

Unbiased Lift-based Bidding System

Daisuke Moriwaki
CyberAgent, Inc.
moriwaki_daisuke@cyberagent.co.jp

Yuta Hayakawa
CyberAgent, Inc.
hayakawa_yuta@cyberagent.co.jp

Isshu Munemasa
CyberAgent, Inc.
munemasa_issu@cyberagent.co.jp

Yuta Saito
Tokyo Institute of Technology
saito.y.bj@m.titech.ac.jp

Akira Matsui
University of Southern California
amatsui@usc.edu

ABSTRACT

Conventional bidding strategies for online display ad auction heavily relies on observed performance indicators such as clicks or conversions. A bidding strategy naively pursuing these easily observable metrics, however, fails to optimize the profitability of the advertisers. Rather, the bidding strategy that leads to the maximum revenue is a strategy pursuing the performance *lift* of showing ads to a specific user. Therefore, it is essential to predict the lift-effect of showing ads to each user on their target variables from observed log data. However, there is a difficulty in predicting the lift-effect, as the training data gathered by a past bidding strategy may have a strong bias towards the winning impressions. In this study, we develop *Unbiased Lift-based Bidding System*, which maximizes the advertisers' profit by accurately predicting the lift-effect from biased log data. Our system is the first to enable high-performing lift-based bidding strategy by theoretically alleviating the inherent bias in the log. Real-world, large-scale A/B testing successfully demonstrates the superiority and practicability of the proposed system.

CCS CONCEPTS

• Information systems → Display advertising.

KEYWORDS

Lift-based Bidding, Real-Time Bidding, Bid Optimization, Display Advertising, Inverse Propensity Score

ACM Reference Format:

Daisuke Moriwaki, Yuta Hayakawa, Isshu Munemasa, Yuta Saito, and Akira Matsui. 2020. Unbiased Lift-based Bidding System. In *Proceedings of AdKDD '20 (AdKDD '20)*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Online advertising has been playing an integral role in the recent business, which accounts for half of media ad spending in the

world [1]. One of the advantages of the online ad is that the performance of the ads can be assessed by simple metrics such as the number of clicks and conversions earned by the ads. These metrics enable ad deliverers, DSPs (Demand-Side Platforms), to charge ad costs for advertisers through the objective billing systems named CPC (cost-per-click), CPA (cost-per-action), and CPM (cost-per-impression). The conventional goal of DSPs is to optimize these metrics by performance-based bidding strategy, which decides bid price based on the probability of users to take the desired action.

Despite its industrial success, there are two major emerging issues with the current bidding process. First, the conventional bidding strategy ignores the probability that a user will convert even without an ad. Such a strategy is suboptimal in the sense that it will not reach users having a high probability of *changing* their action by showing an ad [11]. Moreover, it might destroy the experience of end-users who have high conversion probability even without an ad [6]. The other problematic issue is the inherent bias in training data. Specifically, for supervised machine learning to work, it is necessary for training and test population to follow the same distribution. However, this is not the case in the online advertising domain. This is because the training data is heavily biased by the ad auction selection in the data collection process (i.e., *impression bias*). Impressions that were assigned to a higher bid price by a past bid strategy would have a higher probability of winning in an auction, thereby having a higher chance to appear in the training data. On the other hand, one has to make a prediction for every impression before the ad auction selection in the testing time. As a result, machine learning models trained on that biased data will overly focus on samples with a high winning probability, resulting in a poor performance in the testing time [14].

These two issues have been solved separately. [11] shows that the conventional strategy optimizing observed clicks or conversions is suboptimal and proposes a lift-based bidding strategy. This strategy decides the bid price based on the predicted *lift* effect of showing an ad to a user and considers their conversion probability without the ad. However, their proposed method does not address the impression bias theoretically. In contrast, [14] proposes a method to alleviate the bias and unbiasedly predict the outcome under the impression bias. However, the bidding strategy considered in [14] still pursues the observed metrics, not the lift-effect. Therefore, a theoretically grounded method that simultaneously solves these two challenging problems has not yet been proposed in the literature.

