Online Learning of a Ranking Formula for Revenue and Advertiser ROI Optimization

Or Levi
ebay/Marktplaats
olevi@ebay.com

ABSTRACT

A standard model for sponsored search comprises of ranking ads by their expected revenue, that is, the product of their bid price and estimated click-through rate (CTR). In this work, we introduce two complementary use cases for ranking ads at an online classifieds site and aim to optimize a ranking formula which extends the traditional one.

First, we address the task of ranking ads on the search results page for revenue optimization. While most works address this challenge by improving CTR estimation, we consider the effectiveness of the CTR estimation as a given and presume that if CTR estimation is somewhat ineffective, it can be compensated by applying a larger weight to the bid factor.

Second, we aim to improve advertiser return on investment (ROI) while keeping a similar level of revenues for ads ranking on the home page feed. To this end, we introduce into the standard ranking formula - a factor that favors ads with higher click-out rate and serves as an effective tie-breaker in cases of two competing ads with relatively similar revenue expectations.

To optimize the ranking formula, for each case, we propose an online learning procedure in a multi-armed bandit setting. Empirical evaluation attests to the merits of this approach compared to the existing ranking in production, which is based on the traditional formula, and validates our reasoning, first, regarding the relationship between CTR estimation effectiveness and the learned weights, and second, on the contribution of the click-out factor to increase in advertiser ROI.

KEYWORDS

Online Learning-to-Rank, Sponsored Search

ACM Reference format:

Or Levi. 2017. Online Learning of a Ranking Formula for Revenue and Advertiser ROI Optimization. In *Proceedings of the fisrt International Workshop on LEARning Next gEneration Rankers, Amsterdam, The Netherlands, October 1, 2017, (LEARNER '17),* 4 pages.

1 INTRODUCTION

Sponsored search is a major monetization source for commercial search engines. The ranking of sponsored ads determines which ads will be displayed and in which order, and thus plays a crucial part in optimization of revenues, user experience and advertiser efficiency.

Our work aims to optimize a formula for ranking ads at Marktplaats.nl, one of the largest sites in the ebay classifieds group. The site employs a pay-per-click advertising model, where advertisers bid for ads to be displayed and are charged by the bid amount, also known as cost-per-click (CPC), once a user clicks on an ad, which in turn generates a revenue for the site. Ads can appear on multiple devices, including desktop, mobile applications and tablets, and at different placements on the site, such as the top of the search results page, interleaved between organic results, and also on the home page feed. After clicking on an ad, users visit the view item page (VIP) where they can click-out to the advertiser website (Figure 1). Advertiser return on investment (ROI) is directly related to the cost per user click-out, also known as cost-per-action (CPA), calculated by dividing the total cost by the total number of click-outs. Hence, optimizing advertiser ROI is equivalent to minimizing the CPA. This setting reflects a potential conflict between the interests of the site and advertisers, which we aim to balance.

Similar to the standard model of sponsored search, ads in our system are ranked by multiplying their bid and estimated CTR. The same bid price applies for ranking an ad independent of query keywords and across all devices and placements. An ad's CTR estimation is calculated using past click-through data independent of query keywords, but separately per each device and placement, as those might exhibit inherently different user behavior. While bids are exact and given by advertisers, CTRs are difficult to estimate because clicks are rare events and new ads frequently enter the system. Specifically, in our setting, click-through data for tablets is relatively sparse, which can result in less effective estimation.

We introduce two use cases for optimization of ads ranking. In the first use case, we aim to optimize revenues for ads ranking on the search results page. A major line of research in sponsored search focuses on improving CTR estimation [1, 5, 6], which is also a major focus of our future work. In this paper, however, we take a different view. We treat the CTR estimation effectiveness as a given and acknowledge that it varies across the multiple devices and placements. Consequently, we re-examine the standard ranking formula and presume that applying a weighting scheme, where the bid factor is more dominant than the CTR estimation, can yield superior revenues. Moreover, we presume that, the less effective the CTR estimation, the larger the weight that should be applied to the bid, and inversely smaller weight to the CTR.

