Optimal Real-Time Bidding for Display Advertising

Weinan Zhang, Shuai Yuan, Jun Wang Department of Computer Science, University College London {w.zhang, s.yuan, j.wang}@cs.ucl.ac.uk

ABSTRACT

In this paper we study bid optimisation for real-time bidding (RTB) based display advertising. RTB allows advertisers to bid on a display ad impression in real time when it is being generated. It goes beyond contextual advertising by motivating the bidding focused on user data and it is different from the sponsored search auction where the bid price is associated with keywords. For the demand side, a fundamental technical challenge is to automate the bidding process based on the budget, the campaign objective and various information gathered in runtime and in history. In this paper, the programmatic bidding is cast as a functional optimisation problem. Under certain dependency assumptions, we derive simple bidding functions that can be calculated in real time; our finding shows that the optimal bid has a non-linear relationship with the impression level evaluation such as the click-through rate and the conversion rate, which are estimated in real time from the impression level features. This is different from previous work that is mainly focused on a linear bidding function. Our mathematical derivation suggests that optimal bidding strategies should try to bid more impressions rather than focus on a small set of high valued impressions because according to the current RTB market data, compared to the higher evaluated impressions, the lower evaluated ones are more cost effective and the chances of winning them are relatively higher. Aside from the theoretical insights, offline experiments on a real dataset and online experiments on a production RTB system verify the effectiveness of our proposed optimal bidding strategies and the functional optimisation framework.

Categories and Subject Descriptors

H.3.5 [Information Systems]: Information Storage and Retrieval - Online Information Services

Keywords

Demand-Side Platform, Bid Optimisation, Real-Time Bidding, Display Advertising

1. INTRODUCTION

Real-time bidding (RTB) has recently emerged as a new display advertising paradigm. Unlike traditional sponsored

KDD'14, August 24-27, 2014, New York, NY, USA.

Copyright 2014 ACM 978-1-4503-2956-9/14/08 ...\$15.00.

http://dx.doi.org/10.1145/2623330.2623633.

search or contextual advertising, where an advertiser presets a bid price for each keyword chosen for their campaigns, RTB allows an advertiser to submit a bid for each individual impression (based on the impression level features) in a very short time frame, often less than 100ms [38]. RTB has fundamentally changed the landscape of display advertising because (i) allowing per impression transactions scales the buying process across a large number of available ad inventories including the leftovers; (ii) real-time audience data encourages behaviour (re-)targeting and makes a significant shift towards buying that is focused on user data [31], rather than contextual data.

Demand-Side Platforms (DSPs) are thus created to help advertisers manage their campaigns and optimise their realtime bidding activities. Figure 1 briefly illustrates the role of a DSP in the RTB eco-system. In RTB display advertising, once a user visits a webpage and an ad impression is created, a bid request for the impression is immediately triggered by the publisher (usually the Supply-Side Platform, a.k.a. SSP, a technology platform to manage publishers' inventories) and then sent to the DSPs via an ad exchange. On behalf of an advertiser, the DSP will compute a bid for this impression and return a bid response to the exchange, where a second price auction is usually held to select the winner. Finally the winner is notified and his/her ad is displayed to the user through the publisher. A more detailed introduction to RTB is given in [35, 38].

More specifically, after receiving a bid request, the DSP will find all eligible ad creatives from all campaigns¹ and compute a bid for each of them. The DSP uses both contextual [6] (e.g. domain, webpage, keywords, time and date, geographical location, weather, language, operating system, browser, etc.) and behavioural [22] (e.g. search, browsing, and purchase history, occupation, income, sentiment, etc.) data to compute a bid. It is common and usually encouraged that advertisers buy user interest segments from third data providers [31]. Note that although we confine our work to the Cost-Per-Mille (CPM) pricing model which is commonly adopted in RTB, other less popular models are also available (e.g. Cost-Per-Click and Cost-Per-Acquisition).

This bid calculation (see the bidding engine in Figure 1) is the most important problem for a DSP. The solution to this problem is referred to as a *bidding strategy*. In pure second price auctions [24] for strategic competitors, theoretically the dominant strategy for advertisers is truth-telling: to bid their private values [15]. When facing a bid request, a DSP will evaluate the value of the impression i.e. to estimate the click-through/conversion rate (CTR/CVR) and

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.

¹The eligibility of the creative and campaign means their target combinations match the bid request, such as placement size, user demographics, geographical location, language, etc. It acts as pre-filtering rules before the bidding process.



Figure 1: An illustration of a demand-side platform and its bidding engine in RTB display advertising.

multiply it by the value of a click/conversion [26]. Many advertisers simply set this value as their bid [32, 26] and keep using it throughout a campaign's lifetime. However, when computing a bid, practical constraints need to be taken into account including the bid landscape (the expected impressions against the bid price), total budget and the campaign's remaining lifetime. These considerations enable the DSP to optimise towards the overall performance of a campaign (usually referred to as Key Performance Indicator, KPI, e.g. the number of clicks, conversions or total revenue) using stochastic methods rather than assuming advertisers are strategic and have a private "true" value per impression [3].

In this paper, the impression level bidding strategy is formulated as a function that maps the individual impression evaluation to the bid value. A novel functional optimisation framework is proposed to find the optimal bidding function: (i) given the budget constraint and the campaign's lifetime, and (ii) taking into account various statistics such as the probability of winning the auction and the prior distribution of the impression features. Our analytical solution indicates that the auction winning function plays a more critical role in shaping the bidding function, whereas the distribution of the features is less correlated. Simple winning functions derived from practical bidding data result in optimal bidding functions that are non-linear and in a *concave* form. Unlike a linear function previously proposed [32], our bidding function encourages to higher bids for impressions with low estimated value because compared to higher evaluated ones, those are more cost-effective and the chance of winning them are relatively higher. Apart from the theoretic insights, both offline experiments on a real dataset and online experiments on a production DSP show that the proposed bidding strategies outperform the strong baselines that have been considered.

To summarise, the contributions of our work are listed below. (i) We propose a novel functional optimisation framework to find the optimal bidding strategy in RTB display advertising. (ii) Based on the auction winning function built from the data, the derived optimal bidding function from our framework is in the concave form against the KPI of each impression, which to our knowledge has never been studied in previous literature on RTB display advertising. (iii) Comprehensive offline and online experiments are conducted to verify the practical effectiveness of our proposed bidding strategy.