In this study, we solve the aforementioned issues simultaneously with the aim of maximizing advertisers' revenue while avoiding damaging end-users' experience. We first propose a method to

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.

AdKDD '20, Aug 23, 2020, Virtual

© 2020 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

predict the lift of showing ads to a specific user using only an observable biased impression log. Our method builds on a theoretically validated debiasing method in causal inference called *inverse propensity scoring* and easily implementable with existing machine learning packages such as *scikit-learn*¹ or *XGBoost*². We further describe the details of our resulting *Unbiased Lift-based Bidding System* architecture used in our real production. Our system is scalable and able to optimize advertisers' profit by deciding the bid price for each impression based on our proposed unbiased lift prediction procedure. Finally, we conduct online A/B testing and demonstrate that our proposed system outperforms the naive, conventional performance-based bidding strategy in a real ad campaign.

Our contributions can be summarized as follows:

We propose a theoretically grounded, easily implementable method to predict the lift-effect of showing ads from biased impression log data.

We develop *Unbiased Lift-based Bidding System*, which maximizes the advertisers' profit by bidding based on the predicted lift-effect.

We demonstrate the efficacy of our proposed system in a large-scale online A/B testing on a real-world DSP.

2 MOTIVATION

This study aims at increasing the effectiveness of an advertising campaign by optimizing bid prices under budget constraints. To explain this, we provide a brief overview of the bidding problem.

Marketers craft a whole picture of their advertising campaign and split a fraction of the budget for online ads, spending the rest to other advertisements such as TVs and newspapers. With that specific assigned budget, DSPs join online ad auctions. This predetermined constraint is almost always an equity constraint and the budget DSPs want to consume without any excess and deficiency. Hence the tasks of DSPs are twofold: 1) allocating the campaign budget for online advertisements efficiently to eat up its budget and 2) acquiring an effective impression in RTB (Real-time Bidding).

DSPs are unable to know which or when, even how many users show up in the auction, but only able to predict them. Hence, DSP needs to determine the appropriate bid price for each bid request sequentially by learning the quantity and quality of bid requests to keep budget constraints. Then the DSP's problem at t -th bid request is to choose the best bid price $b_t \in \mathbb{R}_+^0$ to maximize expected profit over τ -period. Let F_t be the winning price of the auction and I_t be the indicator function that tells if the bidder wins or not. F_t is the highest bid for first price auction and second highest bid for second price auction. I_t takes one if the DSP wins and takes zero if not. We treat them as a function of b_t .

Let E_t be the *value* obtained from the ad-slot t . The conventional performance-based bidding strategy uses the probability of conversion with an ad as the value as follows.

$$E_t = \sum_j p_j \theta_j \quad (1)$$

It just uses the observed performance of the ad, but ignores the action changes caused by the ad. In contrast, the lift-based bidding strategy set E_t as the *lift* of conversions caused by the ad impression.

Namely,

$$E_t = \sum_j p_j \theta_j \quad \sum_j p_j \Rightarrow \theta_j \quad (2)$$

Then the DSP solves the following maximization problem.

$$\max_{b_t} \sum_j p_j \theta_j \quad E_t \quad F_t \quad (3)$$

$$s.t. \quad \sum_j p_j \theta_j = \dots \quad (4)$$

where the second equation implies that the expected ad inventory cost F_t must meet remaining budget at t ,

³ With the above optimization problem, DSPs should gain an effective impression having a large lift-effect to maximize the profit. Our focus in this study is to propose a method to predict the lift-effect with the biased impression data and implement a system that controls the bid prices based on the predicted lift-effect.

3 BACKGROUND AND RELATED WORK

We briefly summarize the background of this study and critical related works.

DSPs participate in the online auctions to purchase ad impressions from SSP (Supply-Side Platform). In general, DSPs charges advertisers for their performance measured by an observed metric such as the number of clicks or conversions. Specifically, they claim a fixed price for each click (cost-per-click, CPC), or each conversion (cost-per-action, CPA). As a result, performance-based bidding strategies aim at maximizing the number of clicks or conversions after impression (Eq. (1)).

