



Advertiser Bidding Prediction and Optimization in Online Advertising

Panagiotis Spentzouris¹, Iordanis Koutsopoulos^{1(✉)}, Kasper Grud Madsen²,
and Tommy Vestergaard Hansen²

¹ Athens University of Economics and Business Athens, Athens, Greece
pspentzouris@gmail.com, jordan@aueb.gr

² Intelligent Banker, Odense, Denmark
kaspergsm@gmail.com, tvhansen@gmail.com

Abstract. We study the problem of optimal bid selection across ads and time, with the aim to maximize incoming click traffic to the advertiser's landing page, which is directly translated in maximizing revenue. A major novelty of our approach lies in using Machine Learning (ML) to build regression models out of available data for deriving for each ad the relations, *(i)* cost-per-click (CPC) charged by the platform versus bid, *(ii)* assigned ad position in the ad list versus bid value, and *(iii)* number of ad clicks versus its position. These regression models naturally reveal hidden trends that would have been otherwise unavailable to the advertiser, such as the bidding behavior of competing advertisers and quality scores of their ads. We then incorporate these relations into a convex optimization problem of budget allocation across ads and across time, the solution of which is the optimal bidding strategy of the advertiser. We validate our approach with real data provided by an online advertising company that is active in the banking sector. Our solution leads to substantial increase in the amount of inbound click traffic to the advertiser's landing page compared to other approaches that are either heuristic and data-agnostic or employ simple statistics on data.

1 Introduction

Online Advertising has evolved into a thriving business, with an average annual growth rate of 9.4%, and it is predicted to reach a total market size of \$142.5 billion by 2021. In web-search advertising, advertisers cast bids in order to have their ads displayed in prominent positions in an ad list, next to organic results of a search-engine platform. Each ad contains a number of keywords. Advertisers set a budget they are willing to spend within a certain period of time, along with their bid for each keyword. Then, an auction is run by the platform and determines the ads to project and their ranking in the list, as well as relevant charges to advertisers, in terms of a cost-per-click (CPC). The form of auction that takes place is the Generalized Second-Price (GSP) one [7], and the rank of an ad is determined by the product of its bid and quality score, the latter

being a cumulative estimate of the quality of the ad, including its design, text, graphics and landing pages.

Contrary to the vast amount of existing literature which considers the point of view of the advertising platform, *we take the pragmatic viewpoint of an advertiser* and study the problem of optimal bid selection across ads and across time, with the aim to maximize incoming traffic to the advertiser's page in terms of number of clicks, while we recognize that having data about other advertisers is practically infeasible. Maximization of inbound click traffic in the landing page is directly analogous to advertiser revenue, since a portion of the click traffic will end up completing a transaction and will produce revenue for the advertiser through a lead.

Each advertiser may leverage historical data about a number of quantities that are directly available to her, either through the advertising platform (e.g. Google AdWords), or through simple measurements at the advertiser side. These quantities are *(i)*, her bid values, i.e. how much the advertiser is willing to pay, *(ii)* the associated CPC values, i.e. what the advertiser actually pays in the end of the auction, *(iii)* the number of impressions of an ad, *(iv)* the average position of an ad in the auction over a certain time interval, and *(v)* the number of clicks received by the ad. The advertiser may then decipher the relationships between these measurable quantities and her bid. For example, the amount an advertiser actually pays for a certain ad position, i.e. the CPC value, depends on her bid, the quality score of the ad, as well as on the bids and quality scores of her competitors' ads. Also, the average position of an ad depends on its bid as well as the bids of competing ads. Finally, the number of clicks that an ad receives changes as a function of its average positions.

We capture the relations above by deriving regression models which may provide interesting insights, since they naturally contain in them hidden trends that would have been otherwise unavailable to the advertiser, such as the indirect impact of bidding behavior and quality scores of competing advertisers on the advertiser's ad position, CPC and number of clicks. The derived regression models provide the aforementioned relations in a simple analytic form that are then introduced to a convex optimization problem.

1.1 Our Contribution

Advertisers compete with each other for ad slots, over possibly several parallel auctions. Each advertiser has a set of ads, and each ad contains some keywords. For simplicity, we assume that each ad contains one keyword, and each keyword belongs to one ad. We study the problem of bid selection and allocation across ads and for different time intervals, in order to maximize the total incoming click traffic to the advertiser's landing page subject to a budget constraint (Fig. 1). The contributions of our work to the literature are as follows.

