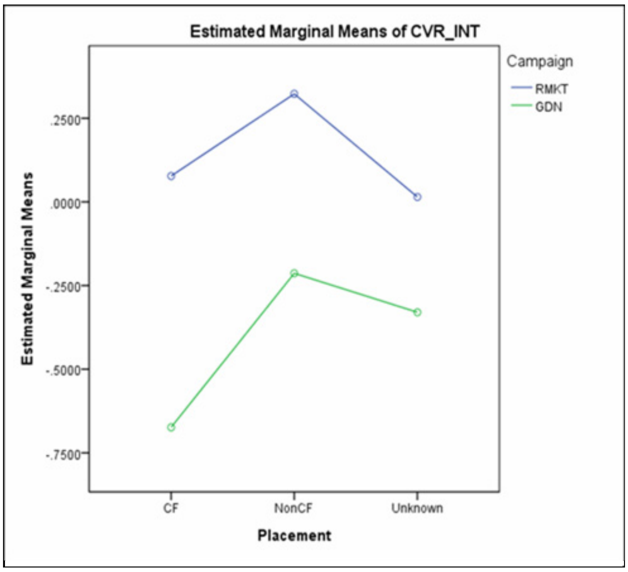

Fig. 3. Interaction plot for estimated means of CVR

The ANOVA results (Table 5) confirmed that the interaction effect between campaign and placement was significant, *F(2, 5661)= 11.631, p < 0.001*, indicating that the campaign difference in overall CVR depends on the particular level of placement. Simple main effects analysis revealed that RMKT campaigns yielded significantly higher CVR than GDN campaigns at all levels of placement ($p < .001$, $p < .001$, $p < .001$, respectively). Meanwhile, for RMKT campaigns, ads appeared on NonCF placements resulted in significantly higher CVR than those shown on the other levels of placements ($p < .001$, $p < .001$, respectively), but there were no differences in CVR between ads placed on CF placements and those on Unknown placements ($p = .57$). Contrastingly, for GDN campaigns, ads shown on both NonCF and Unknown placements produced significantly higher CVR than those appeared on CF placements ($p < .001$, $p < .001$, respectively), but there was no difference between ads placed on NonCF placements and those on Unknown placements ($p = .32$). Table 6 sums up the results of simple main effects.

Table 5. Two-way ANONA table for conversion rate

Dependent Variable: CVR_INT					
Source	Type III Sum of Squares	df	Mean Squares	F	Sig.
Corrected Model	600.103	5	120.021	133.645	.000
Intercept	54.310	1	54.310	60.475	.000
Campaign	224.974	1	224.974	250.513	.000
Placement	137.316	2	68.658	76.452	.000
Campaign*Placement	20.891	2	10.445	11.631	.000
Error	5083.890	5661	.898		
Total	5684.499	5667			
Corrected Total	5683.993	5666			

Table 6. Summary of simple main effects on CVR

Factor	Levels	Post-Hoc Test
Placement	CF	RMKT > GDN
	NonCF	RMKT > GDN
	Unknown	RMKT > GDN
Campaign	RMKT	NonCF > (CF ≈ Unknown)
	GDN	(NonCF ≈ Unknown) > CF

4.5 Two-Way ANOVA on Purchase CVR

This section presented the results that we further examined the effects of campaign and placement on the purchase conversion rate of display ads. To obtain the data on purchase conversions, we need new versions of the placement performance reports segmented by conversion category. Each segmented report downloaded from Google Ads backend has more than one row per placement; more specifically, each row showing all aggregated statistics of conversions for a combination of placement, campaign name, and conversion category. A segmented report differs from a regular one mainly in two folds: first, it aggregates the statistics of conversions by the domain name (e.g., example.com), instead of full URL (e.g., http://www.example.co/path_to_file); second, each of its rows has data on conversion rate and cost per conversion, but no data provided on clicks. Thus, we need to perform some calculation to get the average conversion rate for each conversion category.