In the second use case, we consider ads ranking on the home page feed. The majority of the feed traffic is on the mobile apps, where click-out rates are significantly lower than desktop. Therefore, in this task we aim to improve advertiser ROI while maintaining relatively similar revenues. We propose that in cases where two competing ads have relatively similar revenue expectations, advertiser efficiency will be increased by favoring the ad with a higher historical click-out rate, and introduce this factor into the ranking formula.

LEARNER'17, October 1, 2017, Amsterdam, The Netherlands

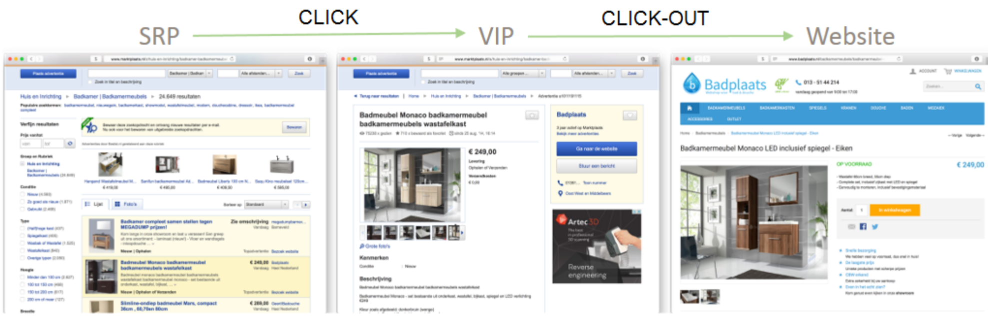

Figure 1: Marktplaats.nl sponsored search. Sponsored results are displayed at the top of the search results page (SRP). Clicking on a result opens the view item page (VIP), where a user might click-out to an advertiser website.

These two use cases reflect our aspiration to optimize near term revenues while improving advertiser ROI to sustain business relationship in the longer term.

Inspired by recent approaches that model online learning-to-rank as a contextual multi-armed bandit problem [2], we treat the challenge of optimizing weights for the ranking formula as an online learning process. Under this model, we attempt to learn the best action, that is, weights for the ranking formula, per each context, namely device and placement on the site, while observing revenues resulting from user clicks.

We show, through an empirical evaluation, that our approach in both use cases outperforms the existing ranking in production, which is based on the traditional ranking formula. The evaluation also validates the underlying premises of our approach. In the revenue optimization task, we point out to the correlation between the CTR estimation effectiveness and the weights learned across the different devices and placements. In the task of optimizing advertiser ROI while keeping revenues unchanged, we show the contribution of the click-out factor to increase in advertiser ROI.

2 RELATED WORK

Works in sponsored search that address the challenge of revenue optimization mostly focus on improving click-through rate estimation [1, 5, 6]. Predicting CTR for ads is typically based on machine-learned models trained using past click-through data. Examples of such models are logistic regression [1, 5], probit regression [6] and boosted trees [3]. These models employ multiple features that might affect the probability of a user clicking on an ad, such as textual match between the user's query and ad content, historical ad performance and personal user preferences. There has also been some work on employing learning-to-rank methods for the CTR estimation task [8], where a statistical model is learned offline.

Our work, on the contrary, treats the CTR estimation effectiveness as a given, and aims to maximize revenues through online learning of a ranking formula, that applies a larger weight to the bid factor to compensate for possibly ineffective CTR estimation.

The most relevant work to ours is that of Lahaie and Pennock [4]. It showed that applying an exponent, substantially smaller than one, to the CTR estimation, can yield superior revenue in equilibrium under certain conditions. There are several differences to our setting, such that it cannot be compared as a baseline. The main difference is that they study keyword auctions, where the bid and CTR estimation are keyword specific, and so is the fine-tuned exponent. In contrast, our ranking function employs weights for both the CTR and the bid factor, which are optimized through online learning, independent of query keywords, and in the context of each placement on our site. Consequently, we demonstrate through an empirical evaluation, that the optimized weights are not affected by keyword specific conditions, but rather by the degree of CTR estimation effectiveness, which varies across the different placements.

