Optimizing AD Pruning of Sponsored Search with Reinforcement Learning

Yijiang Lian¹, Zhijie Chen², Xin Pei¹, Shuang Li¹, Yifei Wang³, Zhiheng Zhang¹, Zhipeng Tao¹, Liang Yuan¹, Hanju Guan¹

¹ {lianyijiang, peixin, lishuang15, zhangzhiheng02,taozhipeng, yuanliang, guanhanju}@baidu.com

² {bytechen.hit}@gmail.com

³ doxa.wang@pku.edu.cn

ABSTRACT

Industrial sponsored search system (SSS) can be logically divided into three modules: keywords matching, ad retrieving, and ranking. During ad retrieving, the ad candidates grow exponentially. A query with high commercial value might retrieve a great deal of ad candidates such that the ranking module could not afford. Due to limited latency and computing resources, the candidates have to be pruned earlier. Suppose we set a pruning line to cut SSS into two parts: upstream and downstream. The problem we are going to address is: how to pick out the best K items from N candidates provided by the upstream to maximize the total system's revenue. Since the industrial downstream is very complicated and updated quickly, a crucial restriction in this problem is that the selection scheme should get adapted to the downstream. In this paper, we propose a novel model-free reinforcement learning approach to fixing this problem. Our approach considers downstream as a black-box environment, and the agent sequentially selects items and finally feeds into the downstream, where revenue would be estimated and used as a reward to improve the selection policy. To the best of our knowledge, this is the first time to consider the system optimization from a downstream adaption view. It is also the first time to use reinforcement learning techniques to tackle this problem. The idea has been successfully realized in Baidu's sponsored search system, and online long time A/B test shows remarkable improvements on revenue.

KEYWORDS

Sponsored search, reinforcement learning, ad selection, global system optimization

ACM Reference Format:

Yijiang Lian¹, Zhijie Chen²,Xin Pei¹, Shuang Li¹, Yifei Wang³, Zhiheng Zhang¹, Zhipeng Tao¹, Liang Yuan¹, Hanju Guan¹. 2020. Optimizing AD Pruning of Sponsored Search with Reinforcement Learning. In *xx*. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION

Web search engine plays a vital role in our daily life for seeking information. Since submitted queries usually express a clear intent, search engine provides a good platform for advertisers to accurately target clients. On this platform, advertisers need to submit keywords, bids and creatives for their products and services. A *Keyword* is a short text to be matched towards query traffic, which

xx, xx, xx 2020. ACM ISBN 978-1-4503-9999-9/18/06...\$15.00 https://doi.org/10.1145/1122445.1122456 can be seen as a bridge to link queries and ads. A *Creative* is some texts that would be shown to users, which generally includes a title and a description (Fig. 1). And the *Bid* is used to express the advertiser's value for this keyword traffic.



Figure 1: A typical shown ad.

There are three main modules in a sponsored search engine (SSE): keyword matching, ad retrieving and ranking. When a query is issued, the *keyword retrieving* module would retrieve all the semantically related keywords, and the *ad retrieving* module would retrieve all the ads of the advertisers who have bought the keywords, filter out geographically or temporal conflicted ones and equip them with compatible styles. Then, these fully equipped ads would be gathered to go through the *ranking* module, where lots of model predictions (like CTR (click through rate), relevance, etc.) and filtering rules are conducted, then an auction (like GSP) is carried out for the remaining ads. Finally the winners would be shown to users.

A serious problem occurring in ad retrieving is that the candidates grow exponentially. For example, a query with high commercial value can retrieve 100 matching keywords, each keyword may be bought by 10 advertisers, then each advertiser might design 10 different creatives for this unit, and 10 different display styles can be chosen. That is, we would get 10⁵ full ad candidates in this case. Directly sending these ads to the *ranking* module would consume a lot of computation time, which is unacceptable for an industrial online system. So we have to prune some candidates earlier during their expansion. Here we set our pruning line between expanding creatives and equipping styles. We refer to all the modules below the pruning line as the downstream system. Then a typical ad selection problem emerges: How to select K items from N candidates to feed into the downstream system such that total revenue can be maximized?