Just focusing on CPC or CPA, however, may yield suboptimal strategy because they omit the probability that the targeted users get converted (e.g., visit a store, purchase a good) without advertising. To incorporate this omitted realm, the DSP needs to predict incremental gain caused by ad impression (Eq. (2)).

The attempts to incorporate this lift-effect in bidding strategy are lift-based bidding [11], incrementality bidding [4], or more broadly uplift modeling [5, 7–9, 12]. We advance this line of literature and address the inherent bias in impression log data in a theoretically grounded and tractable manner. Specifically, we propose a method to unbiasedly predict the lift-effect of an ad impression on a specific user from biased impression data. Moreover, our method is easily implementable with the well-known machine learning libraries, while the previous debiasing method for the performance-based bidding strategy needs additional implementation cost [14]. We believe that our method will broaden the application of the promising lift-based bidding strategy. This is because every online ad system is biased by the ad auction selection [14], and practitioners can implement and try our unbiased lift-effect prediction method immediately. To our knowledge, we are the first to theoretically consider both the lift-effect prediction and ad impression bias in the online advertising literature.

¹<https://scikit-learn.org/stable/>

²<https://xgboost.readthedocs.io/en/latest/>

³In the real-world production the formulation of the constraint depends on advertisers' interest. F_t can be cost-per-click or cost-per-impression.

4 PROPOSED METHOD

In this section, we propose and describe our *Unbiased Lift-based Bidding System*.

4.1 Setup

We consider a bidding strategy that rewards E for each lift g so that a DSP earns E for each extra conversion (as in Eq. (2)). Our algorithm calculates the bid price $1\beta_{gB^1O}$ as:

$$1\beta_{gB^1O} = U E g^{B^1O} j \mathbf{x}_g^{O'} \quad (5)$$

where $U \geq 10 \cdot 1^0$ is cost rate, a budget pace multiplier that will be detailed in Section 4.3. E is some fixed value a DSP earns for conversion lift. In our algorithm, g is the predicted lift-effect of advertising O for user ℓ characterized by a feature vector \mathbf{x}_g . $B^1O \in \mathcal{S}$ represents the ad exposure state of an ad O under consideration, and \mathcal{S} is a set of possible states. We use the number of impressions of O to ℓ as a scalar variable representing the ad exposure state of O to ℓ , and thus $\mathcal{S} = \{0, 1, \dots, \infty\}$ here.

To formally define the lift-effect g , we introduce the essential notation called *potential outcomes* in causal inference [3]. Let $y_{\ell}^{B^1O}$ denote user ℓ 's potential outcome associated with the exposure state B^1O of ad O . Each user ℓ has potential outcomes associated with every possible state, i.e., $y_{\ell}^{B^1O} \in \mathcal{R}$ (g , however, only one of them is observable). The observed outcome for user ℓ is defined as $y_{\ell}^{B^1O} = y_{\ell}^{B^1O}$ where B^1O is a realized exposure state for user ℓ . Note that the potential outcomes associated with every possible state other than the realized one, i.e., $y_{\ell}^{B^1O} \in \mathcal{R}$ (not $y_{\ell}^{B^1O}$) is unobservable or counterfactual to the analysts.

The lift-effect g of showing ad O for each user ℓ is sequentially defined as the difference between the expectation of the potential outcomes given the two consecutive ad exposure states (the number of impressions):

$$g^{B^1O} j \mathbf{x}_g^{O'} = \mathbb{E}[y_{\ell}^{B^1O} j \mathbf{x}_g^{O'} | B^1O = 1] - \mathbb{E}[y_{\ell}^{B^1O} j \mathbf{x}_g^{O'} | B^1O = 0] \quad (6)$$

where $\mathbb{E}[y_{\ell}^{B^1O} j \mathbf{x}_g^{O'} | B^1O = 1]$ is the expected potential outcome of ℓ when the number of impressions is B^1O , in contrast $\mathbb{E}[y_{\ell}^{B^1O} j \mathbf{x}_g^{O'} | B^1O = 0]$ is the expected potential outcome when the number of impressions is $B^1O - 1$. Thus, Eq. (6) is the reasonable definition for the lift-effect of showing an ad O one more time to a specific user ℓ who has been exposed to the ad $B^1O - 1$ times in the past.

