

Fatigue-Aware Ad Creative Selection

Daisuke Moriwaki

Komei Fujita

Shota Yasui

moriwaki_daisuke@cyberagent.co.jp

fujita_komei@cyberagent.co.jp

yasui_shota@cyberagent.co.jp

CyberAgent, Inc.

Takahiro Hoshino

bayesian@jasmine.ocn.ne.jp

Keio University

RIKEN Center for Advanced Intelligence Project

ABSTRACT

In online display advertising, selecting the most effective ad creative (ad image) for each impression is a crucial task for DSPs (Demand-Side Platforms) to fulfill their goals (click-through rate, number of conversions, revenue, and brand improvement). As widely recognized in the marketing literature, the effect of ad creative changes with the number of repetitive ad exposures. In this study, we propose an efficient and easy-to-implement ad creative selection algorithm that explicitly considers user's psychological status when selecting ad creatives. The proposed system was deployed in a real-world production environment and tested against the baseline algorithms. The results show superiority of the proposed algorithm.

CCS CONCEPTS

• Information systems → Display advertising; Personalization; • Applied computing → Marketing.

KEYWORDS

Display advertising, Real-Time Bidding, Advertising fatigue, Advertising repetition, Experiment, User Modeling

ACM Reference Format:

Daisuke Moriwaki, Komei Fujita, Shota Yasui, and Takahiro Hoshino. 2020. Fatigue-Aware Ad Creative Selection. In *SUM '20: State-based User Modelling*, Feb 07, 2020, Houston, TX. ACM, New York, NY, USA, 8 pages.

1 INTRODUCTION

In online display advertising, it is a common practice to prepare multiple ad creatives (a combination of an image and text or a video, see Figure 1 for example) for each advertising campaign. The DSP (Demand-Side Platform) selects the most effective ad creative from the candidates.

Since new ad creatives are continuously added every day, every hour in the typical online advertising scene there is little time to test the effectiveness of each ad creative beforehand. The *bandit algorithm* [13] is best suited to such situation. The bandit algorithms



Figure 1: Variation of ad creatives for a mobile game title. The messages are all different: “Don’t miss the special event!” (left), “Have you already tried these coupons?” (center), and “Waiting for your return” (right). Users jump to login screen of the game when click on these ads.

give us a way to choose one from multiple choices under uncertainty of the outcome by balancing exploration and exploitation.

The effectiveness of ad creative depends not only on the quality of ad creative but also on the audience. For example, one might prefer an informative ad creative while another like an emotional ad creative. The DSP, therefore, often considers *context* as well as the ad creative. The bandit algorithm that learn contexts is called *contextual bandit* and have been extensively studied at a practical level [3, 7, 18, 28].

One serious side-effect of the contextual bandit is that the ad selection becomes deterministic once the algorithm has *learned* enough volume of data. The basic bandit algorithm, regardless of its implementation (e.g. UCB, Thompson sampling, or ϵ -greedy), assumes the reward of arms are unchanged. The degree of exploration is decreasing with the number of plays and eventually goes to zero. As a result, targeted users keep watching same ad creative that is considered as the best by the algorithm given users’ information.

The marketing science literature has been extensively studying the consequences of repetitive ad exposure. A seminal work by Zajonc [31] shows mere repeated exposure can incur positive attitude among audience from the experiment results. On the other hand, many researchers point out there is a threshold above which additional ad exposure has negative effect for users’ attitude due to boredom or advertising fatigue. As surveyed in [21, 25], *Two-factor model* theory is widely supported by the researchers in this domain. This theory claims that repetitive advertising has two phases: in the early stage the attitude toward the product improves with the number of repetition but at some point advertising fatigue dominates the positive effect and the net effect (positive effect minus negative effect) plummets (Figure 2).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SUM '20, Feb 07, 2020, Houston, TX

© 2020 Association for Computing Machinery.




Figure 2: A graphical image of "Two-factor model" for advertising repetition. Previous researches have found that the repetitive advertising has both positive and negative effects on consumers' attitude toward advertised products and the net effect is described as an inverted U-shaped curve[25].

The over-exposure of same ad is especially prevalent in display advertising. The real-time bidding (RTB) enables advertisers to reach out to targeted users whenever user browses web or use apps with advertising slots. The distribution of the counts of each user-ad creative pairs from real RTB data (Figure 3) shows that 47.5% of creative-user pairs appear multiple times within 24 hours, which leads that 64.6% of the users experience multiple exposure to same ad creative.