There are several challenges for this problem. Firstly, the real industrial downstream system is very complicated, which usually deals with ad CTR prediction, blacklist filtering, diversity filtering, ad quality checking etc. Without considering the downstream system, ads selected earlier might be filtered out a lot. Secondly, at this moment, the ad candidate is not complete as style information is unavailable yet. Thus we can not obtain a precise CTR estimation which is a key element for winning the auction. Thirdly, this is a NP-hard combinatorial optimization problem.

Some heuristic approaches are widely employed for tackling these challenges, such as sorting all candidates under the revenue estimation and selecting the top K items. Obvious disadvantages of this kind of solutions are that, it is impossible to design perfect rules due to the super complexity of downstream system, and quality of the whole selection such as diversity is easily overlooked. Even if we adopt supervised learning based deep neural networks to solve it, it is hard to perform well as we are lack of best training samples.

In this paper, we resort to a reinforcement learning approach. Using a model-free learning framework, the complex downstream system can be considered as a black-box environment and the agent is an ad selector. The agent sequentially selects K items and feeds ad queue into the downstream system. The downstream system would estimate the final revenue for the ad queue, then this final eCPM (estimated cost per mille impressions) would be taken as a reward signal to guide the selection policy. In summary, this work offers the following contributions:

- (1) We propose an ad pruning agent which adapts to the downstream system. We hope this idea would shed light on the future design of the industrial sponsored search system.
- (2) We propose a reinforcement learning based approach to tackling the ad selection problem, and this work has been successfully applied to Baidu's real sponsored search system.
- (3) For the concern of training efficiency and safety of the real online system, a simulator system for RL training has been devised and implemented.

2 RELATED WORK

The combination of Deep Learning (DL), known as Deep Reinforcement Learning (DRL), has led to great success, both in academic research and in industrial applications, such as games [9], finance [6], healthcare [6], as well as Google's machine translation system [13]. Recently, the utilization of deep neural networks into sponsored search systems has yielded great benefits, particularly for matching queries and bidwords in the semantic space [5, 12]. Advanced techniques such as generative sequence to sequence models have also been adopted to produce bidwords or match similar short sentences [7].

For the subset selection in ad pruning, previous studies rely on various hand-designed heuristics to approximate the solution. Recent advances in deep learning provides an elegant and efficient method to such combinatorial optimization problems [2, 4]. In particular, we model the ad subset selection as a sequential decision making problem, and learns to maximize the overall reward of the selected subset, which is estimated by probing the downstream system. Similar optimization view to ours, [3] proposes an active question reformulation agent which interacts with a black box QA system and learns to reformulate questions to elicit best answers from the downstream.

3 TERMINOLOGY

For clarity, we declare some terminology and notation here:

- (1) SHOW denotes the total shown ad counts on result pages;
- (2) CTR denotes the average click ratio received by the search engine, which can be formalized as #{clicks} #{searches}.
- (3) CPM denotes revenue received by search engine for 1000 searches, which can be formalized as <u>revenue</u> <u>#{searches}</u> × 1000.
- (4) eCTR denotes the estimated click through rate if an ad is shown.
- (5) eCPM denotes the estimated CPM. It equals to eCTR×Charge.
- (6) Bid is the price provided by an advertiser for a keyword.
- (7) Charge is the true expense after the auction. Using GSP, charge is less than Bid.

4 METHODOLOGY

4.1 **Problem Formulation**

Suppose there are *N* ad candidates provided by the upstream, denoted as $X = \{x_1, x_2, ..., x_N\}$, we are required to select best *K* ads $Y = \{y_1, y_2, ..., y_K\}$ from *X* and feed *Y* to the downstream to maximize the revenue.