To predict g , we train predictors of outcomes for every possible state $B^1O \in \mathcal{S}$ separately and combine their predictions as follows.

$$\hat{g}^{B^1O} j \mathbf{x}_g^{O'} = 5^{B^1O} \mathbf{x}_g^{O'} + 5^{B^1O-1} \mathbf{x}_g^{O'} \cdot 8B^1O \in \mathcal{S} \quad (7)$$

where 5^{B^1O} and 5^{B^1O-1} predict $\mathbb{E}[y_{\ell}^{B^1O} j \mathbf{x}_g^{O'} | B^1O = B^1O]$ and $\mathbb{E}[y_{\ell}^{B^1O} j \mathbf{x}_g^{O'} | B^1O = B^1O - 1]$, respectively. To accurately predict the lift-effect g , it is essential to predict the expected probability of conversion under each ad exposure state well.

4.2 Unbiased Lift-effect Prediction

To obtain a well-performing predictor 5 for each ad exposure state, it is ideal to directly optimize the following generalization error:

$$\mathbb{E}_{B^1O \in \mathcal{S}} \mathbb{E}_{\mathbf{x}_g^{O'}} [5^{B^1O} \mathbf{x}_g^{O'} - g^{B^1O} j \mathbf{x}_g^{O'}]^2 \quad (8)$$

where 5^{B^1O} is a predictor for $\mathbb{E}[y_{\ell}^{B^1O} | \mathbf{x}_g^{O'}]$, specifies a loss function such as mean squared error, \mathcal{P}^{B^1O} is joint probability distribution of the **whole population**, meaning that the population before the ad auction selection or the testing time. We consider optimizing the generalization error defined over the whole population because we apply 5^{B^1O} to predict the potential outcome in the testing time.

In reality, however, it is impossible to directly optimize Eq. (8). This is because we can utilize only a finite size n_{B^1O} of training data $\mathcal{D}_{B^1O} = \{y_{\ell}^{B^1O} j \mathbf{x}_g^{O'} | B^1O = B^1O\}$ for each ad state and cannot take the expectation to obtain Eq. (8). The conventional solution to this is *Empirical Risk Minimization* (ERM), which optimizes the empirical approximation of Eq. (8) as:

$$5^{B^1O} = \arg \min_{5^{B^1O}} \mathbb{E}_{\mathcal{P}^{B^1O}} [5^{B^1O} - g^{B^1O} j \mathbf{x}_g^{O'}]^2 = \arg \min_{5^{B^1O}} \frac{1}{n_{B^1O}} \sum_{\ell \in \mathcal{D}_{B^1O}} (5^{B^1O} - y_{\ell}^{B^1O})^2$$

The ERM principal works well under the situation of the same train-test distribution, however, the ad impression bias breaks this premise of machine learning. Specifically, the simple empirical approximation of the loss function over \mathcal{D}_{B^1O} has a bias, i.e., $\mathbb{E}_{\mathcal{P}^{B^1O}} [5^{B^1O} - g^{B^1O} j \mathbf{x}_g^{O'}]^2 < \mathbb{E}_{\mathcal{P}^{B^1O}} [5^{B^1O} - g^{B^1O} j \mathbf{x}_g^{O'}]^2$ for some given 5^{B^1O} . The bias issue emerges because users assigned to a higher bid price has higher density in the training data than that in the test data (i.e., $\mathcal{P}^{B^1O} < \mathcal{P}^{B^1O}$). As a result, the trained predictor 5^{B^1O} may perform poorly in the testing time, because it mistakenly overfits to the over-represented samples in the training data.