Advertiser efficiency in sponsored search is generally not considered as a separate objective. The common premise is that advertiser revenues are directly related to CTR and thus improving CTR estimation also increases efficiency with respect to advertisers. Wang et al. [7] recognized that the objectives of the site and advertisers are not always consistent and proposed to model ads ranking as a multi-objective optimization problem. However, similar to the common premise, they also used CTR as the objective for optimizing advertiser utility.

In our work, we focus on balancing revenues and advertiser ROI, wherein the latter is related more directly to cost-per-action (CPA) than CTR. CPA is affected to a large degree by the effectiveness of the advertiser view item page, but can also be improved directly through the ranking formula as we demonstrate in the next sections.

3 METHOD

We revisit the traditional ranking formula in sponsored search, where the ranking score of an ad is equal to its bid multiplied by its estimated CTR, and introduce the following weighting scheme:

$$CPC^{w_1} * CTR^{w_2} \tag{1}$$

such that $w_1 + w_2 = 1$.

For the revenue optimization task, we presume that if CTR estimation is somewhat ineffective, it can be compensated by applying a larger weight to the bid factor. Accordingly, we study sets of weights where $w_1 > w_2$. It can be expected that reducing the CTR weight in the ranking would result in a lower click-through rate. However, if the CTR estimation is indeed ineffective, this drop should be relatively mild and should be more than compensated by an increase in average CPC, yielding higher revenues.

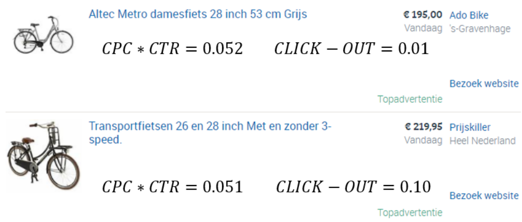

Figure 2: Example of two ads with relatively similar revenue expectation but large difference in click-out rate. In these cases, using the click-out factor we can display more ads with higher click-out rate to improve advertiser ROI.

For the task of balancing revenues and advertiser ROI, we employ the traditional product of CPC and CTR, which captures the expected revenue from an ad, and introduce a third factor into the formula to capture an ad's click-out rate:

$$CPC * CTR * w_1 + CLICK - OUT * w_2 \qquad (2)$$

such that $w_1 + w_2 = 1$. Our premise is that the click-out factor can serve as a tie-breaker in cases where two competing ads have relatively similar revenue expectations, as illustrated in Figure 2. An effective tie-breaker will allow us to keep revenues relatively stable while significantly increasing advertiser efficiency.

For each of these two ranking formulas, our aim is to find an optimal set of weights per each context, namely a certain device and placement on our site. We model this challenge as a contextual multi-armed bandit problem, where each slot represents a set of weights, and our objective is to play the best slot for each context. The best slot would be the one that maximizes the expected reward, that is, revenues in the case of Formula 1, or ratio of advertiser efficiency to revenues in the case of Formula 2. We also consider the effect on the relevancy of the results to users and impose a bound on a click-through rate drop that would be tolerated.

To optimize the weights, we devise a split test setting, with equally sized buckets, representing different sets of weights, and propose the following online learning procedure with a epsilon-first strategy of pure exploration followed by pure exploitation.

First, we set weights as per the above mentioned premises. Specifically, for the revenue optimization task, the weight of the bid factor is set to values in {0.6,0.7,0.8,0.9}, based on our presumption, that the bid factor should have a larger weight than that of the CTR estimation factor. For the task of balancing revenues and advertiser ROI, we set the weight of the click-out rate to values in {0.025,0.05,0.075}. These relatively minor weights reflect our premise that this factor should serve as a marginal tie-breaker.

Next, we observe each slot's performance over a one-week period, to account for seasonal factors, and select the weighting scheme that achieves the maximal improvement with statistically significant difference to the baseline, and such that click-through rate does not drop by more than 10%. Our main performance objective is revenue per mille impressions (RPM) in the revenue optimization task, and the ratio of CPA to RPM for advertiser ROI optimization. If no set of weights meets this criteria, we keep the baseline, which is the traditional sponsored search ranking formula.

