

AiAds: Automated and Intelligent Advertising System for Sponsored Search

Xiao Yang*, Daren Sun, Ruiwei Zhu, Tao Deng, Zhi Guo, Zongyao Ding, Shouke Qin, Yanfeng Zhu
Baidu Inc.

{yangxiao04,sundaren,zhuruiwei,dengtao02,guozhi,dingzongyao,qinshouke,zhuyanfang}@baidu.com

ABSTRACT

Sponsored search has more than 20 years of history, and it has been proven to be a successful business model for online advertising. Based on the pay-per-click pricing model and the keyword targeting technology, the sponsored system runs online auctions to determine the allocations and prices of search advertisements. In the traditional setting, advertisers should manually create lots of ad creatives and bid on some relevant keywords to target their audience. Due to the huge amount of search traffic and a wide variety of ad creations, the limits of manual optimizations from advertisers become the main bottleneck for improving the efficiency of this market. Moreover, as many emerging advertising forms and supplies are growing, it's crucial for sponsored search platform to pay more attention to the ROI metrics of ads for getting the marketing budgets of advertisers.

In this paper, we present the AiAds system developed at Baidu, which use machine learning techniques to build an automated and intelligent advertising system. By designing and implementing the automated bidding strategy, the intelligent targeting and the intelligent creation models, the AiAds system can transform the manual optimizations into multiple automated tasks and optimize these tasks in advanced methods. AiAds is a brand-new architecture of sponsored search system which changes the bidding language and allocation mechanism, breaks the limit of keyword targeting with end-to-end ad retrieval framework and provides global optimization of ad creation. This system can increase the advertiser's campaign performance, the user experience and the revenue of the advertising platform simultaneously and significantly. We present the overall architecture and modeling techniques for each module of the system and share our lessons learned in solving several key challenges. Finally, online A/B test and long-term grouping experiment demonstrate the advancement and effectiveness of this system.

CCS CONCEPTS

• **Information systems** → **Sponsored search advertising; Computational advertising.**

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 the author(s) 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.

KDD '19, August 4–8, 2019, Anchorage, AK, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6201-6/19/08...\$15.00

<https://doi.org/10.1145/3292500.3330782>

KEYWORDS

sponsored search, automated bidding, intelligent targeting, intelligent creation

ACM Reference Format:

Xiao Yang*, Daren Sun, Ruiwei Zhu, Tao Deng, Zhi Guo, Zongyao Ding, Shouke Qin, Yanfeng Zhu. 2019. AiAds: Automated and Intelligent Advertising System for Sponsored Search. In *The 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '19)*, August 4–8, 2019, Anchorage, AK, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3292500.3330782>

1 INTRODUCTION

Sponsored search is an indispensable part of the business model in the modern online advertising market. According to the Statista [38], revenue in the search advertising segment amounts to 96 billion dollars in 2018. By determining the keyword searches that are most relevant to their business's offerings, advertisers create ads and bid on relevant keywords to place their ads in the search results. The display and position of the ads are determined by a real-time auction when users are searching for corresponding terms. Sponsored search provides considerable revenue for general search engine services such as Google and Baidu. It is a huge online market in which tens of billions of auctions are held every day and several different types of business offerings are distributed to various users.

Traditionally, the pay-per-click pricing model and the keyword targeting technology are the two keys to the business success of sponsored search. The pay-per-click model is the most common payment method, in which an advertiser pays a publisher only when the ad is clicked. Keyword targeting provides an accurate match between the search query and the advertising terms, in which advertisers bid on keywords that are related to their products or services to display their ads to the targeted audience on the search result pages. After more than 20 years, the traditional model of sponsored search gradually shows its shortcomings and limits.

Firstly, keyword targeting requires that the advertiser should select plentiful keywords relevant to their business to increase the coverage of related and targeted search traffic. As users can express their search intent in a variety of different queries, it is challenging for an advertiser to find all the terms relevant to their offer from this huge inventory of possible terms. Due to the limit of exploration capability and knowledge about the broad scope of different keywords, most advertisers can only bid on a handful of relevant keywords which lead to insufficient advertising effect. Though the sponsored system has provided some keyword recommendation tools and multiple match types such as exact, phrase and broad and so on, but the keywords and match type set by advertisers still play an important role in ad retrieval, and all the retrieved ads should be subject to these keywords. An ad will not be retrieved even if search

queries are related if the advertiser does not bid on corresponding keywords or set the match type correctly. As the biggest search engine in China, Baidu has billions of search queries and subsequent page views each day and the new queries grow dramatically. Massive search traffic and highly dynamic search intents pose huge challenges for advertisers in their manual selection of keywords. In this situation, the keywords selected manually by advertisers determine the upper-bound of commercial search volumes, sponsored search engines cannot achieve the actual global optimal matching between search queries and ads, the advertiser also cannot get sufficient ad impressions if he doesn't have an expert for selecting keywords.

Secondly, in the pay-per-click pricing model, since the keyword bids play a direct role in the ranking and pricing of ads, advertisers should bid on each keyword cautiously. Due to the trade-off between usability and pertinency, most of the sponsored search systems adopt the keyword-level bidding language, and advertiser sets the bid price for keyword to represent his/her willingness to pay based on the click value of this keyword. To keep and improve campaign performance, advertiser must adjust the bid prices frequently according to the objectives (KPIs) and various bidding feedback signals such as cost-per-click, ad impressions or clicks, budget data, and so on. As sponsored search market has a highly dynamic and competitive environment, it takes a long time to build a stable manual bidding strategy. The aforementioned keywords selection problem also brings more workloads to the bidding optimization task, when the advertiser selects more keywords, then more efforts need to be put into the bidding strategy. Advertisers from big businesses usually recruit search engine marketing (SEM) experts to optimize their campaigns, while maintaining plenty of keywords and their bidding optimization is non-trivial especially for small and local businesses that don't have dedicated marketing staff. Moreover, keyword-level bidding language is too coarse-grained to represent the real value of each search traffic, it only can give an average click value for a cluster of the search volume. Even for the same search keyword, different time or different locations or different users may generate different click values for advertisers. The desired bidding language should support adjusting bid prices based on where, when, and how people search, but it will bring the curse of dimensionality to the manual bidding optimization.