Figure 3: Number of exposure to same ad creative for each user within 24 hours. $N = 27,750$

Recently the concept of advertising fatigue has received growing attention from machine learning and data mining researchers[1, 2, 11, 19]. In particular, Agarwal et al. (2009)[2] and Ma et al. (2016)[19] proposed a fatigue-aware news recommendation system which evidently improved upon existing algorithm. Similar idea is applied to LinkedIn recommendation system[14]. The DSPs also need to consider the psychological status of the audience when choosing an ad creative. To capture users' perception toward each ad creative, we introduce a quantitative measure of advertising fatigue. When we have only one ad creative in hand, fatigue will be well proxied by the number of ad exposures in the past. However, it is not the case in display advertising. When we have multiple creatives, the past exposure to other creatives should be considered because the creatives more or less resemble to each other and expected to increase fatigue in a similar way. We conjecture the gain in fatigue is correlated with the degree of similarity of the ad creative on user's screen and what he has seen in the past. To get a more concrete idea, let's assume user's attitude toward ad creative evolves according to Figure 2. Then the magnitude of the improvement of attitude by the

additional ad exposure (=marginal effect of ad exposure) decreases as user's advertising fatigue increases (Figure 4). Suppose that a user has seen same ad creative repeatedly then the effect of the ad creative is negative due to the high level of fatigue (point A). In this case, we can restore the effect of advertising by using a new ad creative (point B). The new ad creative can be a slight modification of the original or totally different ones. A typical set of creatives is already shown in Figure 1. That is, the DSP needs to switch ad creative when user gets high level of fatigue.




Figure 4: The level of fatigue and effect of additional ad exposure (a graphical image). According to the two-factor model, user who has seen same ad creative many times gets advertising fatigue and the effect of ad exposure is negative (A). But new ad creatives can reduce fatigue level and restore effectiveness (B).

We propose a new algorithm for ad creative selection for DSP that explicitly considers advertising fatigue. The algorithm not only selects the best ad creative based on the context but also estimate the level of fatigue for each user and take it into account.

Since the display advertising environment has strict latency requirements, the DSP server can spare only a few milliseconds for computation since it should deliver thousands of ad creatives to numerous users within unnoticeable duration. To save as much as computation time, we first develop a very simple calculation formula for advertising fatigue which capture the similarity of current ad creative and past ad exposure. Then we propose the implementable algorithm that add fatigue features to the existing reward predictor. To check the feasibility and the performance, we conducted an online experiment on real-world display advertising. The result shows that adding fatigue feature improves CTR (click-through rate, the percentage of being clicked out of impressions) and CVR (conversion rate, the percentage of users that complete desired goal such as log-in) over existing algorithm without any systemic problem. We also conducted post-experiment analysis and find the algorithm indeed reduced the fatigue. Therefore, our experiments show an apparent relationship between fatigue and CTR/CVR.

The rest of the paper is as follows. In the next section, we review existing work in Section 2. Then we formalize the problem and explain the existing system in Section 3. Section 4 proposes the new algorithm that takes the advertising fatigue into account. Section 5 describes our online experiment and the results. Section 6 conducts post-experiment analysis. Finally, Section 7 concludes.

2 RELATED WORKS

We review related works in both marketing science and machine learning.

2.1 User fatigue in Machine Learning Literature

Agarwal et al. (2009)[2] perform explanatory analysis on user's repeat-exposure to recommended articles using data from *Yahoo!* front page. They incorporate decaying factors to their CTR predictor based on the analysis. Ma et al. (2016)[19] also conduct an extensive analysis for user's past exposure on data from *Bing* news recommendation engine. They propose to use users' past exposure and reaction to same news item and items from same category as features for the machine learning model. User fatigue is also considered explicitly in [1, 11]. The present paper extends their researches in a number of ways. First we extend the literature from recommendation to display advertising. Second, we define and calculate item similarity to measure fatigue. Third, we deployed our proposed algorithm in a production environment and proved its efficiency and feasibility. Not explicitly mentioning user fatigue, Lee et al. (2014)[14] improves recommendation algorithms in *LinkedIn* using user's past exposure to the recommended items (other members) and the actions to them.

2.2 Repetitive Ad in Marketing Science

The relationship between the number of repetitions and the effectiveness of ad have been actively studied in marketing science literature for a long time. Extensive reviews and meta-analysis point out that the optimal number of repeated ad exposures depends many conditions and thus is hard to determine[21, 25]. In the context of online advertising, an analysis of large scale *natural experiments* finds that wear-out occurs in a heterogeneous manner[17]. A lab experiment reveals that ad creatives with high quality is immune to wear-out effects[9].

2.3 Effect of Varied Ad

The effectiveness of varied ad creative has been examined in marketing science literature [8, 26, 29]. In particular, Chatterjee et al. (2003)[8] examined effectiveness of (i)repeating same ad creative and (ii)varying ad creative in lab experiments. The point estimate of CTR (click-through rate) of varied ad creative exceeded the same ad creative while their test statistics is not significant due to its small sample size.