2. RELATED WORK

Bid optimisation is a well-studied problem in online advertising [17, 21, 18, 32]. Nonetheless, most research has been so far limited to keyword auction in the context of sponsored search (SS) [15, 4, 28]. Typically, under the scenario of pre-setting the keyword bids (not impression level), the keyword utility, cost and volume are estimated and then an optimisation process is performed to optimise the advertisers' objectives (KPIs) [5, 42, 23, 41]. Given a campaign budget as the cost upper bound, optimising the advertiser performance is defined as a budget optimisation problem [17, 30]. Furthermore, the authors in [7, 16] focus on the bid generation and optimisation on broad matched keywords, where query language features are leveraged to infer the bid price of related keywords. The authors in [40] propose to jointly optimise the keyword-level bid and campaign-level budget allocation under a multi-campaign sponsored search account. Some recent work focuses on periodically changing the pre-setting keyword auction price, taking into account the remaining budget and lifetime. For instance, in [3, 20], Markov decision process is used to perform online decision in tuning the keyword bid price, where the remaining auction volume and budget act as states and the bid price setting as actions. In [23] authors propose to calculate a bid allocation plan during the campaign lifetime, where the bid price on each keyword is set in different discrete time unit by considering the market competition and the CTR on different ad positions. However, none of the work evaluates per-impression auction as in SS, all the bids are associated with keywords and impression level features are seldom considered, especially for advertisers and their agencies. Moreover, in SS bid optimisation, search engines play two roles: setting the keyword bids as well as hosting the auctions. The objective function could be diverted to optimise the overall revenue for the search engine [2, 33, 14, 43], rather than the performance of each individual advertiser's campaigns.

The bid optimisation for RTB display advertising is fundamentally different. First, the bids are not determined by pre-defined keywords [38], but are based on impression level features. Although in general advertisers (or DSPs) are required to set up their target rules, they need to estimate the value of each ad impression that is being auctioned in real time and return the bid price per auction. Second, in RTB, CPM pricing is generally used [38]. Winning an impression directly results in the cost, despite the fact clicks and conversions can now be directly optimised by advertisers and DSPs. Thus, the dependencies over various effectiveness measures such as $eCPC^2$, CPM and the budget constraints need to be studied in a single framework. In [18] the authors propose an algorithm that learns winning bids distribution from full or partial information of auctions in display advertising. The algorithm then makes bidding decisions to achieve the best delivery (number of impressions) within the budget constraint. In [10], the bid price from each campaign can be adjusted by the publisher side in real time and the target is to maximise the publisher side

 $^{^2 \}rm Effective cost per click (eCPC)$ - The cost of a campaign divided by the total number of clicks delivered.

Table 1: Notations and descriptions.

Notation	Description
$oldsymbol{x}$	The bid request represented by its features.
$p_{x}(\boldsymbol{x})$	The probability density function of \boldsymbol{x} .
$ heta(oldsymbol{x})$	The predicted KPI if winning the auction
	of \boldsymbol{x} . It could be the CTR, CVR etc.
$p_{ heta}(heta)$	The probability density function of KPI θ .
B	The campaign budget.
N_T	The estimated number of bid requests during
	the lifetime T of the budget.
$b(heta(oldsymbol{x}),oldsymbol{x})$	The bidding strategy is defined as function
	b(). Assume a generative process:
	$\boldsymbol{x} \to \boldsymbol{\theta} \to b$, so $b(\boldsymbol{\theta}(\boldsymbol{x}), \boldsymbol{x}) \equiv b(\boldsymbol{\theta}(\boldsymbol{x}))$. See
	the text. We occasionally use notation
	b to refer to a specific bid value.
$w(b(heta(oldsymbol{x})),oldsymbol{x})$	The probability of winning the bid request \boldsymbol{x}
	with bid price $b(\theta(\boldsymbol{x}))$. We approximate it by
	the dependency assumption: $\boldsymbol{x} \to \boldsymbol{\theta} \to \boldsymbol{b} \to \boldsymbol{w}$,
	so $w(b(\theta(\boldsymbol{x})), \boldsymbol{x}) \equiv w(b(\theta(\boldsymbol{x})))$. See the text
	for details.

revenue. The most relevant work is in [32], where DSP returns the bid price with a linear relationship to the predicted CTR (pCTR) for each ad impression being auctioned. The analytical solution from our proposed functional optimisation framework shows that an optimal bidding function should be non-linear. The non-linearity is closely related to the probability of auction winning, but is loosely correlated with the prior distribution of the impression features.

As a new advertising paradigm, other problems in RTB have also been studied. In [25] the authors focus on the pacing problem, where the target is to smoothly deliver the campaign budget. From the SSP perspective, the reserve price setting in RTB ad auctions is studied in [36]. In [26] the sparsity problem of conversion rate estimation is handled by modelling the conversions at different selected hierarchical levels. In [13] the authors study the evaluation measures of the display advertising performance and they find the site visit turns to be a better proxy than the user click. In addition, there is some work on the ad exchange communication problem [9, 29]. More discussion on related research of RTB can be found in [38].

3. OPTIMAL REAL-TIME BIDDING

Each user visit of a publisher page triggers an ad auction for each ad slot on that page, which, in real time, generates a bid request for each targeting campaign. Given a bid request, along with features covering user, ad, context, and auction information, the DSP bidding engine decides whether to participate this auction, and if participates, returns a bid for this auction. The bid price depends on many factors. It is not only influenced by the predicted KPI (p-KPI) value of the ad impression being auctioned, such as the CTR and CVR, that the advertiser wants to achieve, but most importantly, related to many other factors such as the budget constraint, the probability of auction winning, and the features and cost of this particular ad impression. In this section, we consider the problem of optimally generating real-time bids as a functional optimisation problem and propose a novel optimisation framework by taking all these factors into account. We show that solving it leads to a practical bidding function.

3.1 Problem Definition

Let us first mathematically formulate the real-time bidding problem. To launch a campaign in display advertising, the advertiser uploads their ad creatives, sets the targeting rules (e.g. the user segmentation, time, location) and the corresponding budget for the lifetime of the campaign.

After the target rules are set, before optimising the bid, the advertiser would first spend a small amount of budget

to bid random impressions in order to learn some statistics. For instance, as studied in [8, 12] the auction volume forecast (e.g. bid landscape prediction) module is usually employed to estimate auction statistics with respect to the current setting and budget constraint. More specifically, we denote the estimated number of bid requests for the target rules during the lifetime T as N_T . Also, each bid request is represented by a high dimensional feature vector \boldsymbol{x} , where its entries consist of two sets of features: one is exacted from the campaign ad itself, and the other is related to the impression being auctioned, e.g., the cookie information of the underlying user, location, time, user terminal, browser, the contextual information about the webpage, etc. We use $p_x(x)$ to denote the prior distribution of the feature vectors which match the campaign target rules. For each campaign, the advertiser can use the historic bidding and feedback data to predict the KPI for the ad impression being auctioned. We denote the predicted KPI of a bid request \boldsymbol{x} as $\theta(\boldsymbol{x})$. Note that different advertisers might consider different KPIs. For example, if the goal of a campaign is to maximise the direct visits, i.e. the total number of clicks, then $\theta(\boldsymbol{x})$ denotes the pCTR for that impression. For CTR estimation, we refer to [19, 34]. If the goal of a campaign is for conversions, then $\theta(\mathbf{x})$ denotes the predicted CVR (pCVR) for that impression. We also denote $p_{\theta}(\theta)$ as the prior distribution of the predicted KPI per bid request. The notation descriptions are given in Table 1.