Table 1: Performance on the revenue optimization task. Change in click-through rate (CTR), average cost-per-click (Avg CPC), revenue per mille impressions (RPM) and click-out rate between ranking using the learned weights and a baseline of the traditional ranking formula, for five different devices and two placements on the SRP, at the top and interleaved with organic results. Overall performance changes that are statistically significant are marked with '*'.

Device	Placement	CTR	Avg CPC	RPM	Click-Out
Desktop	Top	-1.10%	3.43%	2.29%	-1.05%
	Inter.	-2.87%	6.89%	3.82%	1.06%
iPhone-	Top	-1.92%	4.40%	2.40%	0.24%
App	Inter.	-8.41%	10.15%	0.89%	1.59%
Android-	Top	-4.05%	8.93%	4.52%	3.86%
App	Inter.	-7.43%	11.99%	3.67%	3.87%
iOS-	Top	-4.72%	10.17%	4.97%	2.07%
Tablet	Inter.	-2.09%	4.30%	2.13%	4.33%
Android-	Top	-4.39%	10.59%	5.73%	2.05%
Tablet	Inter.	-3.39%	6.19%	2.59%	4.71%
Overall		-3.18% *	6.43% *	3.05% *	1.66%

Finally, we also consider more fine-grained weights in adjacency to the best performing solution. If, for example, a bid factor weight of 0.6 has given the best performance, we consider weights of 0.55 and 0.65, and run another iteration with the best performing solution and the fine-grained weights.

4 EVALUATION

We evaluate our methods using the following online experiments on the classifieds site Marktplaats.nl. Each of the two use cases we introduced In Section 1 is evaluated separately.

First we optimize the weights for each device and placement following the procedure described in Section 3. Learning phase takes two weeks, one week for initial weights and one week for fine-grained weights. Training data overall includes more than 50 million impressions and more than 1 million unique ads. Subsequent to the learning phase, we run an online A/B test to evaluate the ranking formula with the learned weights against a baseline of the traditional sponsored search ranking formula, which is the existing method in production. Each alternative is assigned with an equal size of the traffic divided randomly by user ID. Lastly, we collect data for the evaluation over a one-month period and report the following measures: revenues per mille impressions, click-through rate, average cost-per-click and click-out rate. Statistical significance of performance differences is determined using a two tailed paired t-test with $p = 0.05$.

Table 1 presents the performance on the revenue optimization task. As expected following the overweighting of the bid factor, click-through rate drops for all the devices and placements, but this is more than compensated by an increase in average CPC, such that overall revenues increase. We see an increase in RPM across all the devices and placements, contributing to a statistically significant increase of 3% in overall RPM.

Table 2: CTR estimation effectiveness, measured by AUC, and the best performing weight of the bid factor per each device and placement.

Device	Placement	AUC	Best Peforming Bid Weight
Desktop	Top	0.723	0.55
	Interleaved	0.698	0.6
iPhone App	Top	0.715	0.55
	Interleaved	0.678	0.65
Android App	Top	0.704	0.6
	Interleaved	0.634	0.7
iOS Tablet	Top	0.541	0.8
	Interleaved	0.552	0.8
Android Tablet	Top	0.571	0.8
	Interleaved	0.538	0.8

Next, we study the correlation between the effectiveness of CTR estimation per each device and placement, and the best performing weight of the bid factor. To evaluate the CTR estimation effectiveness we use the AUC measure calculated on past click data. Table 2 presents the AUC and best performing weight per each device and placement. We see a clear correlation between the two (-0.985 Pearson correlation); Specifically, the less effective the CTR estimation (lower AUC), the larger the weight of the bid factor. As expected, CTR estimation for tablets is substantially less effective than for desktop and mobile apps, due to click sparsity. Accordingly, they are assigned with the largest bid factor weights. We also see that the AUC for the interleaved results is generally slightly lower than that of the top results, and the learned bid factor for them is generally larger.

For the task of balancing revenues and advertiser ROI on the home page feed using ranking Formula 2, we present the overall RPM and CPA performance for three weights of the click-out factor (Figure 3). As expected, RPM drops in all the three alternatives, but the drop in CPA is much more substantial, which presents an attractive trade-off for improving advertiser ROI, especially in the low-weight alternative where revenues are essentially unchanged. Moreover, we can see that the larger the weight of the click-out factor, the larger the drop in CPA, or equivalently, the more advertiser ROI is improved.