2.4 Bandit Algorithms in Online Ad/Recommendation System

The bandit algorithms have been already applied to variety of advertising/recommendation system [7, 18, 27]. In particular, an application to ad format selection is closest to ours[28]. Our research is also related to the *non-stationary bandit* problem, in which the reward is changing with time and the optimal arm is not always the same. Among various algorithms, directly estimating the decaying factor of rewards are proved to be effective[12, 16]. We also consider continuous creation and deletion of arms, which is considered in [6]. While existing literature focus on relatively simple

mechanism of reward changing (e.g. decrease with time or pulls), we introduce a numerical measure of ad fatigue of users and tackle the problem. Several studies explicitly consider user abandoning due to the users' psychological status in contextual bandit[5, 15]. In these studies, profitable interventions are considered to have a negative effect on users psychological status and possibly cause user churn. In this scenario, a good contextual bandit algorithm needs to save the frequency of uncomfortable treatments to keep the users engaged. The proposed algorithms solve a constrained optimization problem to find the best policy [15] or calculate the optimal sequence of interventions [5] at each time of intervention. However, both studies are limited to the simulation and not applied to real production environment.

3 PROBLEM SETTING AND EXISTING AD CREATIVE SELECTION SYSTEM

In this section we introduce real-world problem setting and explain current deployed system.

3.1 Ad Creative Selection for Retargeting to Game Users

In this paper, we focus on a DSP that targets the churned game users who have stopped playing for a while. But the same idea can be applied to a variety of settings in advertising and recommendation. The DSP intensively buys ad views of these users through Real-Time Bidding (RTB) and urge them to return to the game. The example of ad creatives has been already shown in Figure 1, which promote new features of the game, issuing coupons, or simply persuading them to return. Some of the users click on the ad, jump to the login screen, and login the game (conversion).

To convert users, the DSP needs to select the most effective ad creative for each user. The existing system first predicts CTR for each ad creative based on contextual information and then chooses one with the highest predicted CTR.

Formally, each time t , the CTR r_t is predicted based on the context vector \mathbf{x}_t for ad creative a_t ,

$$\hat{r}_t(a_t) = \sigma(\boldsymbol{\theta}_0 \cdot \mathbf{x}_t + \boldsymbol{\theta}(a_t) \cdot \mathbf{x}_t), \quad (1)$$

where $\boldsymbol{\theta}_0$ is a d -dimensional weight vector that is invariant to ad creative, and $\boldsymbol{\theta}(a_t)$ is a d -dimensional *action-specific* weight vector. Both vectors include bias term. $\sigma(\cdot)$ represents the sigmoid (logistic) function. The context vector \mathbf{x}_t includes information from bid request (e.g. device, site/app and SSP), and hour of the day. The algorithm selects the ad creative with the highest expected reward.

One serious problem with the algorithm is that it does not generate good training data. The algorithm chooses same ad creative for same context. However, we need the results for the other ad creatives to train the model. The bandit algorithm is an established remedy to the problem that add randomization to selection algorithms. Following [7], the existing system adopts the Thompson sampling. That is, for each time t , action-specific weight vector $\boldsymbol{\theta}(a_t)$ is sampled from the distribution $\mathcal{N}(\boldsymbol{\mu}(a_t), \alpha \Sigma(a_t))$ for each candidate ad creative, where scalar $\alpha \in (0, 1]$ controls the degree of exploitation. Action-invariant vector $\boldsymbol{\theta}_0$ is set to $\boldsymbol{\mu}_0$ since it does not matter for selection.

Then the system chooses the best ad creative from candidates set \mathcal{A}_t , namely, $a_t^* = \arg \max_{a_t \in \mathcal{A}_t} \hat{r}_t(a_t)$. The action space \mathcal{A}_t is different across t . \mathcal{A}_t is determined based on the characteristics of user and ad-slot by the system. Moreover, the set of available ad creative changes over time. As explained in [6], ads are continuously added and deleted from circulation. The algorithm chooses newly-added ad creative with a probability of $1/|\mathcal{A}_t|$ for exploration. The whole procedure is described in Algorithm 1.

Algorithm 1 Baseline Ad Creative Selection Algorithm (Baseline)

Require: $\mathcal{N}(\mu(a), \alpha\Sigma(a))$
for $t = 1, \dots, T$ **do**
 Take feature vector \mathbf{x}_t ,
 for $a_t = 1, \dots, |\mathcal{A}_t|$ **do**
 if $\mu(a_t), \Sigma(a_t)$ is available **then**
 Sample $\theta(a_t) \sim \mathcal{N}(\mu(a_t), \alpha\Sigma(a_t))$
 Calculate $\hat{r}_t(a_t)$ according to eq.(1).
 else if $\mu(a_t), \Sigma(a_t)$ is NOT available due to lack of data **then**
 Assign $\hat{r}_t(a_t)$ very large value at the probability of $\frac{1}{|\mathcal{A}_t|}$
 and very small values at the probability of $1 - \frac{1}{|\mathcal{A}_t|}$
 end if
 end for
 Select $a^* = \arg \max_{a \in \mathcal{A}_t} \hat{r}_t(a)$
 Observe $r_t \in \{0, 1\}$
 if t is the last impression of the day **then**
 Update $(\mu(a), \Sigma(a))$ using data from the previous day
 end if
end for