In the meantime, though the global digital advertising market is still growing, the percentage of digital ad revenue captured by search is falling. There are many emerging advertising forms such as social media advertising, video advertising, contextual or native advertising and so on. As these new digital media have massive active users and can provide novel advertising products, the search engine is no longer the dominated traffic source for online advertising. For the performance advertisers, they optimize for the sales, signups, or other so-called conversions generated directly from their ads, and they will choose to allocate their budget to multiple advertising media based on the return on investment (ROI) metric. In this competitive situation, the ROI or the cost-per-acquisition (CPA) of advertising becomes more and more important to digital media to attract advertisers. For the sponsored search market, the traditional metrics such as cost-per-click, ad clicks are very indirect for performance advertisers to control and optimize their conversions

and ROI, and the keyword-level bidding optimization is inefficient for increasing the advertising performance.

Thirdly, the creation of search advertising plays an important role in ad performance, especially impacts the click-through rate (CTR) of ad. Traditional format of search advertising is the textual ad which has three creation parts: a headline text or title, a display URL, and a description text. Advertisers can optimize their ad creations by using more attractive title or promotional description. Recently, as online ad offerings become increasingly complex, rich ads with features such as larger formats, reviews, maps, sitelink extensions, call or app buttons, images, and many other decorations that result in an advertiser having several possible ads with varying sizes and layouts. When the content and layout for ad creation become more diverse and rich, how to design better creations faces a larger combination and optimization space for advertisers. The sponsored search system provides many ad formats, advertisers should prepare relevant materials and select related formats to set the specific ad creations. Due to multiple content configurations and layouts available to advertisers, the manual selection or optimization of ad creations is very hard to advertisers.

To address these main challenges in sponsored search, we propose and build an automated and intelligent advertising system, called AiAds. As the above three problems are all the bottlenecks of manual optimization for advertisers, we use machine learning models to solve these problems by transforming the traditional manual targeting, bidding, ad creation tasks into automated tasks. Based on this system, automated bidding, intelligent targeting and intelligent creation are integrated to support a more intelligent advertising system, and the advertiser can entrust the performance optimization of ads to this system by only setting their target. In this paper, we focus on sharing our experience in building the AiAds system and report empirical results after deployed it in Baidu. The main contributions of this work are as follows:

- We present a straightforward bidding language and corresponding automated bidding strategy for advertisers to optimize their campaign performance directly. We show the basic data requirement and model architecture used in bidding strategy. The new bidding language and strategy also extend the traditional pay-per-click pricing model and bring new challenges in designing the auction mechanism.
- Based on the new bidding language, we break the limits of traditional keyword targeting method. By using the more straightforward retrieval and matching model, the system can optimize a more optimal matching and selection between search queries and ads in an end-to-end manner.
- We present a componentized framework for designing and generating ad creations which can use the materials to optimize the content and layout of advertising automatically.
- We conduct the online A/B test and long-term grouping experiment on live traffic of Baidu, all results show that the AiAds system can significantly increase the advertising performance, increase the revenue of the search engine and increase the user experience.

The rest of the paper is organized as follows. We start by introducing the related works in Section 2. We then introduce the overall systems architecture in Section 3. The bidding language and

automated bidding strategy are described in Section 4. In Section 5, we talk about the models used to deploy our solutions for intelligent targeting. The intelligent creation framework is described in Section 6. The experimental results are discussed in Section 7. We finally conclude this paper in Section 8.

2 RELATED WORK

As the sponsored search is a hot research direction, there has been much work and an abundant literature on the optimization of different aspects [15, 32] of the sponsored search systems.

For the keyword targeting and ad retrieval task, a vast amount of methods have been proposed. For example, how to improve the broad match of bid keywords for a given query was studied in [7, 9, 14, 46]. An adaptive algorithm was proposed [20] which could utilize arbitrary similarity functions and catch the dynamics in the broad match. Generating bid keywords for some given landing pages of the advertisers was discussed in [33]. The advertisability of tail queries in sponsored search system was studied in [31]. In addition, some works have been proposed for bid keyword recommendation. A novel algorithm for advertising keywords recommendation was presented in [44]. By leveraging the contents of Wikipedia. Recommending a group of relevant yet less-competitive keywords to an advertiser was proposed in [47]. Most of these works have considered keyword relevance as a key factor in their algorithms. However, these methods can not overcome the limitation of keyword-based ad retrieval. In order to enhance sponsored search ad retrieval, a number of extractive summarization techniques for landing pages were explored in [12]. A query-ad semantic matching approach based on embeddings of queries and ads was proposed in [19], and the embeddings were learned on user search session data in an unsupervised manner. Machine translation model [37] was used to translate a natural language query into a keyword. A collaborative filtering algorithm based on the bipartite graph was presented in [2]. Similarly, a network-based [42] ad retrieval framework was proposed. For our intelligent ad retrieval task, we use some models about mining and learning for heterogeneous networks [13], and leverage deep learning techniques for the semantic matching problem [35, 36] in the information retrieval and recommendation systems.