To alleviate this bias issue of ERM in the online advertising situation, we apply the *inverse propensity score* (IPS) estimation technique to debias the estimation of the ideal loss in Eq. (8). Our loss function takes the following form:

$$\hat{\mathcal{L}}_{\mathcal{P}^{B^1O}} [5^{B^1O}] = \frac{1}{n_{B^1O}} \sum_{\ell \in \mathcal{D}_{B^1O}} \frac{1}{4_{B^1O} \mathbf{x}_g^{O'}} (5^{B^1O} - y_{\ell}^{B^1O})^2 \quad (9)$$

where $n_{B^1O} = |\mathcal{D}_{B^1O}|$ is the total number of the training data, and $4_{B^1O} \mathbf{x}_g^{O'} = \mathbb{P}(B^1O = B^1O | \mathbf{x}_g^{O'})$ is the probability of user ℓ being assigned to the ad exposure state B^1O called the *propensity score*. A fascinating property of the IPS loss in Eq. (9) is that it is unbiased for the ideal generalization error as the following proposition states.

PROPOSITION 4.1. *The IPS loss function in Eq. (9) is unbiased for the ideal generalization error in Eq. (8), i.e., for any given 5^{B^1O} , we have*

$$\mathbb{E}_{\mathcal{P}^{B^1O}} [\hat{\mathcal{L}}_{\mathcal{P}^{B^1O}} [5^{B^1O}]] = \mathbb{E}_{\mathcal{P}^{B^1O}} [5^{B^1O} - g^{B^1O} j \mathbf{x}_g^{O'}]^2$$

under standard identification assumptions in causal inference [3, 8, 10].

The above proposition suggests that our IPS loss function successfully alleviates the bias issue of ERM and approximates the ideal loss from only observable data. Therefore, to unbiasedly predict the lift-effect under the ad impression bias, we optimize the IPS loss and use the resulting predictors to obtain the final lift-effect prediction as:

$$\hat{g}^{B^1O} j \mathbf{x}_g^{O'} = 5^{B^1O} \mathbf{x}_g^{O'} + 5^{B^1O-1} \mathbf{x}_g^{O'} \cdot 8B^1O \in \mathcal{S}$$

where $5^{B^1O} = \arg \min_{5^{B^1O}} \hat{\mathcal{L}}_{\mathcal{P}^{B^1O}} [5^{B^1O}]$ is the IPS loss minimizer.

4.3 PID Control for Cost Rate

We use U for budget pacing, and it discounts the predicted lift-effect in the auctions as in Eq. (5). Appropriate U is not known a priori, and thus we use PID (Proportional Integral Derivative) control algorithm following [13] to adaptively update it to keep the spending constant. We tune the parameters γ, δ, β , and parametric specification of the actuator in Algorithm 1 by an offline simulation. We found that the performance is stable when using the exponential actuator, i.e., $4G\beta^1$ as a parametric specification.

Algorithm 1 PID control for cost rate

Require: Default cost rate U , hyperparameters γ, δ, β
 Set $U = U$
while \dot{Y} **do**
 $AAA =$ remaining budget/remaining hours - budget spending
 at 1
 $U_{s,1} = 4G\beta^1 AAA : \gamma, \delta, \beta$
end while

5 IMPLEMENTATION

We implement the proposed algorithm in a unique DSP specialized for online-to-offline marketing provided by CyberAgent, Inc, a Japan-based ad agency. Our DSP uses location data to determine bid prices and treat users' visits to offline stores as conversions.

5.1 Architecture

We summarize the whole architecture of our bidding system in Figure 1. For each bid request, corresponding g multiplied by cost rate U is returned as the bid price. g is a predicted lift-effect for a coming impression. The ad impression count is calculated by scanning the history of an ad impression for each user. The PID controller updates the cost rate U in every hour based on the gap between ideal budget spending and realized budget spending. The whole procedure is scalable and done within a few milliseconds and does not harm the user experience.