- We consider the pragmatic viewpoint of a single advertiser that aims to best utilize data available to her in order to tune her bids for ads so as to maximize incoming click traffic to her landing page.

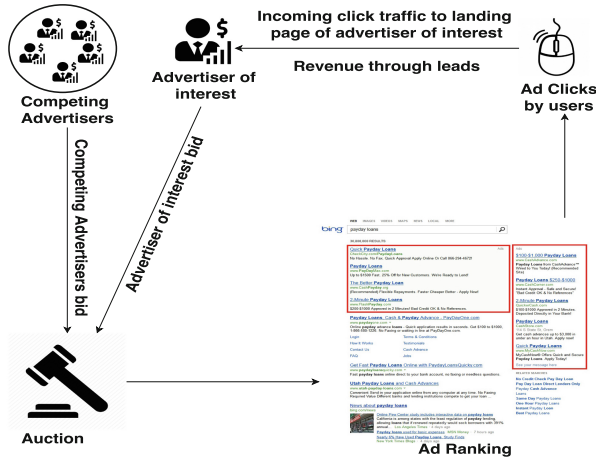


Fig. 1. A pictorial view of the sequence of steps in order for an ad to be projected and possibly clicked by users and produce revenue.

- We follow a data-driven approach, where we build linear regression models about the relations of position versus bid, and CPC versus bid, and number of clicks versus position for each ad. These regression models demonstrate hidden trends about the competitor’s behavior for different ads and different time intervals.
- We incorporate the analytical expressions of the regression models into a convex optimization problem of bid selection across ads and across time, for maximizing the total number of ad clicks (and therefore the click traffic in the landing page), subject to a constraint on a maximum budget to be spent in a specific time period. The solution of the problem reveals interesting insights about the relation of the bid and the parameters of the regression models.
- We validate our approach using real data. Our approach is shown to outperform other approaches that are either heuristic and data-agnostic or employ simple statistics on the data.

To the best of our knowledge, our work is among the first to follow the sequence of steps needed for an ad to be projected and possibly clicked through search advertising, and it uses machine-learning-based generated models to find hidden relations between the core ad attributes that are involved in these steps. We then incorporate them into a convex optimization problem the solution of which is the optimal bidding strategy for the advertiser of interest (Fig. 2).

2 Model

2.1 Setup

We consider an advertising platform (e.g. Google AdWords) and a set of advertisers. We take the perspective of a single advertiser with a set \mathcal{A} of N ads,

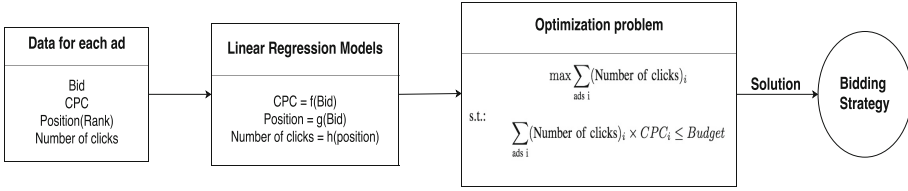


Fig. 2. Our methodology in order to find the optimal bidding strategy in terms of maximizing total click traffic for the advertiser.

where each ad has a quality score and contains some keywords. Without loss of generality, we assume that an ad includes exactly one keyword, so when we refer to an ad, we refer to its keyword. The advertiser has a budget B to be spent over a period of time for bidding for her ads. The advertiser decides on bid vector $\mathbf{b} = (b_1, \dots, b_N)$, where b_i denotes the amount to bid for ad (keyword) i . The advertiser participates in different auction processes, one for each ad, and she competes with the same or similar keywords of other advertisers in order to have her ads displayed in as high a rank as possible in the lists. Each such auction takes as input the bids and quality scores of competing advertisers for the same or similar keywords and decides on the ranking of ads and the CPC to be paid by each ad. Finally, ads are displayed in the list, and they may be clicked by users who view the ad list next to organic search results. Once an ad is clicked, the user is taken to the landing page of the corresponding advertiser, from which a lead (i.e. a product purchase) might occur within some time interval. The advertiser thus earns a certain amount of revenue by each lead event.