Bidding optimization has been well studied in sponsored search. A systematic exploration of a natural class of greedy bidding strategies was undertaken in [10]. Bid optimization and generation for advanced match was studied in [8, 14]. The joint optimization of campaign budget allocation and bid price setting was proposed in [45]. However, most previous works focused on the keyword-level auction paradigm, and the conversion or ROI metrics was not taken into the bidding strategy. A solution to automatically adjust the bid price for advanced matching based on conversion rate prediction was presented in [34]. Similarly, a bid optimizing strategy called optimized cost per click was proposed in [49]. A reinforcement learning based real-time bidding strategy for sponsored search was proposed in [48]. In the industry of sponsored search, there are also some bidding strategies tools [6, 18], such as Enhanced CPC, Target CPA, Maximize Conversions, and so on. The conversion-based bid strategies are more suitable for the performance advertisers to optimize their ROI.

For the rich advertisements in sponsored search, there are a lot of ad extensions or formats in Google [17] and Baidu [4]. The ad creation optimization was less studied by academia. But there was some research work about how to design the combinatorial auction mechanism for rich ads [3, 11, 22]. Two sampled rich search ads in Baidu are presented in Figure 1, there are text creations, images, download links, phone call button, sitelink extensions in the ads.



Figure 1: The formats of ads in Baidu.

3 SYSTEMS ARCHITECTURE

Before diving into how the system is built, let us first introduce the overall system architecture of AiAds. As showed in Figure 2, the whole system is composed of five modules: the unified data center, the basic models of search ads, the intelligent targeting service, the automated bidding engine, and the intelligent creation generator. The unified data center provides the basic data requirement for the downstream tasks and consists of several data sources such as ad click data, ad conversion data and search log data, and so on. The basic CTR and CVR model are built to model and predict the click-through rate and conversion rate for each ad. Some ad relevance models are built to quantize the relevance between queries and ads, and to optimize the user experience. When users submit search requests, the intelligent targeting service retrieves all the related ads directly, and the automated bidding engine produces bid prices for each ad based on the real-time models and the target of advertisers. Finally, based on the candidate materials and contents, the intelligent creation generator makes the combination and the generation of ad creations to display.

3.1 Basic Setting

We now describe the setting and notation we use throughout the paper. Our setting is a standard sponsored search. First, we introduce some notation. Assume that there is a set $N = \{1, \dots, n\}$ of n advertisers. For each advertiser i , he selects K_i related keywords and the k_{ij} indicates the j th keyword selected by advertiser i , and he creates A_i ads and the a_{ij} indicates the j th ad created by advertiser i . As already mentioned, the Pay-Per-Click pricing model is employed. For each k_{ij} , advertiser i has a private click value $v_i(k_j)$, expresses the maximum price per click he is willing to pay. In order to participate in the auction, advertiser i is required to submit a bid $b_i(k_j)$. $b_i(k_j)$ is a proxy of the value $v_i(k_j)$ but might

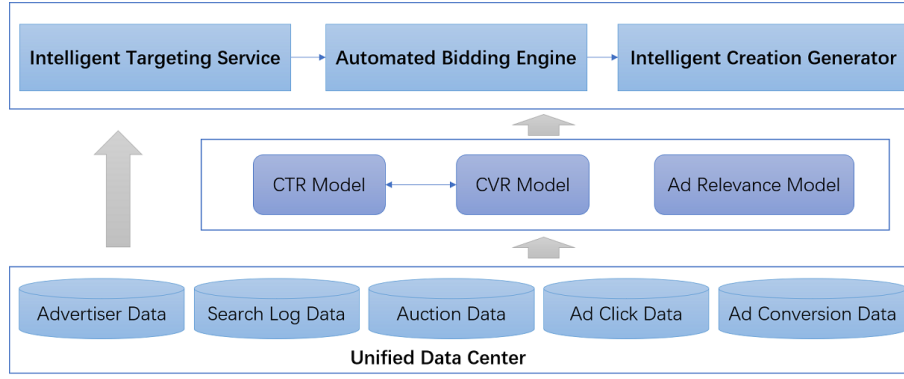


Figure 2: The Systems Architecture of AiAds.

not exactly equal $v_i(k_j)$ due to the strategic behaviors of the bidders. Submitting a bid $b_i(k_j)$ guarantees to advertiser i that he will not be charged a price higher than $b_i(k_j)$ per click. The vector of the advertiser i 's bids, b_i , will usually be referred to as his bidding strategy profile. When the displayed ad of advertiser i associated with k_{ij} is clicked, he will be charged a price $p_i(k_j)$, which is also called cost-per-click(CPC), and we have $p_i(k_j) \leq b_i(k_j)$.

CTR and CVR are two important parameters in the context of sponsored search. The click-through-rate(CTR) is the probability that a given ad will be clicked when displayed, and the conversion rate(CVR) is the probability that a ad conversion (A conversion is an action that's counted when someone interacts with advertiser's ad and then takes an action that defined as valuable to his business such as purchased a product, signed up for newsletter, called advertiser's business, or downloaded app and so on) be acquired after click this ad. So we can have that $CTR = \frac{\text{clicks}}{\text{impressions}}$ and $CVR = \frac{\text{conversions}}{\text{clicks}}$. We use $ctr(a_{ij})$ and $cvr(a_{ij})$ to represent the actual CTR and CVR of the ad a_{ij} , and the $pctr(a_{ij})$ and $pcvr(a_{ij})$ to represent the predicted CTR and CVR. Another important quantity is the cost-per-acquisition (CPA) which represents the average cost for each conversion, so we have $CPA = \frac{\text{cost}}{\text{conversions}}$.