Once the major statistics are gathered, our problem now is to devise an optimal bidding strategy such that a certain KPI objective over the budget will be maximised. For the sake of clarity, we consider the number of clicks as the objective here, while a natural extension to covering alternative KPIs and their experimental results are given later in Section 5.5. Mathematically, our optimal bid generation problem is formulated as a *functional* optimisation problem:

 $b()_{ORTB} =$

$$\underset{b()}{\operatorname{arg\,max}} \quad N_T \int_{\boldsymbol{x}} \theta(\boldsymbol{x}) w(b(\theta(\boldsymbol{x}), \boldsymbol{x}), \boldsymbol{x}) p_x(\boldsymbol{x}) d\boldsymbol{x} \tag{1}$$
while the $N_T \int_{\boldsymbol{x}} b(\theta(\boldsymbol{x}), \boldsymbol{x}) w(b(\theta(\boldsymbol{x}), \boldsymbol{x}), \boldsymbol{x}) p_x(\boldsymbol{x}) d\boldsymbol{x} \in B$

subject to
$$N_T \int_{\boldsymbol{x}} b(\theta(\boldsymbol{x}), \boldsymbol{x}) w(b(\theta(\boldsymbol{x}), \boldsymbol{x}), \boldsymbol{x}) p_x(\boldsymbol{x}) d\boldsymbol{x} \leq B.$$

where $b(\theta(\boldsymbol{x}), \boldsymbol{x})$ denotes the bidding function we intend to obtain. It depends on the feature vector \boldsymbol{x} and the estimated CTR $\theta(\mathbf{x})$. $w(b, \mathbf{x})$ denotes the estimated winning rate for a bid price b given the feature \boldsymbol{x} of the impression auction. In Eq. (1), the product of $\theta(\boldsymbol{x})$ and $w(b, \boldsymbol{x})$ produces the probability of click given an impression auction. Marginalising it over the feature space yields the expected click per impression auction. Note that in practice, the impression auctions arrive sequentially, so one can potentially make a sequential bidding rule by taking a feedback loop and employing a dynamical optimisation model such as partially observable Markov decision processes (POMDPs) [37]. However, generally these models are computationally expensive thus not feasible in our case, where bid decisions usually need to be returned within 100ms. We, thus, take a two-stage approach (learning statistics such as $p_x(x)$ and N_T , then optimising the bids). Without loss of generality, we consider a simple static model and follow a widely accepted assumption in the previous bid optimisation work [20, 10]: each time the feature vector is independently generated from an identical distribution.

The constraint is made by an upper bound of the expected cost. RTB normally applies the second price auction (pay the second highest bid). However, due to the reserve price setting, the cost is quite often higher than the second highest bid [36, 38]. We thus use the bid price $b(\theta(\boldsymbol{x}), \boldsymbol{x})$ as the upper bound of the cost of winning. Specifically, the product of the cost and the winning rate produces the expected cost per impression auction. Marginalising it over the feature space and multiplying by N_T yields the expected cost, which is set to be not greater than the budget B. Note that by maximising the click number with the budget constraint, the eCPC is minimised, which is a much important measure in display advertising.

To make the above problem solvable, we consider sequential dependency among the variables for each auction by making the following assumptions:

- Assume b(θ(x), x) ≡ b(θ(x)). That is: x → θ → b. This allows us to largely reduce the functional decision space for the optimisation, while still gaining the dependency of the impression features through the KPI estimation θ(x). The previous work in [32] also adopts a similar dependency (bid only depends on CTR).
- Assume w(b, x) ≡ w(b). That makes the feature x only influence the winning rate via its generated bid: x → θ → b → w. The assumption as sensible as we found out (shown in Section 4.1) that the dependency over the bid request features is far less than the bid price. Previous SS bid optimisation work [23, 40] also makes such assumption on winning keyword ad slots.

Then the optimisation problem is rewritten as

$$b()_{\text{ORTB}} = \underset{b()}{\operatorname{arg\,max}} \quad N_T \int_{\boldsymbol{x}} \theta(\boldsymbol{x}) w(b(\theta(\boldsymbol{x}))) p_x(\boldsymbol{x}) d\boldsymbol{x}$$
(2)
subject to $N_T \int_{\boldsymbol{x}} b(\theta(\boldsymbol{x})) w(b(\theta(\boldsymbol{x}))) p_x(\boldsymbol{x}) d\boldsymbol{x} \leq B.$

Furthermore, since there is a deterministic relationship between \boldsymbol{x} and $\theta(\boldsymbol{x})$, the relationship between their probability density is also determined:

$$p_{\theta}(\theta(\boldsymbol{x})) = \frac{p_{\boldsymbol{x}}(\boldsymbol{x})}{||\nabla \theta(\boldsymbol{x})||}.$$
(3)

Thus we can focus on θ via performing integration by substitution

$$\int_{\boldsymbol{x}} \theta(\boldsymbol{x}) w(b(\theta(\boldsymbol{x}))) p_{\boldsymbol{x}}(\boldsymbol{x}) d\boldsymbol{x}$$

$$= \int_{\boldsymbol{x}} \theta(\boldsymbol{x}) w(b(\theta(\boldsymbol{x}))) p_{\theta}(\theta(\boldsymbol{x})) || \nabla \theta(\boldsymbol{x}) || d\boldsymbol{x}$$

$$= \int_{\theta(\boldsymbol{x})} \theta(\boldsymbol{x}) w(b(\theta(\boldsymbol{x}))) p_{\theta}(\theta(\boldsymbol{x})) d\theta(\boldsymbol{x})$$

$$= \int_{\theta} \theta w(b(\theta)) p_{\theta}(\theta) d\theta \qquad (4)$$

and the similar substitution for $\int_{\boldsymbol{x}} b(\theta(\boldsymbol{x})) w(b(\theta(\boldsymbol{x}))) p_x(\boldsymbol{x}) d\boldsymbol{x}$. Rewriting the integration with respect to θ leads to our final functional optimisation problem as follows:

$$b()_{\text{ORTB}} = \underset{b()}{\operatorname{arg\,max}} \quad N_T \int_{\theta} \theta w(b(\theta)) p_{\theta}(\theta) d\theta \tag{5}$$

subject to $N_T \int_{\theta} b(\theta) w(b(\theta)) p_{\theta}(\theta) d\theta \leq B.$