For the advertiser of interest, let b_i denote her bid for ad i . We implicitly assume that the auction for ad i is run once, hence b_i is the amount that is bid for ad i . Let p'_i be the position of ad i in the ranked list. Although we know high positions in the ranked list have small absolute value, e.g. positions 1,2, we define the position in such a way that higher positions in the list are associated with high values. The position of each ad i in the list is equal to its negative, i.e. $p_i = -p'_i$. Thus, if ad i is ranked first in the list, then $p_i = -1$, if it is ranked second, then $p_i = -2$, and so on. Also let c_i be the cost-per-click (CPC) that is paid for ad i . Finally, let n_i be the number of user clicks on ad i within some time interval.

There exist hidden dependencies between these quantities that we seek to capture in the sequel. First, the CPC value for ad i , c_i , depends on bid value b_i , i.e. the amount of money that the advertiser declares she is willing to pay for ad i . We denote this relation as $c_i = f_i(b_i)$, where $f_i(\cdot)$ denotes a continuous, and non-decreasing function. Along the same lines, position p_i depends on the advertiser's bid b_i , and let $p_i = g_i(b_i)$ denote that relation, where $g_i(\cdot)$ is a continuous, non-decreasing function. Finally, the number of clicks n_i of ad i depends on its position, i.e. $n_i = h_i(p_i)$ where $h_i(\cdot)$ denotes again a non-decreasing function.

2.2 Dataset and Linear Regression Models

The dataset that is readily available at the advertiser's side for each ad i is of the form:

$$\mathcal{D}_i = \{(b_i^d, c_i^d, p_i^d, n_i^d) : d = 1, \dots, M_i\}, \quad (1)$$

where M_i is number of data points available in the dataset for ad i . For each ad i , we are interested in using dataset \mathcal{D}_i so as to build models for approximating the relations $f_i(\cdot), g_i(\cdot), h_i(\cdot)$. We can use various machine-learning methods to derive models for the three relations above and then try to fit the models to the data in a way that reduces the total approximation error. Several methods are available, e.g. Neural networks, non-linear regression or linear regression models. In this work, we adhere to linear regression models, because our objective is to demonstrate the advantages of fitting a data-driven model in an optimization problem and extract the benefits of optimization, rather than comparing different ML methods. Further, linear regression models provide a simple means to have the relations above in an analytic form so as to feed them in the optimization problem, to be presented in the next section, and derive interesting insights about the solution. Hence we consider the following three models:

$$c_i = f_i(b_i) = \alpha_i b_i + \beta_i, \quad \alpha_i \geq 0 \quad (2)$$

$$p_i = g_i(b_i) = \gamma_i b_i + \delta_i, \quad \gamma_i \geq 0 \quad (3)$$

$$n_i = h_i(p_i) = \lambda_i p_i + \mu_i, \quad \lambda_i \geq 0 \quad (4)$$

where (α_i, β_i) , (γ_i, δ_i) and (λ_i, μ_i) are the parameters of the regression models to be computed from dataset \mathcal{D}_i for each ad i . The approach can be clearly extended in case the dataset is of different form.

The functions $f_i(\cdot), g_i(\cdot), h_i(\cdot)$ above contain some hidden trends about the competing advertisers' bidding behavior as well. In the next section, we will include these models in an optimization problem that will give the optimal bid allocation policy in terms of total number of clicks for the ads of the advertiser.

3 Problem Statement and Formulation

3.1 Bidding for Maximizing Total Number of Clicks for Ads

We are interested in finding the bid allocation policy $\mathbf{b} = (b_1, b_2, \dots, b_N)$ that maximizes the advertiser's total click traffic for all ads. This is a key objective for an advertiser, since a percentage of this traffic will end up completing a transaction and produce revenue through leads. Presumably the advertiser participates in several auctions, one for each ad, and that auctions are independent from each other. We formulate the optimization problem as follows:

$$\max \sum_{i=1}^N h_i(p_i), \quad (5)$$

subject to the constraint:

$$\sum_{i=1}^N h_i(p_i) f_i(b_i) \leq B, \tag{6}$$

with $b_i \geq 0$ for $i = 1, \dots, N$, where B is a fixed amount of budget to be spent over a specific period of time.