Advertisers can be split into two categories: brand and performance. Brand advertisers aim for long-term growth and awareness, they have a mandate to meet a specific business goal—showing impressions to an audience, generating clicks, or maximizing revenue—driven by long-term considerations instead of immediate profit. Performance advertisers optimize the immediate tradeoff between value—measured as sales, sign-ups, or other so-called conversions generated directly from their ads—and cost. Return on investment(ROI) has been the standard metric for measuring this tradeoff across all types of advertising for decades. ROI measures the ratio of the profit obtained (“return”) to the cost or price paid (“investment”), i.e., the density of profit in cost: $ROI = \frac{\text{Revenue} - \text{Cost}}{\text{Cost}}$. Being a density metric, unconstrained maximization of ROI is not sensible, instead, performance advertisers come with an ROI constraint and maximize their revenue.

The performance advertisers are the main customers for sponsored search and many other digital advertising platforms. To prevent advertisers shifting ad budgets, advertising platforms should keep a competitive ROI metric.

4 AUTOMATED BIDDING ENGINE

In this section, we will cover the details of how we design and build the automated bidding engine to address the problems of traditional keyword-level manual bidding optimization.

4.1 The Bidding Language

In sponsored search, advertisers set and optimize their bids to achieve a specific goal for their business. As plenty of keywords should be maintained and highly dynamic and competitive auction environment, the keyword-level bidding language is very inefficient and brings great challenges to manual bidding optimization. Nowadays, for the performance advertisers, their goal is to maximize their profit with an ROI constraint. We can formulate the utility function of performance advertisers as:

$$\begin{aligned}
 U_i &= (\text{revenue} - \text{cost}) \\
 &= (\text{sale value} * \text{sales} - \text{conversions} * \text{CPA}) \\
 &= (\text{sale value} * \text{sale rate} - \text{CPA}) * \text{conversions} \\
 \text{s.t. } ROI_i &\geq \gamma_i
 \end{aligned} \tag{1}$$

And the sale rate = $\frac{\text{sales}}{\text{conversions}}$. For the ROI constraint, we can have:

$$\begin{aligned}
 ROI_i \geq \gamma_i &\Rightarrow \frac{\text{revenue} - \text{cost}}{\text{cost}} \geq \gamma_i \\
 &\Rightarrow \frac{\text{sale value} * \text{sale rate} - \text{CPA}}{\text{CPA}} \geq \gamma_i \\
 &\Rightarrow \text{CPA} \leq \frac{\text{sale value} * \text{sale rate}}{1 + \gamma_i}
 \end{aligned} \tag{2}$$

From the equation (2), the ROI constraint can be transformed into a target CPA constraint. So the utility function of performance advertisers is their profit subject to a target CPA.

In the practice, as the digital advertising market has more and more supplies such as social media and video media and both these media have ample traffic volumes, so the performance advertisers face the budget allocation problems. Rational advertisers will allocate their budgets based on the ROI of supply media, and the media with higher ROI and more conversions will get more budgets from the advertisers.

Therefore, for the real optimization goals of performance advertisers, the keyword-level click value based bidding language is too

indirect and coarse-grained, and advertisers should make lots of efforts in accumulation of clicks and calculation of CVR. To tackle this problem, we provide a straightforward target CPA based bidding language and advertiser can set the target CPA at campaign or ad group level directly, to replace the manual keyword-level bid. The target CPA is a straightforward representation of advertiser's ROI constraint, and the advertising platform should provide as many conversions as possible for the advertiser under this constraint.

4.2 Automated Bidding Strategy

Based on this new bidding language, advertisers only need to set and optimize their global target CPA for conversions, and the concrete real-time bidding for each auction can be done by the automated system. From the equation (1) and equation (2), we can conclude the optimizing objective of the advertiser, which is:

$$\begin{aligned} & \max\{\text{conversions}\} \\ & \text{s.t. } CPA \leq \text{target CPA} \end{aligned} \quad (3)$$

The revenue of sponsored search system can be formulated as:

$$\begin{aligned} \text{Revenue} &= \sum_{i=1}^M \text{click}_i * CPC_i \\ &= \sum_{j=1}^N \text{conversion}_j * CPA_j \end{aligned} \quad (4)$$

From the equation (4) and equation (3), we can conclude the optimizing objective of the automated bidding strategy, which is:

$$\begin{aligned} & \max\{\text{conversions}\} \\ & \text{s.t. } \min|CPA - \text{target CPA}| \end{aligned} \quad (5)$$

And the auction mechanism should ensure the IC(incentive compatibility) for expressing the target CPA, under the setting with multiple sellers and multiple bidders.

The equation (5) is also can be regarded as the optimizing objective of the AiAds system, and we can find that this objective is all-win to advertisers, sponsored search platform and users.

To optimize the equation (5), we propose a multiplicative automated bidding strategy to produce the real-time bidding (RTB) for each auction by combining multiple real-time bidding factor(also can be solved as a feedback control problem), which is:

$$RTB = CPA * pcvr * AF * BF * CF * Alpha \quad (6)$$

The objective of this bidding strategy is to minimize the CPA gap and to maximize the conversions. And each bidding factor is described as follows:

- **CPA**: The target CPA set by the advertiser.
- **pcvr**: The CVR of ad predicted by machine learning model.
- **AF**: The auction factor which is used to quantify the statistical gap between bid and CPC. As the bidding strategy is to optimize the conversions and make the average CPA close to target CPA, the statistical gap information can be used to increase proximity to the target value. We use a gradient boosting framework[25] to model the gap, the label is $\frac{\text{bid}}{\text{CPC}}$, and the features consist related statistics about the query, the advertiser and the auction context, other bidders' information cannot be used due to the incentive problem.