3.2 Optimal Solutions

The Lagrangian of the objective function (Eq. (5)) is

$$\mathcal{L}(b(\theta),\lambda) = \int_{\theta} \theta w(b(\theta)) p_{\theta}(\theta) d\theta -\lambda \int_{\theta} b(\theta) w(b(\theta)) p_{\theta}(\theta) d\theta + \frac{\lambda B}{N_T}, \quad (6)$$



Figure 2: Winning function 1 and corresponding optimal bidding function $b_{ORTB1}(\theta)$.

where λ is the Lagrangian multiplier. Based on calculus of variations, the Euler-Lagrange condition of $b(\theta)$ is

$$\theta p_{\theta}(\theta) \frac{\partial w(b(\theta))}{\partial b(\theta)} - \lambda p_{\theta}(\theta) \Big[w(b(\theta)) + b(\theta) \frac{\partial w(b(\theta))}{\partial b(\theta)} \Big] = 0$$
(7)

$$\lambda w(b(\theta)) = \left[\theta - \lambda b(\theta)\right] \frac{\partial w(b(\theta))}{\partial b(\theta)},\tag{8}$$

where we can see the KPI probability density $p_{\theta}(\theta)$ has been eliminated and the *form* of bidding function $b(\theta)$ only depends on the winning function $w(b(\theta))^3$. This is mainly because both the objective and constraint take the integration over the distribution of $p_{\theta}(\theta)$. Different winning functions result in different optimal bidding functions. Here we present two winning functions which are typical and fit the curves of real-world data. And we derive the optimal bidding function form for each winning function.

3.2.1 Winning & Bidding Function 1

As depicted in Figure 5 from our experiment on a real data⁴, the winning rate w(b) consistently has an (approximately) concave shape: when the bid price is low, adding a unit bid will increase the winning rate more than when the bid is already high. Thus a simple winning function is in the form of

$$w(b(\theta)) = \frac{b(\theta)}{c + b(\theta)},\tag{9}$$

where c is a constant. An illustration of the winning function with different c's is given in Figure 2(a).

Taking a derivative with respect to the bid gives:

$$\frac{\partial w(b(\theta))}{\partial b(\theta)} = \frac{c}{(c+b(\theta))^2}.$$
(10)

Taking Eq. (9) and (10) into Eq. (8) gives:

$$\frac{\theta c}{(c+b(\theta))^2} - \lambda \Big[\frac{b(\theta)}{c+b(\theta)} + c\frac{b(\theta)}{(c+b(\theta))^2}\Big] = 0$$
(11)

$$\left(b(\theta)+c\right)^2 = c^2 + \frac{\theta c}{\lambda}.$$
 (12)

Solving the above equation gives our final optimal bidding function:

$$b_{\text{ORTB1}}(\theta) = \sqrt{\frac{c}{\lambda}\theta + c^2} - c.$$
 (13)

³Later we will show that the optimal value of λ depends on $p_{\theta}(\theta)$, but λ is only a parameter in $b(\theta)$; thus $p_{\theta}(\theta)$ does not change the general form of $b(\theta)$.

⁴All the price numbers are presented with the unit of RMB cent. And the bid price is counted on CPM.



Figure 3: Winning function 2 and corresponding optimal bidding function $b_{ORTB2}(\theta)$.

Under the assumption of the winning function 1 in the form of Eq. (9), the optimal bidding function $b_{\text{ORTB1}}(\theta)$ is in a concave form: a square root function form. Figure 2(b) gives an illustration of this bidding function with different parameter c's, fixing $\lambda = 5.2 \times 10^{-7}$.

3.2.2 Winning & Bidding Function 2

For some campaigns with competitive targets, or the targeted publishers/SSPs setting a high reserve price, the winning probability will not increase rapidly when the bid price is around 0; only after the bid price becomes larger than some non-zero value the winning probability starts to dramatically increase. Such case usually occurs in high-profile ad slots [40]. To get this feature, we slightly alter it and propose an alternative of winning function:

$$w(b(\theta)) = \frac{b^2(\theta)}{c^2 + b^2(\theta)},$$
(14)

where the parameter c controls the increasing point of the curve⁵. An illustration is given in Figure 3(a).

Following the same token, we solve Eq. (8) using the winning function in Eq. (14), i.e.

$$b_{\text{ORTB2}}(\theta) = c \cdot \left[\left(\frac{\theta + \sqrt{c^2 \lambda^2 + \theta^2}}{c \lambda} \right)^{\frac{1}{3}} - \left(\frac{c \lambda}{\theta + \sqrt{c^2 \lambda^2 + \theta^2}} \right)^{\frac{1}{3}} \right].$$
(15)

Fixing $\lambda = 5.2 \times 10^{-7}$, the bidding functions with different c's are shown in Figure 3(b). Again the $b_{\text{ORTB2}}(\theta)$ is a concave function.

Our optimisation framework is a general one: Eq. (8) shows that different winning functions would lead to different optimal bidding functions. The framework can adapt to various ad markets with different winning functions. Here we estimate the winning functions from real data (Figure 5) and limit our study to the RTB markets only.

3.2.3 Discussions on Resulting Bidding Functions

Unlike the linear form bidding function in the previous study [32, 26] (denoted as Lin), our derived bidding functions (denoted as ORTB) Eq. (13) and (15) suggest a nonlinear concave form mapping from CTR estimation to the bid value under a budget constraint for RTB. As shown in Figure 4, compared with Lin, ORTB bids higher when the estimated KPI is low, which means ORTB allocates more budget on the low reward and low cost cases.

The strategy of bidding more low cost impressions comes from the shape of the winning functions. In Figure 5 we find that for all the campaigns, when we increase the bid



Figure 4: Concave bidding function bids higher on the cases with pCTR less than 0.0011 (example) and vice versa, which allocates more budget on the cases with low bid price and thus low cost.

price from zero, the winning probability will have a high growth rate first and after the bid price surpasses a region, the winning probability starts to converge to 1. As such, the strategy of ORTB will earn much more winning probability while only increase a little CPM because of the concavity of winning rate respective to the bid price.