3.2 Solution

Since both the objective function and the function involved in the constraint are increasing over the vector bid \mathbf{b} , the constraint is satisfied with equality at the optimal solution \mathbf{b}^* . Thus, without loss of generality, it makes sense to consider the problem with an equality constraint. Let $\omega_i = \lambda_i \gamma_i \alpha_i$, and $\rho_i = \lambda_i (\gamma_i \beta_i + \delta_i \alpha_i) + \mu_i \alpha_i$.

This is a convex optimization problem, since both the objective function and the constraint are convex. Then we define the Lagrangian function, and after some algebraic manipulations we get the optimal solution as:

$$b_i^* = \frac{1}{2\nu a_i} - \frac{\rho_i}{2\omega_i} \tag{7}$$

Expression (7) shows that bid value b_i decreases as a_i (the slope of the straight line $c_i = f_i(b_i)$) increases. Thus, if the CPC value of an ad increases with a high rate, then our algorithm will raise the bids of alternative ads whose CPC value grows at a slower rate. On the other hand, as the parameter $\omega_i = \lambda_i \gamma_i \alpha_i$ increases, the bid value b_i increases. Parameter ω_i is the product of the slopes of the straight lines $c_i = f_i(b_i)$, $p_i = g_i(b_i)$, and $n_i = h_i(p_i)$. If an ad earns higher positions and its number of clicks increases at a high rate, then our algorithm will “promote” this ad and increase its bid value. A second observation is that the product $\lambda_i \gamma_i$ overpowers parameter α_i , which is inversely analogous to the bid value b_i . Therefore, the rate at which an ad’s number of clicks increases or the rate at which it earns higher average positions weighs more in our algorithm than the rate at which the CPC value increases.

Remark 1: An alternative formulation could take into account the quality q_i for each ad $i \in \mathcal{A}$ and include that as a factor in the objective function, which will now become $\sum_{i=1}^N q_i h_i(p_i)$ and expresses the total weighted number of ads clicks, where the weight is the quality of each ad. By a similar reasoning as above, the optimal bid value for ad i is given as,

$$b_i^* = \frac{q_i}{2\nu a_i} - \frac{\rho_i}{2w_i} \tag{8}$$

which tells us that the bid value of an ad grows proportionally to its quality score.

Remark 2: The optimization problem above considers budget allocation only across ads. We can have an enhanced formulation that would allow the advertiser

to spend different amounts of budget at different time intervals. For example, it is natural to assume that the auction structure (i.e the number of competing advertisers and their bids) and its outcome will be different at different times of the day or at different days of the week (e.g. week-days or week-ends) hence the bid would need to be adjusted as well. Let $t = 1, \dots, T$ denote an index for different time intervals, where T is the total number for such time intervals. We will have another problem definition, where $b_i(t), p_i(t), c_i(t), n_i(t)$ denote the bid, position, CPC and number of clicks of ad i at time interval t . Then, the budget allocation policy is given by vector $(\mathbf{b}(t) : t = 1, \dots, T)$, with $\mathbf{b}(t) = (b_1(t), \dots, b_N(t))$. Hence, the optimization problem becomes one of deciding on the bid policy that maximizes the total number of clicks across ads and across time, subject to a budget constraint.

4 Data Experiments

4.1 Dataset and Regression Models

We use a real dataset that was provided to us by an advertiser. The dataset consists of data from Google AdWords, and among other metrics it contains the bid values, CPC values, average position, and number of clicks for all keywords of the company. These metrics are available per day, for a time period of 6 months between August 2016 and February 2017. In our experiments we use 5 popular keywords/ads from the dataset since we are interested in demonstrating the benefits of the optimal policy. The properties of this policy are expected to hold for a greater number of keywords as well.

For the optimization problem (5)–(6), we need 3 prediction models for each keyword and each time interval, i.e 30 regression models in total. The three regression models that we build for a keyword i and a time interval $t = 1, 2$ are, one for the estimation of the CPC $c_i = f_i(b_i)$, one for the estimation of the average position $p_i = g_i(b_i)$, and one for the estimation of the number of clicks $n_i = h_i(p_i)$.