3.2 Parameter Estimation

We estimate $(\mu(a_t), \Sigma(a_t))$ as posterior Gaussian distribution of the weight vector $\theta(a_t)$ of eq.1. That is, we first estimate $\theta = (\theta_0, \theta(a_t))$ as point estimates from L2-regularized logistic regression and derive $\Sigma_0, \Sigma(a_t) \in \Sigma$ since the inverse of Hessian of the objective function (negative log-likelihood plus L2-penalty) becomes variance-covariance matrix of the weight vector. Since the derivation is same as Chapell and Li (2011)[7] we do not detail here. It is basically using the result that the L2-regularized logistic regression is equivalent to MAP estimation and utilize Laplace transformation (see Bishop (2006)[4]). Due to the intractable size of our Hessian matrix, we only use diagonal elements of Hessian.

To estimate both common weight parameters μ_0 and action-specific $\mu(a)$, $a \in \mathcal{A}$ at once, we first expand feature vector \mathbf{x}_t by interacting a binary action vector $\mathbf{a} \in \{0, 1\}^{|\mathcal{A}|}$ whose k th element takes one when corresponding action is taken while the others take zeros. That is, for $\mathbf{a}_t = (0, 1, 0)^T$, $\mathbf{a}_t \otimes \mathbf{x}_t = (\mathbf{0}_d^T, \mathbf{x}_t^T, \mathbf{0}_d^T)^T$ where $\mathbf{0}_d$ denotes a zero vector of length d . Finally, by combining action-invariant component \mathbf{x}_t , we have $\mathbf{z}_t = (\mathbf{x}_t^T, \mathbf{a}_t \otimes \mathbf{x}_t)^T$. We use \mathbf{z}_t for both training and predicting procedures. The actual implementation is detailed in Section 5.

4 FATIGUE-AWARE AD CREATIVE SELECTION

As explained in Section 1, we need to consider user's psychological status when predicting reward. Formally, our *fatigue aware* reward predictor becomes

$$\hat{r}_{i,t}(a_t) = \sigma(\mathbf{x}_t \cdot \theta_0 + \mathbf{x}_t \cdot \theta(a_t) + b_1 \kappa_{i,t}(a_t) + b_2 \kappa_{i,t}(a_t)^2). \quad (2)$$

The equation 2 departs from eq. (1) in two ways. First, the new predictor uses the metric for the level of advertising fatigue $\kappa_{i,t}(a_t)$. $\kappa_{i,t}(a_t)$ measures the estimated ad fatigue when a user i is exposed to ad creative a_t at time t . The specific calculation of $\kappa_{i,t}(a_t)$ is detailed below. Second, the new predictor take into account user i while eq. (1) does not contain subscript i . In the sense, the fatigue-aware predictor can personalize the ad creative selection.

κ is included in the function as a quadratic form. However, both b_1 and b_2 can take any values. While the marketing literature find inverse-U shape relationship described in Figure 2, we **do not** presume the relationship between the two. The actual values of weights (b_1, b_2) for fatigue metric are learned from the data. With the new predictor the DSP server chooses ad creative a_t from a set of candidates \mathcal{A}_t . The whole procedure is Algorithm 2.

Algorithm 2 Fatigue-Aware Ad Creative Selection (FA)

Require: $\mathcal{N}(\mu(a), \alpha\Sigma(a)), \mathbf{h}_{i,t}, s$
for $t = 1, \dots, T$ **do**
 Take feature vector \mathbf{x}_t and history for i at t $\mathbf{h}_{i,t}$
 for $a_t = 1, \dots, |\mathcal{A}_t|$ **do**
 if Action specific distribution is available **then**
 get a vector of similarity scores for a_t , i.e. $\mathbf{s}(a_t)$
 Calculate $\kappa_{i,t}(a_t) = \mathbf{h}_{i,t} \cdot \mathbf{s}(a_t)$
 Sample $\theta(a_t) \sim \mathcal{N}(\mu(a_t), \alpha\Sigma(a_t))$
 Calculate $\hat{r}_t(a_t)$ according to eq.(2).
 else if Action-specific distribution is NOT available **then**
 Same as Algorithm 1
 end if
 end for
 Select $a^* = \arg \max_{a \in \mathcal{A}_t} \hat{r}_t(a)$
 Observe $r_t \in \{0, 1\}$
 if t is the last impression of the day **then**
 Update $\mathcal{N}(\mu(a), \alpha\Sigma(a)), b_1, b_2$ on data from previous day
 end if
end for

4.1 The Calculation of Fatigue

Here, we detail the calculation of fatigue metric κ . The basic idea is that the level of fatigue is determined by the number of past exposure to the same or similar ad creatives. Suppose the similarities among three ad creatives are as shown in Figure 5. Then the level of fatigue for each ad creative candidates can be sum of similarity between each candidate and creatives to which the user are exposed in the past. Figure 6 shows an example. This user has seen the blue ad creative three times and the yellow creative one time. Then the ad fatigue becomes approximately 3.4 (3 for the same creative and .39 for yellow creative) when he sees the blue.