We employ a model-free reinforcement learning approach for dealing with this problem, where the complicated downstream system is treated as a black box environment, and the agent sequentially selects K ads with one ad at each step. In particular, a Markov Decision Process is defined as follows (see Fig. 2). For each step t $(1 \le t \le K)$ the state variable s_t occupies the current unselected ad list X_t as well as the already selected ad list Y_t . All ads in X_t and Y_t are drawn from the original set X. At start time, i.e. t = 1, the unselected set equals to the origin input while the selected set is empty, i.e. $X_1 = X$ and $Y_1 = \emptyset$. Based on the current state s_t , each action a_t is to choose one candidate ad from the unselected list X_t , and append it to the selected list. The policy of our agent is a selection probability distribution over the whole unselected candidate set X_t . It is approximated by a neural network parameterized by θ , and denoted as $\pi_{\theta}(a_t|s_t)$. When the selection is finished, these K ads in Y would be sent to the downstream system for further ranking and auction. Finally, ads won out in this competition would be shown to users. We denote the selection route as τ , and the final winner ads as $Z = \{z_1, z_2, \dots, z_M\}$. Then eCPM of Z are used as the whole reward of this episode, that is,

$$r(\tau) = \sum_{i=1}^{M} \text{eCPM}_{z_i}.$$
 (1)

The objective function we are going to optimize is:

$$\max_{\theta} \left\{ L(\theta) = \mathbb{E}_{\tau \sim \pi_{\theta}} \left[r(\tau) \right] \right\}.$$
(2)



Sequential Decision Process

Figure 2: The sequential decision process of the pruning agent: The state comprises of the selected ads and unselected ones, and the action is to select one ad from the unselected set. In this illustration, the upstream provide N = 5 ad candidates, and K = 3 ads are picked out. At Step 1, the agent selects ad2, then ad2 is moved to the selected ad set. Similarly, ad4 and ad5 are selected at Step 2 and Step 3. Then output set $Y = \{ad2, ad4, ad5\}$ would be sent to the downstream.

4.2 Training Algorithm

 ∇

We use Policy Gradient (PG) [10] to solve the above problem. The optimization direction is decided by

and parameters are updated according to

$$\theta \leftarrow \theta + \alpha \nabla_{\theta} L(\theta) \tag{4}$$

where α is the learning rate. The training procedure is explained in Algorithm 1.

The state of the s	Algorithm	1:	Training	algorithm
--	-----------	----	----------	-----------

while Agent's parameter θ is not converged do					
sample a mini-batch of queries Q ;					
for each query $q \in Q$ do					
go through the upstream to get X_q ;					
if $ X_q > K$ then					
Agent: sequentially pick out K ads Y_q with					
parameter θ ;					
add selection path τ_q as $< q, \tau_q >$ to selection					
logs;					
send Y_q to the downstream ;					
Env: calculate the reward $r(\tau_q)$;					
add reward $\langle q, r(\tau_q) \rangle$ to reward logs;					
join $\langle q, \tau_q, r(\tau_q) \rangle$ and add to training dataset;					
end					
end					
sample (state, action, reward) from training dataset ;					
update parameters θ according to Equ (4) ;					
end					

Specially, we resort to our agent to pruning ads only if counts of the candidates provided by upstream is greater than *K*. The selection path τ_q by our agent would be logged in the form of $< q, \tau_q >$, and respective reward calculated by downstream is logged in the form

of < $q, r(\tau_q)$ >. The training samples are produced after joining these two logs.

One more problem is that the agent takes K selecting actions for the downstream, but only one reward for the whole sequence is available. If we make no distinction between these actions and give equal reward for each selection, the training will be inefficient. Here we adopt reward shaping [8] is to overcome this shortcoming, which has been applied in Machine Translation [1, 11]. In our scenario, the reward is set according to the eCPM contributions. It means that the reward for selecting a shown ad $z \in Z$ is set to eCPM_z, while the reward for selecting an unshown ad is zero. Additionally, since the eCPM differs a lot among different ads, to reduce this variance, logarithm transformation is adopted to smooth the original eCPM.