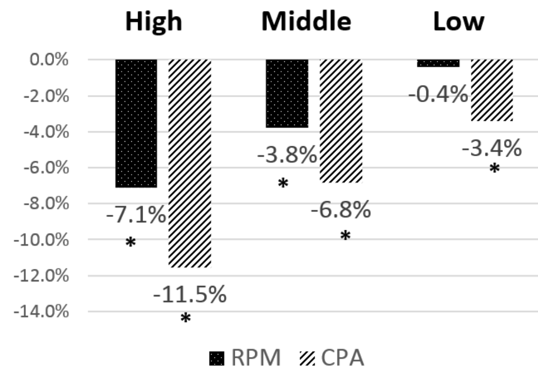

Figure 3: Performance on the task of balancing revenues and advertiser ROI. Change in overall revenue per mille impressions (RPM) and cost-per-action (CPA) for three weights of the click-out factor in Formula 2, compared to a baseline with no click-out factor. Note that CPA drop is equivalent to improvement in advertiser ROI. Performance changes that are statistically significant are marked with '*'.

5 CONCLUSIONS

We introduced two use cases for ads ranking at a classifieds website that complement each other as part of our aspiration to optimize revenues in the near term while improving advertiser ROI to sustain long term business. To address these challenges, we re-examine the traditional sponsored search ranking formula, introducing a weighting scheme where the bid factor is more dominant to compensate for CTR estimation ineffectiveness, in the first case, and introducing a click-out factor as a tie-breaker, in the second one. Subsequently, we proposed an online learning procedure in a multi-armed bandit setting to optimize the ranking formula for each case. Online experiments showed that our ranking formula with the learned weights outperforms a baseline of the traditional ranking formula, which is currently implemented in production. Furthermore, we demonstrated that the underlying premises of our approach are evident in the results. First, the less effective the CTR estimation for a specific device or placement, the larger the learned weight of the bid factor. Second, setting minor weights for the click-out factor can serve to increase advertiser ROI directly through the ranking formula, and the larger the weight, the bigger the increase to advertiser ROI.

As avenues for future work, we plan to extend the learning method to online Bayesian bandits. In addition, we plan to test more advanced CTR estimation methods and subsequently revisit this analysis. It can be expected that once more effective CTR estimation is reached, it would be less beneficial to outweigh the bid factor, if at all.

REFERENCES

[1] H. Cheng and E. Cantu-Paz. 2010. Personalized click prediction in sponsored search. In *Proceedings of the third ACM international conference on Web search and data mining, WSDM.*

[2] A. Grotov and M. de Rijke. 2016. Online Learning to Rank for Information Retrieval. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval.* 1215–1218.

[3] A. Kornetova I. Trofimov and V. Topinskiy. 2012. Using boosted trees for click-through rate prediction for sponsored search. In *Proceedings of the Sixth International Workshop on Data Mining for Online Advertising and Internet Economy, ADKDD.*

[4] S. Lahaie and D. M. Pennock. 2007. Revenue analysis of a family of ranking rules for keyword auctions. In *Proceedings of the 8th ACM Conference on Electronic Commerce.* 50–56.

[5] E. Dominowska M. Richardson and R. Ragno. 2007. Predicting clicks: estimating the click-through rate for new ads. In *Proceedings of the 16th international conference on World Wide Web (WWW-07).* 521–530.

[6] T. Borchert T. Graepel, J. Q. Candela and R. Herbrich. 2010. Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft's bing search engine. In *Proceedings of the 27th International Conference on Machine Learning.*

[7] J. Yan Y. Wang, B. Wei, Z. Chen, and Q. Du. 2012. Multi-objective optimization for sponsored search. In *Proceedings of the Sixth International Workshop on Data Mining for Online Advertising and Internet Economy.*

[8] J. Yang D. Wang J. Yan J. Hu Y. Zhu, G. Wang and Z. Chen. 2009. Optimizing search engine revenue in sponsored search. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval.* 588–595.