In order to build the prediction models we used linear regression techniques from Machine Learning. Let us assume that our training dataset, for the regression model $c = f(b)$ of a specific ad and a specific time interval is $L = \{b_n, c_n\}_{n=1}^M$, where each b_j (bid) is the data input and c_j (CPC) the corresponding output, j is an index for each data point in the data set, and M is the size of the dataset, i.e. the number of different entries available to train each model. We will choose a polynomial of first degree as a model:

$$f(b_j, w_0, w_1) = w_0 + w_1 b_j, \quad (9)$$

and in order to find the parameters w_0, w_1 that minimize the approximation error, we minimize over \mathbf{w} the cost function of least squares. We measure the performance of the model by means of Root-mean-square-error. Finally, we use cross-validation in order to choose an appropriate value for the regularization parameter.

4.2 Results

We run our experiment and compare the results for 4 different policies. We use the same set of 5 keywords for each policy, and initialize all bids at the value 0.1. We run experiments for different values of budget $B = \{100, 200, 500, 1000, 1500, 2000\}$. The policies we study and compare are:

- **Random bid allocation policy** (Policy 1) is a baseline approach which allocates randomly a portion of the budget across ads, without taking into consideration any information about the performance of the ad. For this policy, we choose uniformly at random one of the five ads, and then we raise this ad's bid by 5%. We repeat this procedure until the portion of the budget is exhausted. We run experiments for portions equal to 10%, 20%, and 30% of the budget. For each portion we run the experiment 10 times, and keep as result the average number of clicks from these runs.
- **Inversely proportional to CPC allocation policy** (Policy 2) allocates at each ad a portion of the budget that is inversely proportional to the average CPC value of ads, i.e. ads with high CPC get a smaller increase in their bids than ads with a lower CPC. Specifically, given the portion of the budget ϕ , which is the maximum amount of money that can be allocated across ads and across time without earning more clicks than the advertiser can pay, each ad's bid is raised by an amount $\frac{1/c_i}{\sum_i 1/c_i} \times \phi$.
- **Proportional to number of clicks allocation policy** (Policy 3) allocates a portion of the budget ϕ , which is the maximum amount of money that can be allocated across ads and across time without earning more clicks than the advertiser can pay, that is proportional to the average number of clicks that ads get, i.e. ads which get a large amount of clicks have their bids increased by a bigger amount than ads which get less clicks. The increase of each ad's bid is $\frac{n_i}{\sum_i n_i} \times \phi$.
- **Optimal policy** (Policy 4) which is the outcome of the optimization problem (5) and (6).

Comparison of Bid Allocation Policies. Figure 3 shows the total number of clicks to the advertiser's landing page under policy 1, for different budget amounts $B = \{100, 200, 500, 1000, 1500, 2000\}$, and different portions of the budget to be allocated across ads and across time equal to 10%, 20%, and 30%. The total number of clicks seems to be an increasing function over budget B for all cases except when the portion of budget is equal to 20%. Due to the randomness of the allocation, for this case and different budget amounts $B = 1000$, and $B = 1500$, it seems that the advertiser could earn more clicks with the smaller budget. Also, because of the random fashion in which we allocate a portion of the budget across ads and across time, it occurs that for a smaller portion of the budget we could have a better bid allocation, which could result in a larger amount of clicks to the advertiser's landing page.

In Fig. 4, we depict the total number of clicks to the advertiser's landing page under three different policies. As expected, the total number of clicks seems to

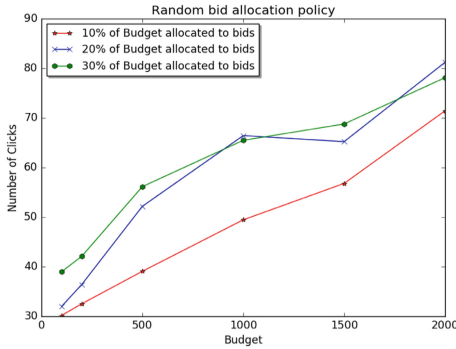


Fig. 3. Total number of clicks to the advertiser’s landingpage under random bid selection policy, for 3 different portions of the budget to be allocated across ads and time.