4.3 Training Architecture

People may argue that: why not use the real CPM as the reward rather than eCPM. RL algorithms commonly require a large number of interactions with the environment. On one hand using real CPM means that we have to wait for the real user's feedback, which may take a long while. And on the one other hand policy exploration on the real traffic probably greatly damnify the daily revenue especially for a commercial system. For efficiency and safety concern, we have built an industrial sponsored search system simulator for our training, which can offer us a reliable and dense reward estimation without doing any harm to actual revenue.

Fig. 3 illustrates the training architecture with the simulator. Firstly, the whole system has been cloned as a simulation environment, and each query issued to the real online system will be copied to the simulated upstream. Secondly, the simulator agent will exploringly select ads according to the current policy, action trajectory of the selection procedure as well as the estimated reward by the simulated downstream will be logged and stored as training samples. Thirdly, the offline trainer update the policy parameters with the training data, and push to the simulator agent in nearly real-time. With such a training architecture, the online system, simulator system, and the trainer are decoupled, high-throughput and low-latency for efficient training can be assured. xx, xx, xx



Figure 3: Training architecture with simulator.

5 EXPERIMENTS

Setup Experiments are conducted on Baidu's sponsor search system. The model to approximate the candidate ad selection score at each step is a two-layer fully connected network. Final selection probabilities are obtained through a *softmax* layer. Specifically, in our experiment, maximum number *N* of candidates is set to 1000, and the selection count *K* is at most 100. Through grid search of related parameters, batch size during training is set to 128, and for each query *q* in the batch, we sample 50 states from the decision route τ_q . Adam optimizer is adopted to update the agent parameters with $\alpha = 0.99$, $\beta = 0.999$, and initial learning rate 10^{-4} .

For input features, two categories for each candidate ad are designed: static and dynamic. Static features describe either the ad itself, or the query/user itself, or ad-query properties, and dynamic features characterize attributes related to the already selected ad collection. That is, the difference lies in whether it changes with the sequential selecting process. The dynamic features aim to help maximize cumulative revenues, as well as maintain the diversity of advertisers. Main features we used are listed in Table 1.

Table 1: The main features.

Туре	Feature Name
static	estimated click through rate
static	query-ad relevance
static	bid
static	pre-trained query embedding
static	pre-trained ad embedding
static	if this ad is compatible with strong style
dynamic	if the same advertiser's ad has been selected
dynamic	accumulated eCPM within same advertiser

Baselines We compare our approach with two baselines. They are commonly used in industry and also adopted in our sponsored search system before this method. The first ranking baseline is by expected charge from the advertiser which equals to estimated clickthrough rate (eCTR) times bid. The second is ranked by estimated achievement from the downstream system, which is calculated by expected charge times the show probability (srq). The show probability is added due to the limited advertising position and multiple factors effect the eventual show results. Here we approximate the show probability of an ad to its accumulated show proportion in

the past 30 days. In the following section, we denote these two baselines by **eCTR*bid** and **eCTR*bid*srq**.

Results We show the key online A/B testing results in Table 2. Three most important indicators of our method are compared with the baselines. They are click count per thousand searches (CTR), shown ad count per thousand searches (SHOW), and revenue per thousand searches (CPM). The results demonstrate that our method achieves positive improvements over all three concerned indicators with both the baseline methods. Considering that ad supplies are stable, increment in SHOW shows that our agent does select better ad candidates which are more compatible with the downstream system. In addition, the proposed method also gains significant CTR improvements, namely, 1.11% over **eCTR*bid*srq** and 2.21% over **eCTR*bid**. It denotes that the newly shown ads are greatly accepted by users. Owing to the consistency of our pruning agent and the down-stream system, as well as the end user preference, we finally achieve a dramatic 1.95% improvement of CPM.

Table 2: Online A/B Test Results.

baselines	CTR	SHOW	СРМ
eCTR*bid	2.21%	1.27%	1.95%
eCTR*bid*srq	1.11%	1.17%	0.99%

It is also worth mentioning that this agent assumes the downstream system to be a completely black box, whereas it may be not. For example, in our system, it is clear that the downstream part would select at most x ads from a same advertiser. Even though that our model succeeds in learning to select diverse ads based on the given features and rewards, we attempt to directly apply this prior knowledge during both the training and online inferring. The result shows that this application brings us a significant return of performance (that is nearly saving 10% time in selecting actions) while without lessening the CPM revenue. We believe that making full use of similar prior knowledge in modeling is a valuable practice in industry systems.