- **BF**: The budget factor which is used to spend budget smoothly over the time in order to reach a wider range of audience accessible throughout a day. A smart pacing approach[1, 26, 41] is used to adjust the bid price based on the prior performance distribution in an adaptive manner by distributing the budget optimally across time.
- **CF**: The calibration factor which is used to make the predicted CVR more close to the real CVR. The optimized isotonic regression[43] model is used to calibrate the *pcvr* values.
- **Alpha**: The alpha factor is a dynamic parameter generated by a reinforcement learning model. As the highly dynamic auction environment in a day, real-time feedback information must be used to boost the achievement of target CPA. We split the whole day into eight time buckets, and the optimization of daily achievement of target CPA can be modeled as an MDP. We design a general representation for states as $s = \langle \text{tid}, \text{sr}, \text{pg}, \text{cg}, \text{cd} \rangle$, where *tid* denotes the bucket number, *sr* denotes the spent ratio of budget, *pg* denotes the *pcvr* gap between the current bucket and the last one, *cg* denotes the CPC gap between the current bucket and the last one, and *cd* denotes the relative diff between current CPA and target CPA. The reward function $r = \min\{15, \frac{\text{target CPA}}{|\text{real CPA} - \text{target CPA}|}\}$ is used, and the action space is 31 discrete numerical ratios from [-3, 3] such as -3, -2.8, ..., 2.8, 3. We adopt a DQN algorithm similar to[29, 30] which employs a deep neural network with weights θ to approximate *Q* value. By the Alpha factor, we can model the bidding strategy as a dynamic interactive and sequential control process in a complex environment rather than an independent prediction or optimization process.

Based on this automated bidding strategy, fine-grained and auction-time bidding comes true in sponsored search, and abundant signals can be used in bidding optimization.

4.3 The CVR Model

From the equation (6), the *pcvr* play an important role in the automated bidding strategy. If the CVR models are only trained with samples of clicked impressions, the model will have poor generalization ability in plenty of un-clicked impressions such as new ads. To tackle this problem, we put the training of the CVR model and the training of CTR model together by using a multi-task model architecture and shares the lookup table for the embeddings of common features. The overall model architecture is presented in Figure 3. Moreover, as the conversion types are diverse, the network of CVR model also adopts a multi-task architecture.

4.4 Auction Design for ROI-constrained Advertisers

The basic setting and assumption of traditional auction design in sponsored search(such as GSP, VCG, Myerson, etc.) are that:

- One seller, multiple bidders.
- The utility function of bidder is quasilinear for single auction $U_i = (\text{click value} - \text{CPC}) * \text{CTR}$.
- One-shot auction or the different auctions are independent.

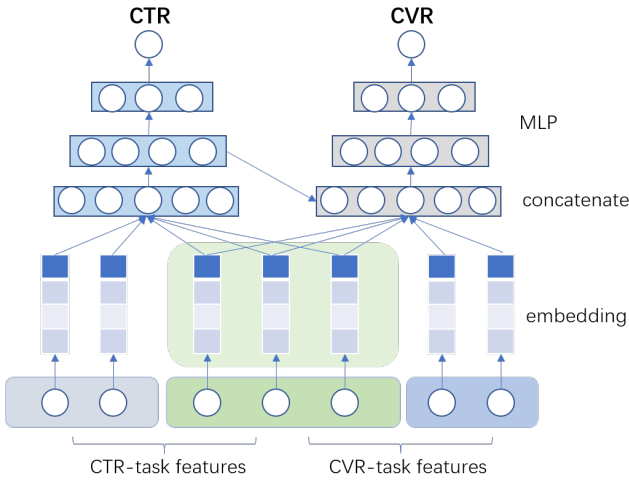


Figure 3: Architecture overview of CVR model.

In fact, since the utility function of performance advertisers is close to the equation (1), the traditional quasilinear utility function for single auction is no more hold. The new bidding language and bidding strategy also bring new challenges in designing the auction mechanism for sponsored search. Since in practice the ROI is computed in average over time, it creates a dependency between auction. Bidders are less sensitive to their obtained ROI in a single auction but are instead chiefly concerned about their expected ROI across many auctions. The real condition of auction design for ROI-constrained advertisers is that:

- Multiple sellers, multiple bidders.
- The utility function of the bidder is equation (1).
- Repeated and sequential auctions, auctions are dependent and context-aware.

For the ROI-constrained bidder, the optimal mechanism design of ad auction should be aware of the ROI metric. There are some work show[16, 23, 40] that the different points in designing auctions for ROI-Constrained bidders.

For the performance advertisers, the two main changes in our auction mechanism are:

- According to[16], keyword-based reserve prices are removed, instead, personalized and CVR-based reserve prices are set.
- To model and optimize the sequential auctions problem, we adopt the bank account mechanism[27, 28] framework to increase the efficiency of the allocation and the revenue of sponsored search platform. We keep a state variable(balance) b_t for each buyer based on his outcome of historical auctions, and use the b_t in the next auction with the balance-independence property, dynamic reserves are also managed via bank accounts.

5 INTELLIGENT TARGETING MODEL

The traditional keyword targeting model restricts the range of ad retrieval, given a search query, the system must retrieve the related keywords, and then retrieve the ads related to these keywords. This

two-stage retrieval procedure may lose many candidates, even for the advanced broad match type, the set of keywords selected by advertisers is still the bottleneck for ad retrieval.

Is the keyword really necessary for sponsored search? Based on the target CPA bidding language, the keyword-level bids are needless, so we can break the limits of keyword targeting method. With the help of automated bidding strategy, the keyword targeting and the match types are no longer necessary for advertisers. If the advertiser adopts the new bidding language, first, all the match type of his keywords will be extended to the advanced broad match, and then models that can retrieve related ads directly from the query can be used for intelligent targeting to break the limits of his selected keywords.

The intelligent targeting service is an end-to-end ad retrieval framework and can achieve direct matching from query to related ads. The overall architecture of the intelligent targeting service is presented in Figure 4, and two types of ad retrieval models are used in this service which can provide more ad candidates by utilizing diverse data sources.