Figure 5: An example of similarity matrix.



Figure 6: Calculation of fatigue. The values of fatigue depends on the history of ad exposure and chosen ad creative. In this case, blue ad creative has been seen three times and incur the highest fatigue.

Formally, let $\mathbf{h}_{i,t}$ be a vector representation of history of ad exposures to user i at time t . Each element in $\mathbf{h}_{i,t}$ represents the number of exposures to each creative. In the above example, $\mathbf{h}_{i,t} = (3, 1, 0)^T$. We record the histories for each user-advertiser pair. Hence, each user has multiple history vectors. The expected level of fatigue to ad creative a_t is calculated as a weighted sum of history: $\kappa_{i,t}(a_t) = \mathbf{h}_{i,t} \cdot \mathbf{s}(a_t)$, where the weight vector $\mathbf{s}(a_t)$ measure how similar a_t and the other creatives are. That is,

$$\mathbf{s}(a_t) = (s(a_t, 1), s(a_t, 2), \dots, s(a_t, |\mathcal{A}_t|))^T, s(j, k) \in [0, 1]. \quad (3)$$

4.2 Calculation of Similarity

Each similarity score $s(a, a')$ is calculated as a weighted average of *text similarity* and *image similarity*. The text similarity is calculated by the cosine similarity of the bag of words (BoW) representations of the description texts. We extracted the words from the texts by using Mecab[20], which is the most commonly used morphological analyzer for the Japanese language, with NEologd dictionary[24], which is a frequently updated Japanese dictionary containing up-to-date words. We used the gensim library[22] for the BoW calculation.

The image similarity is also calculated by the cosine similarity of the vector representation of the images. The vector representation is extracted from each image by using pre-trained MobileNetV2[23] implemented in Keras[10]. Here, we use the output from the last pooling layer as the representation. We put three times higher weight to the text similarity because the text similarity (i.e., BoW representation) is much easier to interpret than the image similarity (i.e., weight of the neural network).

In a real world online display advertising environment, complicated models such as DNN are hard to be implemented because

DSPs are required extremely short response time not to deteriorate user experience. By calculating similarity in advance, we can compress information from image and text into scalar variables it drastically reduces computing cost.

4.3 Implementation

The online advertising is an computationally complex environment. To deliver thousands of ad creatives to numerous users within unnoticeable duration, the system can spare only a few milliseconds to calculate fatigues for all the candidates of each impression. To save time, we first calculate the similarity scores for each pair of creatives off-line and store in the server. We update the histories for each user-advertiser pair in **real-time**. Owing to the extremely high frequency of impressions, recording all the ad creative exposures for all the users could lead to an explosion of the data volume, which needs to be capped at some point. Therefore, the system records the first impression of a user to each ad creative every minute. That is, if a user watches the same ad creative more than once within a minute, it is counted as one impression.¹ To further reduce the data volume and consider forgetting, the histories more than 24 hours before are deleted from the database. The whole system is described in Figure 7.

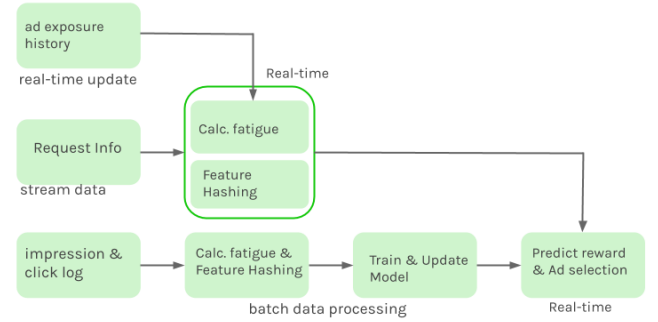


Figure 7: Summary of fatigue-aware ad creative selection system.

5 ONLINE EXPERIMENTS

Since the reward is changed by the past intervention by the system, offline evaluation is not very meaningful in this case. Notice that replay method needs the exact same intervention history to each user in this case. Hence we run an online experiment in a production environment to examine the effectiveness of the proposed algorithm. In this section, we first describe the setting of the experiment and then show the results.

5.1 Setting

We deploy our Fatigue-Aware (FA) selection algorithm in a DSP provided by CyberAgent, inc., a Japan-based major online advertising company. Three advertising campaigns for three different mobile game titles are selected for the experiment according to various conditions². A part of impressions for these campaigns are used

¹This data reduction affect less than 40% of the users.