The calculation started with a simplification of the many conversion categories into three main categories: Purchase/Sale, Signup and Others. We then calculated the average conversion rate for each conversion category by the following steps:

(1) Filter out data rows with zero conversion.
(2) For each remaining row, divide conversions by conversion rate to get clicks.
(3) Group rows by conversion category and get the sums of conversions and clicks for each category.
(4) For each category, get the average conversion rate by dividing its total conversions by total clicks.

The statistics of conversion category across accounts in Table 7 indicated that more than 50 percent of conversions were making purchases; however, its average conversion rate was only 3.07 percent, far below that of the other two conversion categories. We further classified each placement into three different levels with the following rules: (1) Assign a placement that begins with either "anonymous.google," or "mobileapp" as "Unknown." (2) Assign a placement as "CF" if it matches one domain name on our CF blacklist; otherwise assign it as "NonCF". Table 8 shows the statistics of purchase conversions by placement category. We can see that purchase conversions occurred from many more NonCF domains with a better average conversion rate.

Table 7. Statistics of conversion category across accounts

Conversion category	Conversions share	CVR
Purchase/Sale	51.64%	3.07%
Signup	38.68%	5.01%
Others	9.68%	7.22%

Table 8. Summary of purchase conversions by placement level

Placement level	Domain share	Purchase conversion rate
CF	12.37%	2.64%
NonCF	80.31%	3.39%
Unknown	7.32%	1.66%

We then conducted a two-way ANOVA to understand whether the campaign differences on purchase CVR depends on the level of placement. We organized and normalized the six treatment groups as described above. Examination of the interaction plot (Fig. 4) revealed that the possibility of a significant interaction exists indicating the differences in purchase CVR between different levels of campaign depend on the level of placement. More specifically, the placement affects RMKT differently than GDN campaigns at NonCF and CF levels, but there were slight differences between RMKT and GDN campaigns at the Unknown level.

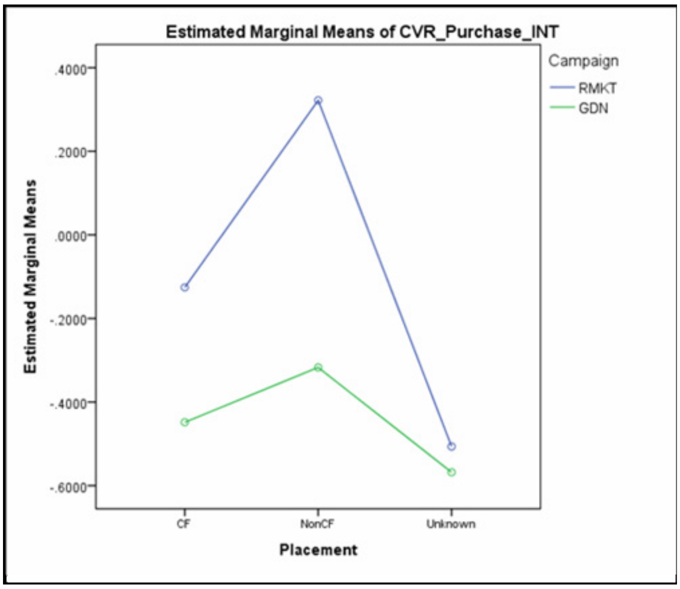

Fig. 4. Interaction plot for estimated means of purchase conversion rate

The ANOVA results (Table 9) confirmed that there was a statistically significant interaction between the effects of campaign and placement on purchase CVR, $F(2, 2623)=11.140$, $p < .001$. Simple main effect analysis showed that RMKT campaigns yielded significantly higher purchase CVR than GDN campaigns at the NonCF and CF levels of placement ($p < .001$, $p < .001$, respectively), but the purchase CVR of RMKT and GDN campaigns did not significantly differ at the Unknown level ($p = .47$). Additionally, ads in RMKT campaigns produced significantly greater purchase CVR on NonCF placements than on CF ($p < .001$) and Unknown placements ($p < .001$), and yielded significantly greater purchase CVR on CF placements than on Unknown placements ($p < .001$). Similarly, for GDN campaigns, ads shown on NonCF placements had significantly higher purchase CVR than on Unknown placements ($p = 0.021$), but the purchase CVR did not differ significantly between ads on CF and NonCF or Unknown placements ($p = .19$, $p = .41$, respectively). We summarized the results in Table 10.