5.2 Preprocessing of Impressions

To attribute each conversion to ad impressions, we use *user-level learning* rather than *log-level learning* used in existing research. Log-level learning implies randomly sampling data points from users' timeline by windows (the rectangles in Figure 2). The method used in [11] can be classified into this group. On the other hand, user-level learning firstly composes user level data and learns the relationship between impressions and conversions for each user. Note that the data size of user-level learning is the number of users, whereas log-level learning generates larger data.

User-level learning has several advantages over log-level learning. First, it does not require any hyperparameter for data splitting such as window size. Second, user-level learning incorporates whole data points including those omitted in log-level learning. One drawback of user-level learning is that the data size is small. This, however, poses little problem for our setting, because DSPs usually reaches out to millions of users for each campaign.

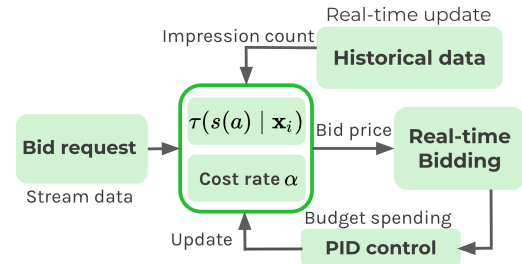


Figure 1: The lift-based bidding system architecture

Note: For each bid request, the DSP server confirms how many impressions the coming user has been exposed to the ad in the past (impression count). The PID controller updates the cost rate U every hour based on the gap between the ideal and realized budget spendings. Then DSP calculates the bid price by combining the cost rate U and pre-predicted lift-effect of exposing the ad to the user one more time g .




Figure 2: Comparing log-level and user-level learning

Note: The log-level learning captures the higher granularity logs than the user-level learning. The log-level learning samples the attributions from the users' timeline giving multiple attribution ids (attr_ids) to the users, and the user-level learning summarizes the whole history of the timeline for each user. The table of log-level learning can have multiple rows (i.e., multiple attr_ids) for a single user, whereas the table of user-level learning has one row per user.

5.3 Lift-effect Model Training

Data Sources. We take the training data from two types of sources. When there are similar advertising campaigns in the past, we use the data for training. When no similar advertising campaign is conducted before we introduce lift-based bidding in the middle of the ongoing campaign and use the data generated at that point.

Ad sizes. The online ad takes various forms, including creative, template, format, and size. Size is especially critical as it is standardized by IAB⁴. To account for the effect of ad size on target variables, Essentially, we group ads into four by size and train size-specific predictors.

Propensity Score Estimation. To use our IPS loss function, we have to estimate the propensity score $4\beta^1$ in Eq. (9). The estimation of the propensity score is the multi-class classification problem. We use multi-class classifier from XGBoost library [2] and train it using the whole training data $f^1 x_{\beta} \cdot \beta_{\beta=1}^0$ and classifies the ad exposure state β_{β} from the feature vectors. The feature vector to estimate

⁴<https://www.iab.com/>

the propensity score includes estimated CVR used for the existing bidder and the number of impressions before the advertising campaign of training period as well as the geographical features of user u . The estimated CVR and the number of impressions before the advertising campaign improve the estimation quality of the propensity score because the higher estimated CVR means higher bid price and the number of impressions in the training period indicates how easy the ad-slots of the user is obtained. For each ad size, we split users into eight classes by the number of impressions $B_{\beta} = \{0, 1, 2, 3, 4, 5, 9, 11, 20, 20\}$. The window size of bins gets larger for the higher number because the distribution of the number of impressions is highly skewed to the right. We train four propensity score predictors for each ad size.

Outcome Prediction. We train outcome predictors for each pair of ad exposure states and ad sizes. We use XGBoost regressor, since the outcome variable y is the number of users' visits to an offline store. To train a regressor, we optimize the IPS loss function in Eq. (9). XGBoost provides a weighting option, thus our IPS loss minimization procedure is easily implementable. The feature vector X includes recency (the last time the user visited the stores), frequency (how many times the user visited the stores), visited POIs (point-of-interests), the number of location logs, census statistics of the user's residential areas such as the share of age groups, averaged size of the households, and land prices.