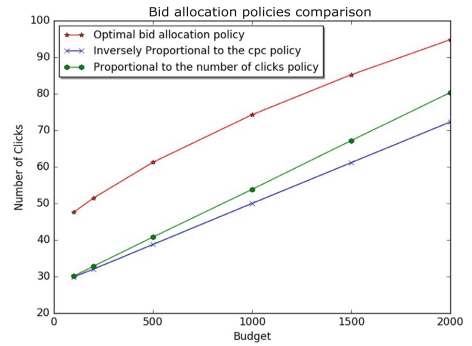


Fig. 4. Total number of clicks to the advertiser’s landing page under 3 different bidding policies.

	B = 500		B = 2000	
	Week-day	Week-end	Week-day	Week-end
keyword1	5.04	0.00	24.97	13.89
keyword2	19.63	19.99	41.49	43.12
keyword3	0.00	0.00	0.00	0.00
keyword4	22.13	0.00	41.31	15.57
keyword5	4.21	18.83	33.99	42.38

Fig. 5. Optimal bid selection for different amounts of budget B and different time intervals t .

be an increasing function over the budget B for all policies. We see that by building regression models about the relations and by integrating them into the optimization problem (5) and (6) we get the largest number of clicks, and thus larger revenue for the advertiser, compared to the other policies in the figure and the random policy in Fig. 3. In Fig. 5, we present the optimal bid selection for different amounts of budget $B = \{500, 2000\}$, and different time intervals $t = \{1, 2\}$, i.e. week-day and week-end.

We observe that policy 2 tries to balance bids across all ads, and as a result it increases the bad performing ads’ bids as well. We also see that under policy 3, there may be some cases of ads that are not currently getting many clicks but they have the potential to get more with less money than the current top performing ads. Nevertheless, this policy will ignore them because of the better performing in terms of clicks, ads. Despite this fact, policy 3 seems to perform better than policy 2. Also, we notice that policy 1 is performing similarly to the two other policies.

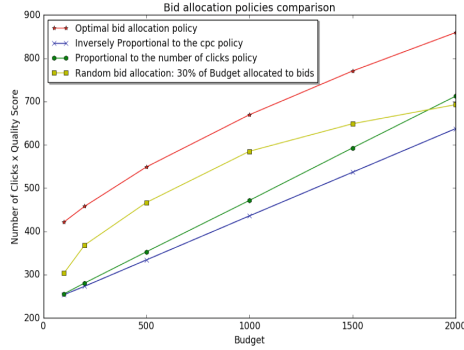


Fig. 6. Total number of clicks weighted by the quality score of each ad under 4 different bidding policies.

Bid Allocation Policies Weighted by Quality Score. In Fig. 6 we show the total number of clicks to the advertiser’s landing page weighted by the quality score of each ad, under 4 different policies. The ads’ quality scores are 9, 10, 8.8, 8, and 10 respectively. We see that the optimal policy outperforms the other 3 policies. Further, due to randomness, the random allocation can be benefited from increasing the bids of the ads with the biggest quality score values and thus outperform policy 3. Nevertheless, as the budget gets bigger policy 2 closes the gap in performance and finally outperforms policy 1 for a budget amount larger than 1700. This shows that for small amounts of budget, it is possible for a quality score-agnostic policy to increase the bids of the non-top performing ads more than the bids of the top performing ones. If the budget is big enough though to increase the bids of all ads, all policies achieve better performance. Finally, policy 2 which tries to balance bids across all ads is again the worst performing policy.

Optimal Bid Allocation for Maximum Total ad Quality. Figure 7, depicts the optimal bid selection for different budget values $B = \{500, 2000\}$, and different time intervals $t = \{1, 2\}$, i.e. week-day and week-end, for the objective that wants to maximize the total ad quality for the advertiser. We observe that for a small

	Quality Score	B = 500		B = 2000	
		Week-day	Week-end	Week-day	Week-end
keyword1	9	10.76	0.00	24.03	12.84
keyword2	10	28.52	29.39	44.69	46.51
keyword3	8.8	0.00	0.00	0.00	0.00
keyword4	8	25.33	0.00	36.68	10.89
keyword5	10	16.32	28.38	38.39	45.75

Fig. 7. Optimal bid selection for different amounts of budget B and different time intervals t , for the objective that takes into account the quality score of ads.

budget value $B = 500$ all the ads' bid values are increased compared to their values in Fig. 5. However, for a bigger budget value $B = 2000$, the bid values of the ads with the best quality scores, i.e. the two ads with quality score $q = 10$, are significantly increased compared to their bid values in Fig. 5, while the bid values of the rest are decreased. As expected, the optimal bidding policy for the objective that wants to maximize the total ad quality, decides to increase the bids of the ads that have the highest quality scores, while bid less for the rest. Furthermore, we observe that keyword 3 has a bid value $b = 0$ in both Figs. 5 and 7. The main reason for that is the slope parameter λ of its regression model $n = \lambda p + \mu$, which show that keyword 3 gets clicks at a slower rate than the other keywords and thus our algorithm will not decide to increase its bid.