3.2.4 Optimal Solution of λ

The bidding functions in Eq. (13) and (15) also take λ as the parameter; let us explicitly denote them as $b(\theta, \lambda)$. To calculate the optimal λ , the Euler-Lagrange condition of λ from Eq. (6) is

$$\int_{\theta} b(\theta, \lambda) w(b(\theta, \lambda)) p_{\theta}(\theta) d\theta = \frac{B}{N_T}.$$
(16)

Given the formula of $b(\theta, \lambda)$, the solution λ can be found. However, in many cases such as using our winning functions, there is no analytic solution of λ . Also we can see the solution depends on $p_{\theta}(\theta)$. Alternatively, one can find the numeric solution using the bidding log data and practically solve it using efficient numeric calculation. In this paper, we take a rather pragmatic approach by regarding λ as a tuning parameter for the bidding functions and learn it from the data. From Eq. (13), (15) and (16), it is easy to see that when λ decreases the value of $\int_{\theta} b(\theta, \lambda)w(b(\theta, \lambda))p_{\theta}(\theta)d\theta$ monotonically increases. Therefore, with larger per-case budget B/N_T , the solution of λ becomes smaller, which corresponds to a higher bid price. Our experiment will demonstrate the trend of optimal λ corresponding to different per-case budget B/N_T (Figure 11 and 12).

4. EXPERIMENT SETUP

Our methods are tested both by offline evaluation (Section 5) using a real-world dataset and via online evaluation (Section 6) over a commercial DSP with real advertisers and impressions. In this section, we introduce the experiment setup and report the results from our data analysis.

4.1 Dataset and Analysis

Dataset description. We use the real-world bidding feedback log from a well-known DSP company as our dataset⁶. It records more than 15 million impressions and the user feedback of 9 campaigns from different advertisers during ten days in 2013. For each bid request, the log contains the information from the user (e.g. the user segmentation), advertiser (e.g. the creative format and size), publisher (e.g.

⁵Actually we can take a more general form of the winning function: $w(b(\theta)) = b^{\alpha}(\theta)/(c^{\alpha} + b^{\alpha}(\theta))$. We investigate the case of $\alpha = 1, 2$ in this paper. When α is larger than 2, there is no analytic solution of $b_{\text{ORTB}}(\theta)$.

 $^{^{6}\}mathrm{The}$ dataset has been publicly released on our website: http://data.computational-advertising.org



 Table 2: Dataset statistics.

Figure 5: Relationship between auction winning rate and bid value for different campaigns.

the auction reserve price, ad slot, page domain and URL) and the context (e.g. the time, region, the browser and operation system). For each bid request, there is an auction winning bid price, and the user feedback (click, conversion) is recorded if the advertiser won the auction. More details of the dataset is shown in Table 2.

Data analysis on winning bids. Figure 5 depicts the winning rate with respect to the bid price for campaigns 1-6 (due to space limit, campaigns 7-9 are not shown). We observe that all the campaigns follow a similar pattern: as the bid price increases, the campaign winning rate increases dramatically; when the bid price gets larger (e.g. more than 100), the increasing of winning rate slows down and finally the winning rate converges to 1. Thus, it is reasonable to employ concave functions like Eq. (9) and (14) to model the relationships. For each campaign, we fit the winning functions with the parameter c leading the least square error with the real curve.

Next, we study the dependency between the bid request features and the winning bid prices (also named as market prices in [3]). Figure 6 gives the box plot [27] of winning price distribution against the features such as hour, weekday, user browser, operation system and location regions of bid requests to campaign 1 (other campaigns follow the similar pattern). Compared with the clear changing relationship with the bid price shown in Figure 5, Figure 6 shows that the winning price distributions do not have obvious dependency on the categorical feature value. It suggests that the bid price is the key factor influencing the campaign's winning rate in its auctions. Once we have known the bid value, the winning rate is less sensitive to the other bid request features. Thus it is practically reasonable to simplify $w(b, \mathbf{x}) \equiv w(b)$ as proposed in Section 3.1.

Training/test split. For each campaign, we perform a 2:1 train/test data split by the time sequence. The training data is mainly used to train the CTR estimator and tune the bidding function parameters. The test data is used to evaluate the compared DSP bidding strategies.



Figure 6: Winning bid distribution against different features for campaign 1.

4.2 Evaluation Measure

The task of DSPs is to optimise each campaign's KPI (such as clicks, conversions etc.) given the budget. Therefore, the KPI is the primary evaluation measure in our experiment. Specifically, we have considered clicks as our primary KPI, while other statistics such as CPM, and eCPC are also monitored. In addition, we also test on an alternative KPI which is the combination of achieved clicks and conversions, as will be discussed in Section 5.5.

4.3 KPI Estimator Training

For each campaign, we use its impression/click/conversion log data to train a KPI estimator for each bid request. Particularly, if the KPI is click, then this task turns to be the well-known CTR estimation [19, 34]. Regression models such as random forest and gradient boosting regression tree can be applied here. Since this work mainly focuses on the bidding strategy instead of the KPI estimator model, we apply the Logistic regression as our CTR estimator as it is a widely used choice [34]. The loss is the cross entropy between the predicted click probability and the ground-truth result. In addition, L2 regularisation is used. For an alternative KPI in our experiment, the KPI estimator training will be discussed in Section 5.5.

Features are extracted from the log data to train the C-TR estimator model. Specifically, we extract 29,408 first-order binary features and based on that generate 657,756 second-order binary features, which yields the total 687,164 features for our training.

4.4 Test Evaluation Setup

Evaluation flow. The evaluation flow is depicted in Figure 7. Given the bidding strategy and a budget for the test period for a particular campaign, we can go through its test data. The test data is a list of records. Each record consists of the features of one bid request, the auction winning price and the user feedback information. Specifically, receiving the bid request features of each record by the timestamp, the bidding strategy generates a bid price for it (if the cost is beyond the budget, just returns 0, i.e. skips the remaining bid requests). If the bid price is higher than the auction, gets its ad shown. The corresponding user feedback (click) and the corresponding charged price of the record are then referenced to update the performance and cost. After that,



Figure 7: Evaluation flow chart.

if there is no more bid requests in the test data, then the evaluation is over, with the final performance returned.

It is worth mentioning that using user feedback logs for evaluation has limitation. In our case, user feedback only occurs for the winning auctions (having ad impressions); there is no user feedback for the losing bids. We thus do not know whether the user will click or convert even if we bid enough high to win the originally losing ad auction in our offline experiment. In this paper, we follow the convention of the offline evaluations from sponsored search [19], recommender systems [39] and Web search [11] that the objects (auctions) with unseen user feedback are ignored (i.e., considered as auction failure cases). To complement with the offline evaluation, we will further show the online test performance on a production DSP in Section 6.

Budget constraints. It is easy to see that if we set the budget the same as the original total cost in the test log, then just simply bidding as high as possible for each case will exactly run out the budget and get all the logged clicks. In our work, to test the performance against various budget constraints, for each campaign, we respectively run the evaluation test using 1/64, 1/32, 1/16, 1/8, 1/4 and 1/2of the original total cost in the test log as the budget.