5.4 The flooring and smoothing g

The predicted lift-effect can be negative. While it is natural for the lift-effect prediction framework to predict the negative effect of the ad, bidding negative price is not rational and could cause a problem in the system that does not expect negative bid price. Hence we first set g to zero for $g < 0$. The floored values are still very volatile across the impression count. It is possible that the expected value of the second impression is very high, while the first impression has zero expected value. Then the impression count for the user never gets two because the bidding for the first impression always loses. To deal with the problem, we smoothed the values with 3-impression count backward-moving average so that the value of the first impression gets high enough in the above example.

6 ONLINE EXPERIMENT

6.1 Experimental Design

To evaluate our unbiased lift-based bidding system, we compared its performance indicators (impressions-per-user, etc.), business KPIs (key performance indicators), and its pacing efficiency with a conventional bidding system through an online A/B testing. The performance-based system employs the community standard bidding system, which decides bid price by Eq. (1). Note that our lift-based bidding system determines the bid price by the predicted lift-effect of showing one additional ad to a user and is described in Eq. (2) and Eq. (5). In the A/B testing, we randomly assigned the two systems to users on an online ad campaign that aims to promote the app released by a major consumer electronics retailer in Japan⁵. The primary aim of the campaign is to increase the number

⁵The experiment duration was June 9-14 2020, where COVID-19 had been a serious issue across the world. Although we understand that the COVID-19 could

Table 1: A/B testing; Lift-based vs Performance-based

	Lift-based	Performance-based	Difference
#Impressions-per-user	1.28	1.0	0.28***
#Clicks-per-user	0.54	1.0	-0.46**
Reach rate	1.71	1.0	0.71***
#Visit-per-user	1.01	1.0	0.0093
Share of visitors	1.004	1.0	0.0036
#Users	467,180	3,310,182	

Note: The lift-based bidding system obtains more impressions and achieves larger reach rate than the performance-based bidding system. We divided each metric by the results of the performance-based strategy for the normalization purpose. **#Impressions-per-user** is the average number of impressions by users; **#Clicks-per-user** is the average number of clicks by users; **Reach rate** is the ratio of the users who got one or more impressions; **Share of visitor** is the ratio of the users who visited one or more offline stores.; The number of stars represents the statistical significance level (p-value $< 0.05^*$, p-value $< 0.01^{**}$, p-value $< 0.001^{***}$).

Table 2: Business KPIs; Lift-based vs Performance-based

	Lift-based	Performance-based
Cost-per-impression	0.27	1.0
Cost-per-reach	0.24	1.0
Cost-per-visit	0.30	1.0

Note: The proposed lift-based strategy efficiently won the impressions and conversions compared to the performance-based one, and thus has positive impacts on our business. We divided each metric by the results of the performance-based strategy for the normalization purpose. The table shows the three essential business KPIs. **Cost-per-impression** is the number of impressions per dollar spent for the ad inventories; **Cost-per-reach** is the number of users with one or more impressions per dollar; **Cost-per-visit** is the number of visits per dollar.

of app users and visitors to the real stores located across Japan. To alleviate the impacts of the experiment on the business, we make a relatively small group for the lift-based group, which results in the unbalanced experiment groups. However, we keep a sufficiently large group size for statistical analysis.

6.2 Results and Discussion

We summarize the results of the A/B testing in Table 1. Compared to the performance-based bidding system, the lift-based bidding system achieved a higher impressions-per-user (28% more) and a higher reach rate (71% more) but obtained fewer clicks (46% less). Consequently, the lift-based bidding system invited more visitors to the real stores on average (0.9% more) and reached a higher share of visitors (0.4% more). In other words, the lift-based bidding system successfully reached out to potential customers and encouraged their visits to the real stores. The lift-based system demonstrates its superiority also in the essential key business indicators (KPIs) described in Table 2. The lift-based strategy won

influence our experiment, we conclude that it did not pose serious issues. First of all, we conducted after the Japanese government called off the state of emergency. Also, the state of emergency even depends voluntarily rather than compulsory, and, unlike other several countries, strict measurements such as "lockdown" were not taken.