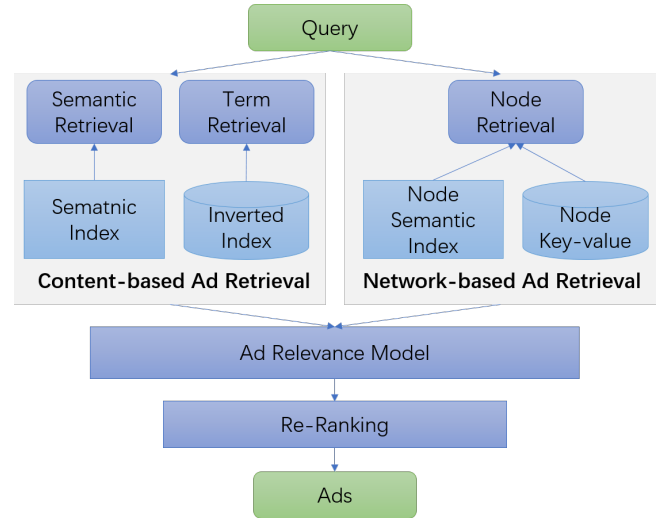


Figure 4: Architecture overview of the intelligent targeting service.

5.1 Network-based Ad Retrieval

To retrieve related ads for search queries, a straightforward method is to utilize the multiple relationships between queries and ads. Based on the historical click data, a heterogeneous network can be constructed to encode structured information of multiple types of nodes and links. A snippet of this multi-relational network is presented in Figure 5. There are four types of nodes and seven types of relationships in this heterogeneous network. Based on the structure of this network, we can find the potential paths from queries to unconnected ads, such as the node Query4 and the node Ad4.

A heterogeneous network is defined as a graph $G = (V, E, T)$ in which each node v and each link e are associated with their mapping functions $\phi(v) : V \rightarrow T_V$ and $\phi(e) : E \rightarrow T_E$, respectively. T_V and

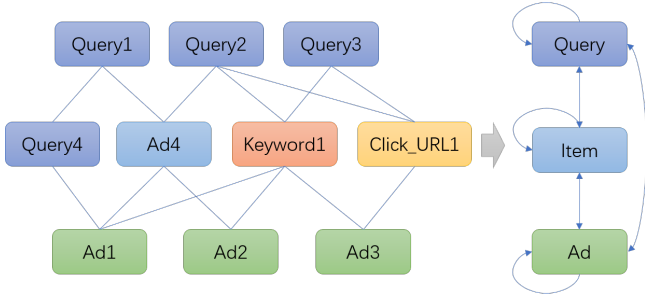


Figure 5: A snippet of the multi-relational networks.

T_E denote the sets of object and relation types, where $|T_V| + |T_E| > 2$. A meta-path[39] P is a path defined on the network schema $T = (T_V, T_E)$ which represents a compositional relations between two given types. Examples of meta-paths from query to ad defined in network schema Figure 5 include $query1 \rightarrow query4 \rightarrow ad1$, and $query2 \rightarrow keyword1 \rightarrow ad2 \rightarrow ad4$.

Firstly, to search top-K similar ad nodes for a given query node, we use the PathSim[39], a meta-path-based similarity measure. Given a symmetric meta path P , PathSim between query q and ad a is:

$$s(q, a) = \frac{2 * |p_{q \rightarrow a} : p_{q \rightarrow a} \in P|}{|p_{q \rightarrow q} : p_{q \rightarrow q} \in P| + |p_{a \rightarrow a} : p_{a \rightarrow a} \in P|} \quad (7)$$

where $p_{q \rightarrow a}$ is a path instance between q and a , $p_{q \rightarrow q}$ is that between q and q , $p_{a \rightarrow a}$ is that between a and a . $s(x, y)$ is defined in terms of two parts: their connectivity defined by the number of paths between them following P , and the balance of their visibility, where the visibility is defined as the number of path instances between themselves. We implement a distributed computation of the PathSim between queries and ads based on MapReduce framework, and the top-k results for each query can be used to build a node key-value index for online targeting service.

In the meantime, the network embedding model also can be used in node retrieval in the heterogeneous network. In our system, the metapath2vec++[13] model is used to learn the latent vector representation for multiple types of nodes. The goal of metapath2vec++ is to maximize the likelihood of preserving both the structures and semantics of a given heterogeneous network, which can be formulated as:

$$\arg \min_{\theta} \sum_{v \in V} \sum_{t \in T_V} \sum_{c_t \in N_t(v)} \log p(c_t | v; \theta) \quad (8)$$

where $N_t(v)$ denotes v 's neighborhood with the t th type of nodes and $p(c_t | v; \theta)$ is commonly defined as a softmax function,

$$p(c_t | v; \theta) = \frac{e^{X_{c_t} \cdot X_v}}{\sum_{u_t \in V_t} e^{X_{u_t} \cdot X_v}} \quad (9)$$

where X_v is the v th row of X , representing the embedding vector for node v , and V_t is the node set of type t in the network. The optimization objective is:

$$O(X) = \log \sigma(X_{c_t} \cdot X_v) + \sum_{m=1}^M E_{u_t^m \sim P_t(u_t)} [\log \sigma(-X_{u_t^m} \cdot X_v)] \quad (10)$$

The derived gradients and pseudo code are listed in[13].

Further than that, since the metapath2vec++ model can't learn good representations for the long-tail nodes and can't sufficiently utilize the rich attributes of nodes, we also adopt the GraphSAGE[21] learning framework and take a meta-path-guided aggregation manner to solve the cold-start problem.