4.5 Compared DSP Bidding Strategies

We compare the following baseline and state-of-the-art bidding strategies in our experiment. The parameters of each bidding strategy are tuned using the training data.

Constant bidding (Const). Bid a constant value for all the bid requests to the campaign. The parameter is the specific constant bid price.

Random bidding (Rand). Randomly choose a bid value in a given range. The parameter is the upper bound of the random bidding range.

Bidding below max eCPC (Mcpc). As discussed in [26], given the advertiser's goal on max eCPC, which is the upper bound of cost per click, the bid price on an impression is obtained by multiplying max eCPC and pCTR. Here we calculate the max eCPC for each campaign by dividing its cost and achieved number of clicks in the training data. No parameter for this bidding strategy.

Linear-form bidding of pCTR (Lin). In the previous work [32], the bid value is linearly proportional to the pC-TR. The formula can be generally written as

$$b_{\rm Lin}(\theta) = b_0 \frac{\theta}{\theta_0},\tag{17}$$

where θ_0 is the average CTR under a target condition (e.g., a user-inventory pair) and b_0 is the basic bid price for this target condition. We tune b_0 in our experiment.

Optimal real-time bidding (ORTB1 and ORTB2). These are our derived optimal bidding strategies in our



Figure 8: Overall performance comparison.

framework. As shown in Eq. (13) and (15), the parameters are c and λ , where c is obtained by fitting the winning probability and λ is tuned using the training data.

Table 3 summarises the attributes of the different strategies. Mcpc is not budget-aware to make sure it spends all the budget. Mcpc, Lin and ORTB perform bidding based on impression-level evaluation. Taking winning functions into account, ORTB is the most informative strategy. In Section 5 we will analyse the impact of these attributes in the final performance.

5. OFFLINE EVALUATION

In the experiment, we answer the following questions. (i) Does our derived non-linear bidding function outperform the state-of-the-art linear one? (ii) What are the characteristics of our ORTB algorithms and how do the parameters impact the performance and what are their relationships with the budget conditions?

5.1 Performance Comparison

The performance comparison on total clicks and eCPC under different budget conditions are reported in Figure 8. We observe that (i) under every budget condition, our proposed bidding strategies ORTB1 and ORTB2 have the best performance on total clicks, which verifies the effectiveness of the derived non-linear forms of the bidding function. (ii) Except ORTB, Lin is the best algorithm in the comparison. This algorithm represents the widely used DSP bidding strategies in product [32]. (iii) Mcpc is aware of the upper bound cost for each bid request, and dynamically changes its bid according to the estimated CTR. However, compared to ORTB and Lin, Mcpc has no adaptibility to different budget conditions. For example, when the budget is relatively low for the bid request volume, Mcpc will still bid based on the originally set max eCPC, while ORTB and Lin can adaptively lower the bid to earn the impressions and clicks with higher RIO. (iv) Rand and Const provide very low performance even though their parameters are tuned under different budget conditions. (v) Also from the eCPC performance we can see Rand and Const spend much more money to get one click than the case-value-aware strategy Mcpc, Lin and ORTB. The last two points suggest the benefit of real-time bidding based display advertising: evaluating the

Table 4: Click improvement of **ORTB1** over Lin for each campaign under different budget conditions.

Cpg.	1/64	1/32	1/16	1/8	1/4	1/2
1	0.68%	1.97%	-0.46%	1.04%	1.25%	0.26%
2	0.70%	0.00%	0.44%	3.38%	0.48%	0.90%
3	27.16%	2.33%	1.45%	0.43%	6.50%	1.11%
4	24.46%	-1.19%	3.11%	2.84%	0.06%	3.14%
5	49.54%	10.85%	0.33%	5.17%	3.82%	1.42%
6	84.07%	49.89%	9.13%	-1.99%	5.67%	3.51%
7	85.67%	51.10%	18.50%	6.27%	2.67%	0.00%
8	58.13%	23.78%	8.82%	1.93%	0.59%	0.65%
9	428.26%	118.78%	103.91%	63.68%	18.82%	6.50%
		•				
	₩ 40% -					
	ner					
	ē	\				



Figure 9: Overall click improvement of ORTB1 over Lin under different budget conditions.

value for each bid request (impression level) plays a significant role in the performance.

Table 4 gives detailed performance improvement on total clicks of ORTB1 over Lin under different campaigns and budget conditions. Among the listed 54 settings, ORTB1 wins Lin in 49 (90.7%) settings, ties in 2 (3.7%) settings, and lose in 3 (5.6%) settings. This shows ORTB1 is fairly robust and the outperformance is stable.

5.2 The Impact of Budget Constraints

It is of great interest to investigate how the bid strategy adapts the changing of the budget constraints. In our experiment, we set the test budget as 1/64, 1/32, 1/16, 1/8, 1/4 and 1/2 of the original total cost in the history log respectively. Figure 9 depicts the percentage improvement on total clicks of ORTB1 over Lin with respect to the budget constraints. As we can see, (i) when the budget is quite low (e.g. 1/64 of the original total cost), the click improvement of ORTB1 over Lin is quite high (more than 45%). This indicates that our proposed bidding strategy performs particularly well under very limited budget conditions. Intuitively, when the budget is quite low (budget on per bid request B/N_T is low), a good bidding strategy should spend relatively low on each bid request. Compared with the linear Lin, ORTB1 allocates more budget on the low cost cases due to the concave form of the bidding function (see Figure 2(b)). This is sensible because from the winning rate functions in Figure 5 we know that lowering the high bid actually does not drop the winning probability too much. By contrast, highering the low bid a little will increase the winning probability a lot. (ii) When the test budget gets higher, the improvement percentage gets lower. This is reasonable: when there is more budget per bid request, the strategy will appropriately reallocate budget from the low cost cases to high cost cases because the high cost cases also mean high value (CTR). Thus the concave degree of the curve in Figure 2(b) will be lower. The curve will relatively approximate to (but not fully change to) the linear form. An extreme case is that when the test budget is set the same as the original total cost in the test log, the improvement is zero. This is because under such condition



Figure 10: Performance on different measures with different budget conditions.

for every bidding strategy (except Mcpc), just bidding as high as possible will exactly run out the budget and get every impression and click in the test log.