Table 9. Two-way ANONA table for purchase conversion rate

Dependent Variable: CVR_Purchase_INT					
Source	Type III Sum of Squares	df	Mean Squares	F	Sig.
Corrected Model	314.259	5	62.852	70.451	.000
Intercept	98.115	1	98.115	109.979	.000
Campaign	36.023	1	36.023	40.378	.000
Placement	75.942	2	37.971	42.562	.000
Campaign*Placement	19.877	2	9.939	11.140	.000
Error	2340.059	2623	.892		
Total	2654.657	2629			
Corrected Total	2654.317	2628			

Table 10. Summary of simple main effects for purchase conversion rate

Factor	Levels	Post-Hoc Test
Placement	CF	RMKT > GDN
	NonCF	RMKT > GDN
	Unknown	RMKT ≈ GDN
Campaign	RMKT	NonCF > CF > Unknown
	GDN	NonCF > Unknown

4.6 Discussions

The results of two-way ANOVA tests on CTR, overall CVR, and purchase CVR all indicated that the two main effects for campaign and placement, respectively and the interaction effect between campaign and placement were statistically significant.

The significant main effects revealed that there were significant differences in performance metrics between RMKT and GDN campaigns and between placements at CF, NonCF, and Unknown levels That is, RMKT campaigns consistently over-performed GDN campaigns in all three performance metrics of interest no matter at what level of placements their ads appeared. Also, ads appearing on CF placements had better CTR than their counterparts on NonCF placements; while ads shown on NonCF placements had better CVR (overall and purchase) than CF ones. Finally, ads on the Unknown placements often underperformed than those on other placements.

The simple main effects indicated that RMKT campaigns yielded better CTR on CF placements, but had better CVR (both overall and purchase) on NonCF placements; GDN campaigns had better overall CVR on NonCF placements and did not perform differently in CTR and purchase CVR when their ads appeared either on CF or NonCF placements.

Google display network (GDN) does provide mechanisms for advertisers to have more control over which webpages, videos, and apps in the Display Network to show their ads. We can select the placements (for example NonCF placements for better CVR) as “managed placements.” The advertisers can also add CF placements in a placement exclusion list so Google will not show their ads on these excluded placements.

5 Conclusions

This study examined the effect of campaign and placement on the performance of display ads using two-way ANOVA. The performance metrics of interest were: click-through rate, overall and purchase conversion rates. The results of two-way ANOVA showed that there was an interaction between the effects of campaign and placement on each of the performance metrics indicating that the campaign difference in ads performance depends on the particular placement level. More specifically, RMKT campaigns achieved larger CTR but smaller CVR on CF placements than on NonCF placements. Contrastingly, GDN campaigns yielded no significant differences in CTR and purchase CVR between CF and NonCF placements but had smaller overall CVR on CF placements than on NonCF ones. For Unknown placements, the results were mixed as campaign type changed. Additionally, RMKT campaigns performed significantly better than GDN campaigns at all levels of placement; that is, campaign difference had a significant impact on ads performance.

We concluded to suggest that: (1) If the focus is on getting clicks to achieve sufficient click-through rate, then running display ads campaigns (RMKT for past visitors and GDN for new visitors) and delivering ads on any types of websites; (2) If obtaining conversions is the primary goal, then running remarketing display campaigns and delivering ads only on non-content farm websites.

References

[1] eMarketer, Google still dominates the world search ad market. <https://www.emarketer.com/Article/Google-Still-Dominates-World-Search-Ad-Market/1014258>, 2016 (accessed 03.03.17).