²We used four campaigns but it turns out that one of them has only one creative.

for the experiment. The sample size of the training data reaches a maximum of one million a day with the hundreds of thousand of features. To deal with the size and the sparsity of the training data, we employ negative down-sampling and the hashing trick[30]. We set sampling rate for negative data to 5% and limit the length of features to 2^{24} for the system without fatigue parameters (Baseline) and $2^{24} + 2$ for FA by hashing-trick. The scalar parameter α that governs the degree of exploration of the algorithm are set to 0.01 for both FA and Baseline. With the small value of α , the Thompson sampling becomes *optimistic*[7]. We apply stochastic gradient decent (SGDC) classifier in sklearn package of Python) to obtain μ with regularization parameter $\lambda = 0.0011$. The parameter vector is updated daily basis using a batch of data from the previous day. These hyper-parameters are tuned using the replay method [18] for Baseline. FA shares all the above setting except for the inclusion of the fatigue parameter. To deal with cold start problem, FA is trained on data from the existing algorithm in the pre-experiment period. During the experiment, it learns only data that it produces.

Baseline (Alg.1) and FA (Alg.2), are tested against the result of randomly-chosen ad creative (Random) in a standard A/B testing. That is, they share the impressions equally based on the users' id and update their parameters based on the logs they generate. The experiment ran for one week.

5.2 Main Results

Table 1 shows the overall results of the experiments. Each metric is normalized with respect to the result of random algorithm (e.g., CTRs for FA and Baseline are divided by the CTR for Random). FA out-performed both baseline and random algorithm in CTR (click-through rate), CVR (conversion rate, the share of post-click conversions in all impressions), and Post-impression CVR including non-click conversions. The results clearly show that the proposed algorithm successfully increased both clicks and conversions. Baseline algorithm outperforms FA for post-click CVR. That is, while baseline collects less clicks than FA, baseline gets conversions more likely than FA once it gets click. But the difference is not statistically significant.

Alg.	Impressions	CTR	CVR	Post-Click CVR	Post-Imp CVR
FA	1,097,261	1.08*	1.12	1.04	1.09***
Baseline	1,081,393	1.04	1.10	1.05	1.04
Rand	1,087,931	1.00	1.00	1.00	1.00

Table 1: The results of the experiments. The numbers are normalized except for the number of impressions. Normalization is performed with respect to the random algorithm. The results are for the aggregate of all three campaigns. Superscripts implies FA is statistically better than Baseline. $P < 0.1$:*, $P < 0.05$:, $P < 0.01$:*****

5.3 Heterogeneity in campaigns

To further examine the results, Table 2 shows the CTRs for each campaign. The detailed results reveal that the baseline is not always superior to random. As for conversion rates, the performance of

the baseline is poor in campaign A. On the other hand, our fatigue-aware algorithm stably outperforms the other two. The instability of baseline possibly comes from the data volume as campaign A has least size of data. With combination of small size of training data and sparsity, the estimates could be very unstable. In real-world production environment, the data volume cannot be controlled by experimenters. Still, fatigue bandit algorithm successfully learned the data.

Alg	Campaigns		
	A	B	C
FA	1.21	1.04	1.32
Baseline	0.95	1.03	1.11
Rand	1.00	1.00	1.00

Table 2: CTRs for three campaigns. Numbers are again normalized with respect to random algorithm.

6 DISCUSSION

In this section, we further investigate the results focusing on the fatigue. Figure 8 shows the distribution of user fatigue at time that the algorithm chooses the ad creative. While the distributions of fatigue are similar among three campaigns, Campaign A has the mean highest similarity score (Table 3) and higher fatigue level. The distribution is dense around 0 to 10, which is equivalent to 0 to 10 times exposure to same ad creative within past 24 hours.




Figure 8: The distribution of fatigue level at each impression for three campaigns used in the experiment. The fatigue level is calculated according to Section 4.1. Bin width is 5.

Campaign	# creatives	mean	sd
A	21	0.48	0.19
B	12	0.39	0.09
C	5	0.34	0.05

Table 3: Summary of similarity scores for each campaign

Figure 9 shows the relation between the number of ad exposures and the level of fatigue. Apparently, users feel least fatigue with random algorithm (dashed line). Random algorithm just distribute

ad creatives evenly regardless of their performance while two other algorithms choose “better” ad creative to maximize clicks. We see the clear difference between fatigue-aware (dotted line) and baseline (solid line). Fatigue-aware algorithm tries to save the accumulation of fatigue by changing ad creatives when the level of fatigue is high. The fatigue-aware algorithm indeed works as we expected.




Figure 9: The relationship between the number of ad exposures in past 24 hours for each user (frequency) and the level of fatigue. Fitted values by local linear regressions of fatigue on frequency.