5 Related Work

Real-time bidding (RTB) represents the cutting-edge frontier of the computational advertising research, and an thriving research area in advertising together with the display ad network and search-based keyword advertising. In RTB, advertisers bid for an ad impression when a user visits a webpage. A repository with a nice taxonomy of recent works can be found at [1].

Budget and Bid Allocation. In the context of display advertising, the authors in [4], model the state transition via auction competition, and they build a Markov Decision Process framework for learning the optimal bidding policy to optimize the ads performance. In [6], the authors define a revenue maximization problem, on an account level, by incorporating a probabilistic model to approximate the probability of winning a position given a price, and then they convert it into an integer optimization programming one. In work [3] they study the problem of finding a bidding strategy in real time mobile advertising. First, they model the win rate using a logistic regression model, and then take the derivative of win rate estimation to generate the distribution of the winning price, and use the expected value of the distribution under the bid price as the winning price estimate. Then a bidding strategy is actually an optimization function that takes the input of expected revenue if winning the auction, win rate and winning price estimate, and generate the final bid price according to some pre-defined objective functions.

In another relevant work [5], the authors try to find the optimal bidding function that maximizes key performance indicator (KPI), i.e. the total number of clicks or revenue in Real Time Bidding (RTB) display advertising. They find a function that returns the probability of winning given a bid value, based on historic data, and then based on the form of the winning rate function they derive empirically a simple function that returns the bid value. Then, they feed these functions into an optimization problem that returns the optimal bid allocation for the advertiser. In a recent work [2], the authors try to use simple heuristic bidding policies to increase the number of clicks, and they set the bid for each ad impression proportionally to the increase of user's conversion rates.

6 Conclusion

We study the problem of bid selection for a single advertiser across ads, each of which is represented by linear regression models that are derived from real data. Our method showcases the benefit of feeding data-driven models into optimization problems, and to the best of our knowledge constitutes the major novelty of this paper.

Possible future steps in this work include studying the problem of maximizing advertiser's revenue across the whole path up to conversion, and pursuing sensitivity analysis to model and take into consideration the inaccuracy of the models.

References

1. Paper Collection of Real-Time Bidding, Weinan Zhang's github page. <https://github.com/wnzhang/rtb-papers>
2. Xu, J., Shao, X., Ma, J., Lee, K., Qi, H., Lu, Q.: Lift-based bidding in ad selection. In: Proceeding of the Association for the Advancement of Artificial Intelligence (AAAI) Conference (2016)
3. Li, X., Guan, D.: Programmatic buying bidding strategies with win rate and winning price estimation in real time mobile advertising. In: Tseng, V.S., Ho, T.B., Zhou, Z.-H., Chen, A.L.P., Kao, H.-Y. (eds.) PAKDD 2014. LNCS (LNAI), vol. 8443, pp. 447–460. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-06608-0_37
4. Cai, H., Ren, K., Zhang, W., Malialis, K., Wang, J., Yu, Y., Guo, D.: Real-time bidding by reinforcement learning in display advertising. In: Proceedings of the ACM Web Search and Data Mining (WSDM) Conference (2017)
5. Zhang, W., Yuan, S., Wang, J.: Optimal real-time bidding for display advertising. In: Proceedings of the Knowledge Discovery and Data Mining (KDD) Conference (2014)
6. Zhang, W., Zhang, Y., Gao, B., Yu, Y., Yuan, X., Liu, T.: Joint optimization of bid and budget allocation in sponsored search. In: Proceedings of the Knowledge Discovery and Data Mining (KDD) Conference (2012)
7. Varian, H.: Online ad auctions. *Am. Econ. Rev.* **99**(2), 430–434 (2009)