5.3 Clicks vs. Impressions

Figure 10(a) and 10(b) show the total clicks and eCPC of each bidding strategy. We observe that both click number and eCPC increase as the budget increases. This is reasonable because when the budget is low, the optimal setting for each bidding strategy will push the budget to the low cost cases, while the budget is higher, the left budget will be allocated on the relative expensive cases (but with lower ROI), which will higher the eCPC. Mcpc is unaware of the budget condition, and its slight eCPC fluctuation purely depends on the data. Another point is that when the budget is set very low (1/64, 1/32, 1/16 of the original spend), the eCPC of Lin and ORTB is lower than Mcpc, and when the budget increases over 1/4 of the original spend, the eCPC of Lin and ORTB starts to be higher than Mcpc.

Figure 10(c) and 10(d) plot the total impressions and CPM of each bidding strategy. We see that while our OR-TB strategies generate the highest clicks, they also produce comparable numbers of impressions against the others. This certainly benefits advertisers who aim to maximising their traffic (clicks) while still want to maintain a good exposure (impressions).

5.4 Parameter Tuning

As explained previously, although parameter λ can be directly solved numerically, for efficiency, in our experiment, we tune it for each campaign using the training data. Figure 11 and 12 shows the corresponding ORTB1 and ORTB2 click performance for the campaign 1⁷ when tuning its parameter λ . For each λ , we also try different value of the other parameter c, which makes multiple points to each x-value for volatility checking. However, practically we direct learn c to best fit the winning rate data for each campaign. From Figure 11 and 12 we see that when the given budget is low, the optimal value of λ is high. This is much reasonable. From Eq. (13) and (15) we can see the parameter λ controls the general scale of the bidding price: when λ is higher, the bid price gets lower. Thus it is understandable that when the budget is more limited, the general bidding price level

 $^{^7\}mathrm{The}$ trends on other campaigns are quite similar.



Figure 11: Parameter tuning on λ of ORTB1 under different budget conditions.



Figure 12: Parameter tuning on λ of ORTB2 under different budget conditions.

should get lower, which corresponds to the higher optimal value of λ .

5.5 Results for An Alternative KPI

As discussed in Section 3.1, our framework is flexible to incorporate different KPIs as our optimisation target. Apart from the number of clicks as the main objective studied in this paper, we also test an alternative KPI as the target by considering conversions, namely, a linear combination of the click number and conversion number, with the parameter k controlling the importance of conversion:

$$KPI = \#click + k \cdot \#conversion.$$
(18)

This objective is practically useful [1] since the conversion is a quite important measure for advertisers [32], and such linear combination can address the sparsity problem of conversion counting [26]. In our dataset, only campaign 2 and 8 have sufficient conversion records. Thus we choose them as our optimising campaigns and we set k = 5 in our experiment. Specifically, we use two logistic regression models to learn and predict the CTR and CVR (pCTR and pCVR) for each bid request and the pKPI can be calculated by pKPI = pCTR + $k \cdot pCVR$, which is the value of θ in the bidding function.

Figure 13 gives the overall KPI performance and the specific clicks/conversions obtained by each bidding strategy⁸. From the results we see our ORTB strategies still outperform other compared bidding strategies, which verifies the effectiveness of ORTB on the alternative KPI optimisation. Particularly, on 1/64 budget condition, ORTB2 achieves higher KPI and conversion numbers than ORTB1, this is mainly because the winning function 2 fits these two optimised campaigns better than winning function 1.





Figure 13: Performance comparison with the new KPI as optimisation target.



Figure 14: Relative performance for online test.

6. ONLINE EVALUATION

With the strategy of emphasising more on lower CPM impressions than the linear strategy, we conduct an online experiment on iPinYou Optimus platform, which is currently the largest DSP in China. We test on three campaigns during consecutive three days in December 2013. Another two compared bidding strategies are: (i) White list (White), which keeps a list of feature rules and bids much high only when the case satisfies any of the white list rules; (ii) Lin as we discussed before. The target KPI in the optimisation is click. Compared with Lin, White acts like the strategy with a step function in [32]. For each bid request, each of the three algorithms has the equal possibility to perform the bidding of iPinYou DSP. The performance comparison with various measures was reported in Figure 14. Due to the data sensitivity, we follow [40] to only present the relative performance here.

From the comparison we can have following observations. (i) ORTB bids much more cases than the other two algorithms (bids lower than the auction reserve price are not counted), and gets the most impressions, clicks and conversions. Also ORTB achieves the lowest eCPC, which indicates it is also the most cost effective algorithm. (ii) ORTB obtains the lowest CPM. These two points show ORTB allocates more budget to the cheap cases. As a result, ORTB bids more cases with lower CPM. (iii) Because of the low CPM on low pCTR cases, ORTB has the lowest auction winning rate and CTR, but this is not that important since our target is the total click number. (iv) White acts oppositely to ORTB. It only bids high on a subset of cases, which results in its low bidding number and high CPM. Due to the carefully made white list, the cases matching the white list do have high CTR. (v) Lin mostly plays moderately, which is in our expectation. In sum, the online test supports our proposed bidding strategies which allocate more budget on the cheap cases to perform more bids with lower CPM.

7. CONCLUSIONS AND FUTURE WORK

In this paper, we proposed a novel functional optimisation framework to find the optimal real-time bidding strategy for display advertising. The derived optimal bidding functions are non-linear and concave with respect to predicted KPI, which to our knowledge has not been studied in the previous literature about real-time bidding strategies. In our experiments, we compared our bidding strategies with other baselines and state-of-the-art bidding strategies under different budget and KPI settings. Apart from the theoretical soundness, the offline and online experiments show that our proposed optimal bidding strategies are also practically the most effective.

In the future work, we plan to further investigate on the bidding function which directly covers the bid request features, not the KPI alone. We also intend to extend our optimisation framework to cover the following three scenarios: (i) As DSPs or ad agencies, how to optimise the overall performance over several advertisers they are managing? Combining auction theory [24] with our functional optimisation framework would possibly provide a sound solution. (ii) Dynamic approaches [3] which tune the bid according to the specific current performance will be considered in our framework. (iii) Modelling the uncertainty and the risk of the bidding. The uncertainty in the CTR prediction, bidding and auction process should be carefully studied to build risk-aware bidding strategies.

8. **REFERENCES**

[1] DSP bidding optimization task.