Next, we check the association between the level of fatigue and the metrics (CTR and CVR). Figure 10 shows the local linear regressions CTR/CVR on fatigue. To eliminate bias generated by algorithms we use logs from random algorithm only. The two panels shows the difference between CTR and CVR. While CTR constantly decreases with the fatigue, CVR shows more complex relationship. In other words, we see only negative effects of fatigue on CTR but find both positive and negative effects for CVR. Possible explanation is that while audience tend to click on new and fresh ad creative for simple curiosity, they do not login game until they understand the message brought by ad creative after repetitive exposures. Hence, it is possible that the DSP should change the algorithm for different goals. This could explain why Baseline is better than FA in terms of post-click CVR. Baseline tends to show the same ad creatives consistently and gets more chance to get conversions once it gets clicks. But overall performance (CTR and CVR) is in favor of FA.




Figure 10: The relationship between the level of fatigue and CTR (left) and post-impression CVR (right). Local linear regression of CTR/CVR on fatigue. Data is from random algorithm. Each line represents each campaign.

7 CONCLUSION

In this paper, we proposed a fatigue-aware ad creative selection algorithm that explicitly considered the changes in users’ attitude

by the repetitive ad exposures. We introduced a numerical measure of the level of fatigue and set up easy-to-implement and efficient algorithm that changes ad creative according to the level of the fatigue. The proposed algorithm was applied to a running DSP in a production environment for one week. The results show the superiority of the proposed algorithm over the baseline algorithm.

ACKNOWLEDGMENTS

The authors thank the Dynalyst team of CyberAgent, especially Hidetoshi Kawase, Yuta Kurosaki, and Shuhei Kimura for providing immense support to perform the experiments, and Kazuki Taniguchi, Masahiro Nomura, Kota Yamaguchi, Mayu Otani (AILab), Takanori Maehara (RIKEN), and Junpei Komiyama (NYU) for helpful advice.