6 CONCLUSION

In this paper we devise a reinforcement learning method to address the efficient ad pruning problem in the real sponsored search engine environment. We cut the system into upstream and downstream, and consider the complicated downstream as a black box environment. Then RL agent is trained to adaptively select *K* ads to maximize eCPM reward. Online long term A/B test on Baidu's sponsored search engine has showed that it greatly outperforms the CPM rule based selection approaches.

In our mind, making ad selection adapted to the downstream system is crucial. Many online search or recommendation service has a coarse-ranking and refined-ranking part, and the coarseranking's job is to select limited candidates from a vast mount of items and feed them into the downstream part. The proposed method can be easily transplanted to these scenarios.

REFERENCES

 Dzmitry Bahdanau, Philemon Brakel, Kelvin Xu, Anirudh Goyal, Ryan Lowe, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2016. An actor-critic algorithm for sequence prediction. arXiv preprint arXiv:1607.07086 (2016). Optimizing AD Pruning of Sponsored Search with Reinforcement Learning

- [3] Christian Buck, Jannis Bulian, Massimiliano Ciaramita, Wojciech Gajewski, Andrea Gesmundo, Neil Houlsby, and Wei Wang. 2017. Ask the right questions: Active question reformulation with reinforcement learning. arXiv preprint arXiv:1705.07830 (2017).
- [4] J. R. Caldwell, Richard A. Watson, C. Thies, and Joshua D. Knowles. 2018. Deep Optimisation: Solving Combinatorial Optimisation Problems using Deep Neural Networks. CoRR abs/1811.00784 (2018). arXiv:1811.00784 http://arxiv.org/abs/ 1811.00784
- [5] Jelena Gligorijevic, Djordje Gligorijevic, Ivan Stojkovic, Xiao Bai, Amit Goyal, and Zoran Obradovic. 2018. Deeply Supervised Semantic Model for Click-Through Rate Prediction in Sponsored Search. arXiv preprint arXiv:1803.10739 (2018).
- [6] JB Heaton, NG Polson, and Jan Hendrik Witte. 2017. Deep learning for finance: deep portfolios. Applied Stochastic Models in Business and Industry 33, 1 (2017), 3–12.
- [7] Mu-Chu Lee, Bin Gao, and Ruofei Zhang. 2018. Rare query expansion through generative adversarial networks in search advertising. In Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. ACM, 500–508.
- [8] Andrew Y Ng, Daishi Harada, and Stuart Russell. 1999. Policy invariance under reward transformations: Theory and application to reward shaping. In ICML,

Vol. 99. 278-287.

- [9] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature* 529, 7587 (2016), 484–489.
- [10] Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. 2000. Policy gradient methods for reinforcement learning with function approximation. In Advances in neural information processing systems. 1057–1063.
- [11] Lijun Wu, Fei Tian, Tao Qin, Jianhuang Lai, and Tie-Yan Liu. 2018. A study of reinforcement learning for neural machine translation. arXiv preprint arXiv:1808.08866 (2018).
- [12] Wenjin Wu, Guojun Liu, Hui Ye, Chenshuang Zhang, Tianshu Wu, Daorui Xiao, Wei Lin, Kaipeng Liu, and Xiaoyu Zhu. 2018. EENMF: An End-to-End Neural Matching Framework for E-Commerce Sponsored Search. arXiv preprint arXiv:1812.01190 (2018).
- [13] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google's neural machine translation system: Bridging the gap between human and machine translation. arXiv preprint arXiv:1609.08144 (2016).
- [14] Jun Zhao, Guang Qiu, Ziyu Guan, Wei Zhao, and Xiaofei He. 2018. Deep reinforcement learning for sponsored search real-time bidding. In Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, ACM, 1021–1030.