Firstly, to utilize the information of node attributes, we extend the basic heterogeneous network to heterogeneous information network(HIN), which means that we take the text content as the attribute of different types of node. To capture the semantic information of each nodes and reduce the dimension of node embedding parameters, we use the term embedding to compose the node embedding vector with an average function $W_i = \text{average}(w_{i1}, w_{i2}, \dots, w_{in})$. The term embeddings are pre-trained with the unsupervised skip-gram model and fine-tuned with the representation learning of node embeddings.

Secondly, for the HIN, the aggregation process should follow the guide of different meta paths. Given a meta-path $t = t_1, t_2, \dots, t_n$, if $T(e1) = t_1$, so the first-order neighbors of $e1$ for the t is $N_t^1(e1) = E_t(t_1 \rightarrow t_2)$, in the same way, we can get the $N_t^k(e1)$ follow the meta-path. Following different meta-paths for queries and ads nodes, we can get the aggregation representation for query and ad from different paths in bottom-up manner, each layer can use shared weights, then we can use a fully-connected layer as the output $p_i = \text{sigmoid}(F(W_q, W_a))$. Then we can learn the embedding parameters with a log-loss function, and the positive instances are the query-ad pairs with clicks or conversions, the negative instances are the query-ad pairs without clicks. The model architecture of GraphSAGE for HIN is presented in Figure 6.

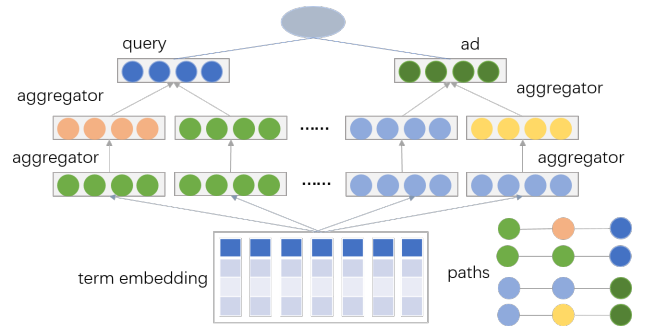


Figure 6: The architecture of GraphSAGE for HIN.

The GraphSAGE for HIN can also be used to generate node embeddings for previously unseen queries and ads nodes.

Based on the metapath2vec++ model and GraphSAGE model for HIN, we can learn the low-dimensional and latent embeddings for query nodes and ad nodes in the heterogeneous network. And the learned latent vectors for query and ad nodes can be used in building the node semantic index based on the Approximate Nearest Neighbor(ANN) retrieval platform.

Based on the structure of heterogeneous network built from historical click data, we can utilize the heterogeneous network mining method and heterogeneous network embedding model to find the potential relationship between queries and ads.

5.2 Content-based Ad Retrieval

Another straightforward retrieval model is content-based. As we can represent the query or the ad as a document, so the ad retrieval task can adopt a text-retrieval manner. In other words, we can retrieve related ads for target query just like the retrieving methods of organic search results.

At first, we should organize the queries and ads as the corresponding document as follows:

- **Query Document:** The query document consists of query text and its extensions such as clicked queries in the same session, clicked titles and snippets in search result and so on.
- **Ad Document:** The ad document consists of ad text and its extensions such as contents extracted from landing page, clicked titles, clicked keywords, clicked creations and so on.

Based on query documents and ad documents, we developed term retrieval and semantic retrieval service. For the term retrieval, an inverted index is built based on the ad documents, and given a request query, we use the traditional term retrieval metrics to return top-k related ad documents as a traditional information retrieval(IR) task with relevance ranking.

For the semantic retrieval, deep learning models for the semantic matching problem can be used to learn the semantic representations of query docs and ad docs. In our system, we adopt the CDSSM[35, 36] model architecture which incorporates a convolutional-pooling structure over word sequences to learn low-dimensional, semantic vector representations for search queries and ad documents. By using the convolution-max pooling operation, local contextual information at the word n-gram level is modeled first. Then, salient local features in a word sequence are combined to form a global feature vector. Finally, the high-level semantic information of the word sequence is extracted to form a global vector representation. Some other advanced neural matching models are also can be adopted to learn the semantic representations.

The semantic matching models are trained on click-through data by maximizing the conditional likelihood of clicked documents given a query. And the learned latent vectors for queries and ads can be used in building the semantic index based on the Approximate Nearest Neighbor(ANN) retrieval platform.

6 INTELLIGENT CREATION FRAMEWORK

The sponsored search system provides many ad formats, and the advertisers have many materials for ad creation such as image, product links, phone number, app package and so on. Since the content and layout for ad creation is diverse and rich, the search space of ad creation is huge and it's very dizzy for advertisers to choose the ad formats or combinations of materials.

To solve this problem, we design a componentized framework to design and generate ad creations intelligently. The overall architecture of the intelligent creation framework is presented in Figure 7. We treat the content and layout of each ad as a hierarchical 5-tuple(Material, Component, Template, Format).

Based on this framework, the advertiser only needs to prepare the materials such as text description about their business, promotional products, images in different sizes, app package, phone number, and so on. The system will also extract and generate some materials from advertiser's landing page or other external resources. Based

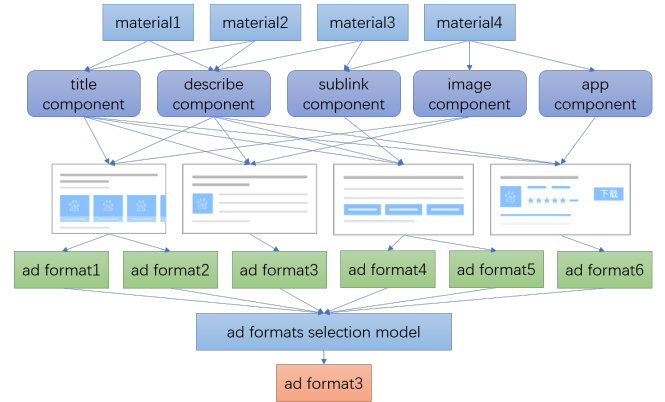


Figure 7: The intelligent creation framework.

on these materials, the various ad components are created by the system to transform the materials into functional units such as a title component, a description component, a call button or a call link component, an image component, and so on. Then based on the available candidate sizes and layout, different templates are generated by searching the combination space automatically under the constraint of UE design rule. The ad template consists of many slots, each slot can be filled with a suitable ad component. So based on the generated ad templates, plenty of ad formats can be generated by filling with ad components.