[2] eMarketer, US digital ad spending to surpass TV this year. <https://www.emarketer.com/Article/US-Digital-Ad-Spending-Surpass-TV-this-Year/1014469>, 2016 (accessed 03.01.17).

[3] P. Bakker, Aggregation, content farms and Huffinization: the rise of low-pay and no-pay journalism, Journalism Practice 6(5-6)(2012) 627-637.

[4] A. Carr, Low quality websites: content farms, ACC library guides. <http://researchguides.austincc.edu/contentfarms> (accessed 15.12.16).

[5] Content farm. In Wikipedia, The Free Encyclopedia. <https://en.wikipedia.org/wiki/Content_farm> (accessed 11.11.16).

[6] Google, Search quality rating guidelines. <https://static.googleusercontent.com/media/www.google.com/en//insidesearch/howsearchworks/assets/searchqualityevaluatorguidelines.pdf > (accessed 11.12.16).

[7] Google, Webmaster guidelines. <https://support.google.com/webmasters/answer/35769> (accessed 11.12.16).

[8] N. Spirin, J. Han, Survey on web spam detection: principles and algorithms, ACM SIGKDD Explorations Newsletter 13(2)(2012) 50-64.

[9] A. Ntoulas, M. Najork, M. Manasse, D. Fetterly, Detecting spam web pages through content analysis, in: Proc. the 15th International Conference on World Wide Web, 2006.

[10] L. Becchetti, C. Castillo, D. Donato, R. Baeza-Yates, S. Leonardi, Link analysis for web spam detection. ACM Transactions on the Web 2(1)(2008) Article No.2.

[11] T. Urvoy, E. Chauveau, P. Filoche, T. Lavergne, Tracking web spam with html style similarities, ACM Transactions on the Web (TWEB) 2(1)(2008) Article No.3.

[12] C. Dong, B. Zhou, Effectively detecting content spam on the web using topical diversity measures, in: Proc. the 2012 IEEE/WIC/ACM International Joint Conferences on Web Intelligence and Intelligent Agent Technology, 2012.

[13] A. Singhal, More guidance on building high-quality sites. <https://webmasters.googleblog.com/ 2011/05/more-guidance-on-building-high-quality.html> (accessed 10.12.16).

[14] N. Panda, A.R. Lehman, T.G. Upstill, Selectively generating alternative queries, Google Patents, US Patent 9,135,307. <https://www.google.com/patents/US9135307>, 2012 (accessed 20.12.16)

[15] Fishkin, R. Google's Farmer/Panda Update: Analysis of Winners vs. Losers. <https://moz.com/blog/googles-farmer-update-analysis-of-winners-vs-losers> (accessed 11.12.16).

[16] Techopedia Inc. Content Farm. <https://www.techopedia.com/definition/16897/content-farm> (accessed 29.12.16).

[17] B. Dalessandro, R. Hook, C. Perlich, F. Provost, Evaluating and optimizing online advertising: forget the click, but there are good proxies, Big Data 3(2)(2015) 90-102.

[18] P. Kireyev, K. Pauwels, S. Gupta, Do display ads influence search? attribution and dynamics in online advertising. International Journal of Research in Marketing 33(3)(2016) 475-490.

[19] P. Barford, I. Canadi, D. Krushevskaja, Q. Ma, S. Muthukrishnan, Adscape: harvesting and analyzing online display ads, in: Proc. the 23rd international conference on World Wide Web, 2014.

[20] A.C. Wu, C.J. Luh, Identifying content farms and evaluating their display ads performance, in: Proc. the 28th International Conference on Information Management, 2017

[21] A. Field, Discovering Statistics Using IBM SPSS Statistics, 4th Edition, Sage, 2013.

[22] G. F. Templeton, A two-step approach for transforming continuous variables to normal: implications and recommendations for is research, Communications of the Association for Information Systems 28(4)(2011) 41-58.