- http://contest.ipinyou.com/rule.shtml. Accessed: 2014-02-06.
 [2] Z. Abrams. Revenue maximization when bidders have budgets. In SODA, pages 1074-1082, 2006.
- [3] K. Amin, M. Kearns, P. Key, and A. Schwaighofer. Budget optimization for sponsored search: Censored learning in mdps. In UAI, 2012.
- [4] A. Animesh, V. Ramachandran, and S. Viswanathan. Online Advertisers Bidding Strategies for Search, Experience, and Credence Goods: An Empirical Investigation. In Second Workshop on Sponsored Search Auctions. EC, 2005.
- [5] C. Borgs, J. Chayes, N. Immorlica, K. Jain, O. Etesami, and M. Mahdian. Dynamics of bid optimization in online advertisement auctions. In WWW, pages 531–540. ACM, 2007.
- [6] A. Broder, M. Fontoura, V. Josifovski, and L. Riedel. A
- semantic approach to contextual advertising. In *SIGIR*, 2007. [7] A. Broder, E. Gabrilovich, V. Josifovski, G. Mavromatis, and
- [7] A. Broder, E. Gabrilovich, V. Josifovski, G. Mavromatis, and A. Smola. Bid generation for advanced match in sponsored search. In WSDM, pages 515–524. ACM, 2011.
- [8] S. Cetintas, D. Chen, and L. Si. Forecasting user visits for online display advertising. *Information Retrieval*, 16(3):369–390, 2013.
- [9] T. Chakraborty, E. Even-Dar, S. Guha, Y. Mansour, and S. Muthukrishnan. Selective call out and real time bidding. In Internet and Network Economics, pages 145–157. 2010.
- [10] Y. Chen, P. Berkhin, B. Anderson, and N. R. Devanur. Real-time bidding algorithms for performance-based display ad allocation. *KDD*, 2011.
- [11] N. Craswell, O. Zoeter, M. Taylor, and B. Ramsey. An experimental comparison of click position-bias models. In WSDM, pages 87–94. ACM, 2008.

- [12] Y. Cui, R. Zhang, W. Li, and J. Mao. Bid landscape forecasting in online ad exchange marketplace. In *KDD*, pages 265–273. ACM, 2011.
- [13] B. Dalessandro, R. Hook, C. Perlich, and F. Provost. Evaluating and optimizing online advertising: Forget the click, but there are good proxies. 2012.
- [14] N. R. Devenur and T. P. Hayes. The adwords problem: online keyword matching with budgeted bidders under random permutations. In *EC*, pages 71–78. ACM, 2009.
- [15] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. Technical report, National Bureau of Economic Research, 2005.
- [16] E. Even Dar, V. S. Mirrokni, S. Muthukrishnan, Y. Mansour, and U. Nadav. Bid optimization for broad match ad auctions. In WWW, pages 231–240. ACM, 2009.
- [17] J. Feldman, S. Muthukrishnan, M. Pal, and C. Stein. Budget optimization in search-based advertising auctions. In EC, 2007.
- [18] A. Ghosh, B. I. P. Rubinstein, S. Vassilvitskii, and M. Zinkevich. Adaptive bidding for display advertising. In WWW, 2009.
- [19] T. Graepel, J. Q. Candela, T. Borchert, and R. Herbrich. Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft's bing search engine. In *ICML*, pages 13–20, 2010.
- [20] R. Gummadi, P. Key, and A. Proutiere. Optimal bidding strategies in dynamic auctions with budget constraints. In *CCC*, pages 588–588. IEEE, 2011.
- [21] K. Hosanagar and V. Cherepanov. Optimal bidding in stochastic budget constrained slot auctions. In EC, 2008.
- [22] J. Jaworska and M. Sydow. Behavioural targeting in on-line advertising: An empirical study. In WISE. 2008.
- [23] B. Kitts and B. Leblanc. Optimal bidding on keyword auctions. *Electronic Markets*, 14(3):186–201, 2004.
- [24] V. Krishna. Auction theory. Academic press, 2009.
- [25] K.-C. Lee, A. Jalali, and A. Dasdan. Real time bid optimization with smooth budget delivery in online advertising. ADKDD, 2013.
- [26] K.-C. Lee, B. Orten, A. Dasdan, and W. Li. Estimating conversion rate in display advertising from past performance data. In *KDD*, pages 768–776. ACM, 2012.
- [27] R. McGill, J. W. Tukey, and W. A. Larsen. Variations of box plots. The American Statistician, 32(1):12–16, 1978.
- [28] A. Mehta, A. Saberi, U. V. Vazirani, and V. V. Vazirani. AdWords and Generalized On-line Matching. In FOCS, 2005.
- [29] S. Muthukrishnan. Ad exchanges: Research issues. In Internet and network economics, pages 1–12. Springer, 2009.
- [30] S. Muthukrishnan, M. Pál, and Z. Svitkina. Stochastic models for budget optimization in search-based advertising. In Internet and Network Economics. 2007.
- [31] L. Olejnik, T. Minh-Dung, C. Castelluccia, et al. Selling off privacy at auction. 2013.
- [32] C. Perlich, B. Dalessandro, R. Hook, O. Stitelman, T. Raeder, and F. Provost. Bid optimizing and inventory scoring in targeted online advertising. In KDD, pages 804–812, 2012.
- [33] F. Radlinski, A. Broder, P. Ciccolo, E. Gabrilovich, V. Josifovski, and L. Riedel. Optimizing relevance and revenue in ad search. In *SIGIR*, pages 403–410, 2008.
- [34] M. Richardson, E. Dominowska, and R. Ragno. Predicting clicks: estimating the click-through rate for new ads. In WWW, pages 521–530. ACM, 2007.
- [35] J. Wang, S. Yuan, X. Shen, and S. Seljan. Real-time bidding: A new frontier of computational advertising research. In *CIKM Tutorial*, 2013.
- [36] S. Yuan, B. Chen, J. Wang, P. Mason, and S. Seljan. An Empirical Study of Reserve Price Optimisation in Real-Time Bidding. In *KDD*, 2014.
- [37] S. Yuan and J. Wang. Sequential selection of correlated ads by pomdps. In CIKM, pages 515–524, 2012.
- [38] S. Yuan, J. Wang, and X. Zhao. Real-time bidding for online advertising: measurement and analysis. In ADKDD, 2013.
- [39] W. Zhang, T. Chen, J. Wang, and Y. Yu. Optimizing top-n collaborative filtering via dynamic negative item sampling. In *SIGIR*, pages 785–788. ACM, 2013.
- [40] W. Zhang, Y. Zhang, B. Gao, Y. Yu, X. Yuan, and T.-Y. Liu. Joint optimization of bid and budget allocation in sponsored search. In KDD, pages 1177–1185. ACM, 2012.
- [41] Y. Zhang, W. Zhang, B. Gao, X. Yuan, and T.-Y. Liu. Bid keyword suggestion in sponsored search based on competitiveness and relevance. *IPM*, 2014.
- [42] Y. Zhou, D. Chakrabarty, and R. Lukose. Budget constrained bidding in keyword auctions and online knapsack problems. In Internet and Network Economics, pages 566–576. 2008.
- [43] Y. Zhu, G. Wang, J. Yang, D. Wang, J. Yan, J. Hu, and Z. Chen. Optimizing Search Engine Revenue in Sponsored Search. In SIGIR, 2009.