REFERENCES

- [1] Zoë Abrams and Erik Vee. 2007. Personalized Ad Delivery When Ads Fatigue: An Approximation Algorithm. In *International Workshop on Web and Internet Economics*. Springer, 535–540.
- [2] Deepak Agarwal, Bee-Chung Chen, and Pradheep Elango. 2009. Spatio-Temporal Models for Estimating Click-through Rate. In *Proceedings of the 18th International Conference on World Wide Web - WWW '09*. ACM Press, Madrid, Spain, 21. <https://doi.org/10.1145/1526709.1526713>
- [3] Alberto Bietti, Alekh Agarwal, and John Langford. 2018. A Contextual Bandit Bake-off. *arXiv:1802.04064 [cs, stat]* (Feb. 2018). <http://arxiv.org/abs/1802.04064> arXiv: 1802.04064.
- [4] Christopher M. Bishop. 2006. *Pattern Recognition and Machine Learning*. springer.
- [5] Junyu Cao and Wei Sun. 2019. Dynamic Learning of Sequential Choice Bandit Problem under Marketing Fatigue. *Proceedings of the AAAI Conference on Artificial Intelligence* 33 (July 2019), 3264–3271. <https://doi.org/10.1609/aaai.v33i01.33013264>
- [6] Deepayan Chakrabarti, Ravi Kumar, Filip Radlinski, and Eli Upfal. 2009. Mortal Multi-Armed Bandits. In *Advances in Neural Information Processing Systems 21*, D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou (Eds.). Curran Associates, Inc., 273–280. <http://papers.nips.cc/paper/3580-mortal-multi-armed-bandits.pdf>
- [7] Olivier Chapelle and Lihong Li. 2011. An Empirical Evaluation of Thompson Sampling. In *Advances in neural information processing systems*. Granada, Spain, 2249–2257.
- [8] Patrali Chatterjee, Donna L. Hoffman, and Thomas P. Novak. 2003. Modeling the Clickstream: Implications for Web-Based Advertising Efforts. *Marketing Science* 22, 4 (Nov. 2003), 520–541. <https://doi.org/10.1287/mksc.22.4.520.24906>
- [9] Jiemiao Chen, Xiaojing Yang, and Robert E. Smith. 2016. The effects of creativity on advertising wear-in and wear-out. *Journal of the Academy of Marketing Science* 44, 3 (May 2016), 334–349. <https://doi.org/10.1007/s11747-014-0414-5>
- [10] François Chollet and others. 2015. Keras. (2015). <https://keras.io>
- [11] Neha Gupta, Abhimanyu Das, Sandeep Pandey, and Vijay K. Narayanan. 2012. Factoring Past Exposure in Display Advertising Targeting. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '12)*. ACM, New York, NY, USA, 1204–1212. <https://doi.org/10.1145/2339530.2339719>
- [12] Junpei Komiya and Tao Qin. 2014. Time-Decaying Bandits for Non-stationary Systems. In *Web and Internet Economics*, Tie-Yan Liu, Qi Qi, and Yinyu Ye (Eds.). Vol. 8877. Springer International Publishing, Cham, 460–466. https://doi.org/10.1007/978-3-319-13129-0_40
- [13] Tor Lattimore and Csaba Szepesvári. 2019. Bandit Algorithms. <https://torlattimore.com/downloads/book/book.pdf>
- [14] Pei Lee, Laks V.S. Lakshmanan, Mitul Tiwari, and Sam Shah. 2014. Modeling Impression Discounting in Large-Scale Recommender Systems. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '14*. ACM Press, New York, New York, USA, 1837–1846. <https://doi.org/10.1145/2623330.2623356>
- [15] Huitian Lei, Ambuj Tewari, and Susan A. Murphy. 2017. An Actor-Critic Contextual Bandit Algorithm for Personalized Mobile Health Interventions. *arXiv:1706.09090 [cs, stat]* (June 2017). <http://arxiv.org/abs/1706.09090> arXiv: 1706.09090.
- [16] Nir Levine, Clammer Koby, and Shie Monner. 2017. Rotting Bandits. *Advances in Neural Information Processing Systems* 30 (2017). <http://papers.nips.cc/paper/6900-rotting-bandits>
- [17] Randall A Lewis. 2015. Worn-Out or Just Getting Started? The Impact of Frequency in Online Display Advertising. Boston, Massachusetts, USA.
- [18] Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. 2010. A Contextual-Bandit Approach to Personalized News Article Recommendation. *Proceedings of the 19th international conference on World wide web - WWW '10* (2010), 661. <https://doi.org/10.1145/1772690.1772758> arXiv: 1003.0146.
- [19] Hao Ma, Xueqing Liu, and Zhihong Shen. 2016. User Fatigue in Online News Recommendation. In *Proceedings of the 25th International Conference on World Wide Web (WWW '16)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 1363–1372. <https://doi.org/10.1145/2872427.2874813>
- [20] Kitauchi A. Yamashita T. Hirano Y. Matsuda H.-Takaoka K. Asahara M. Matsumoto, Y. [n. d.]. Japanese morphological analysis system ChaSen version 2.0 manual. ([n. d.]).
- [21] Corneilia Pechman and David W. Stewart. 1988. Advertising Repetition: A Critical Review of Wearin and Wearout. *Current issues and research in advertising* 11, 1-2 (1988), 285–329.
- [22] R. Rehkrek and P. Sojka. 2010. Software Framework for Topic Modelling with Large Corpora. Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks. <http://is.muni.cz/publication/884893/en>
- [23] Howard A. Zhu M. Zhmoginov A. Chen L.-C. Sandler, M. 2018. MobileNetV2: Inverted Residuals and Linear Bottlenecks. *arXiv:1801.04381 [cs]* (2018).
- [24] Hashimoto T. Okumura M Sato, T. 2017. Implementation of a word segmentation dictionary called mecab-ipadic-NEologd and study on how to use it effectively for information retrieval (in Japanese). In *Proceedings of the Twenty-three Annual Meeting of the Association for Natural Language Processing*.
- [25] Susanne Schmidt and Martin Eisend. 2015. Advertising Repetition: A Meta-Analysis on Effective Frequency in Advertising. *Journal of Advertising* 44, 4 (Oct. 2015), 415–428. <https://doi.org/10.1080/00913367.2015.1018460>
- [26] David W. Schumann, Richard E. Petty, and D. Scott Clemons. 1990. Predicting the Effectiveness of Different Strategies of Advertising Variation: A Test of the Repetition-Variation Hypotheses. *Journal of Consumer Research* 17, 2 (1990), 192–202.
- [27] Liang Tang, Yexi Jiang, Lei Li, and Tao Li. 2014. Ensemble Contextual Bandits for Personalized Recommendation. In *Proceedings of the 8th ACM Conference on Recommender Systems - RecSys '14*. ACM Press, Foster City, Silicon Valley, California, USA, 73–80. <https://doi.org/10.1145/2645710.2645732>
- [28] Liang Tang, Romer Rosales, Ajit Singh, and Deepak Agarwal. 2013. Automatic ad format selection via contextual bandits. In *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management - CIKM '13*. ACM Press, San Francisco, California, USA, 1587–1594. <https://doi.org/10.1145/2505515.2514700>
- [29] H. Rao Unnava and Robert E. Burnkrant. 1991. Effects of Repeating Varied Ad Executions on Brand Name Memory. *Journal of Marketing Research* 28, 4 (Nov. 1991), 406–416. <https://doi.org/10.1177/002224379102800403>
- [30] Kilian Weinberger, Anirban Dasgupta, John Langford, Alex Smola, and Josh Attenberg. 2009. Feature hashing for large scale multitask learning. In *Proceedings of the 26th Annual International Conference on Machine Learning - ICML '09*. ACM Press, Montreal, Quebec, Canada, 1–8. <https://doi.org/10.1145/1553374.1553516>
- [31] Robert B. Zajonc. 1968. Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology* 9, 2, Pt.2 (1968), 1–27. <https://doi.org/10.1037/h0025848>