Given the materials provided by advertisers, the componentized framework can generate different combinations and layouts automatically. The task of designing and generating ad templates can be formulated as the 2D rectangle packing problem which is NP-hard, we take a heuristic search algorithm[24] and train an evaluation network to select valid templates with high CTR.

Finally, given plenty of generated ad formats, we design a selection model to predict the CTR and CVR for each ad format in different contexts. Given the target CPA, CTR and CVR, we then rank the ad formats based their expected cost-per-mille($eCPM = CPA * CVR * CTR$), and the ad format with the highest score will be selected to display.

7 EXPERIMENTS

To evaluate and demonstrate the effectiveness of our proposed system, we conduct the online A/B test and long-term grouping experiment on live traffic of Baidu. As the AiAds system will take over the main campaign optimization task such as bidding, targeting, ad creation and the system relies on that the advertiser should use conversion tracking tools provided by Baidu[5], we need the authorization from advertisers.

We collect 13350 advertisers from Baidu sponsored search platform, and set up the conversion tracking tools for them to get the conversions data. Then we choose 670 advertisers to use the AiAds system as the experimental group, and the others can be treated as the controlled group. Meanwhile, for the experimental group, we can also conduct the online A/B test to verify the performance of the AiAds system and each module. The statistics of the experimental group are presented in Table 1.

Table 1: Information of the experiment group

Item	Numbers
Advertisers	670
Daily impressions	101528334
Daily clicks	6528431
Daily conversions	378267
Nodes	119490684
Edges	359530239

For the experimental group, the advertisers adopt the target CPA bidding language and the automated bidding strategy, their keyword-level manual bidding is invalid. Their keyword match types are changed to broad match, and the intelligent targeting service provides additional ad retrieval results. They also can use the intelligent creation framework to optimize their ad creations by submitting multifarious materials, and the system generates more and more ad formats for them.

Firstly, the experimental results of the online A/B test are illustrated in table 2. We compared the AdAds system, the auto bidding strategy and the intelligent creation framework with the traditional baseline of manual optimization. The intelligent targeting service cannot be tested alone because it must be used in conjunction with the automated bidding strategy. The AiAds system can get more than 56% improvement in conversions and more than 47% improvement in the revenue of sponsored search system. Compared with the traditional advertising system, the AiAds has an overwhelming advantage due to its technical advancement.

Table 2: Experimental result of the online A/B test

Metric	AiAds	Auto-Bidding	Intelligent Creation
Revenue	43.51%	24.49%	6.83%
Click	26.88%	10.93%	6.77%
Conversion	56.91%	38.35%	5.64%
CVR	23.67%	24.72%	-1.06%
CPA	-8.54%	-10.02%	1.13%
Ad Quality	+0.89%	+1.21%	+0.03%

Secondly, as the AiAds system has deployed and launched for a long time at Baidu, we can conduct the long-term grouping experiment. We trace and compare the metrics of advertisers in experiment group with the advertisers in the controlled group for a year. The comparison between the performance during the last week in 2018 with the performance during the last week in 2017 for these two groups is presented in table 3. From the result we can find that the revenue of sponsored search and budget of the advertisers in the experimental group have significantly increased.

Both the online A/B test and the long-term grouping experiment demonstrate the advantage and the effectiveness of the AiAds system. For Baidu, there are more and more advertisers are switching to using this new advertising system. Since the advertising platform is an auction market, the performance of AdAds will gradually abate as the increase of more advertisers who choose the AiAds,

Table 3: Experimental result of the long-term grouping experiment

Metric	Experiment Group	Controlled Group
Revenue	27.22%	10.85%
Budget	30.20%	8.83%
Conversion	22.67%	5.20%
CVR	18.24%	-2.91%
CPA	3.71%	5.37%

but it still has a distinct advantage compared with the baseline due to the lifting in CVR and finding new query-ad matching space.

8 CONCLUSIONS

In this paper, we introduce the automated and intelligent advertising system deployed at Baidu. As the manual optimizations of keyword-level bidding, keyword selection and ad creation are highly inefficient and time-consuming, We design and develop automated bidding strategy, intelligent targeting model and intelligent creation framework to combine 3 optimization technologies to take the labor and guesswork out of targeting, bidding, and ad creation.

The automated bidding strategy can replace the traditional keyword-level and manual bidding strategy, and given a simple and straight bidding language for advertisers. The intelligent targeting model can break the limits of keyword targeting and achieve straightforward ad retrieval from queries to ads, which will bring more commercial traffic to advertisers and sponsored search platform. The intelligent creation framework provides a simple way to have a global optimization for the ad creation and layout.

By replacing manual tasks with automated and intelligent models, we think that the AiAds system opens a new door, and is bringing a revolution for sponsored search. For the whole digital advertising market, the intelligent technologies will reform or rebuild most parts of the market and will bring new opportunities and challenges. The underlying principle is letting the machines do what they do best. We hope that the lessons learned in building the AiAds system can also be helpful to the entire advertising market.

Due to the limits of pages, we give the guideline of this system but can't fully introduce the details for each model. Furthermore, the technologies behind the AiAds are still improving, we will separately introduce each part of AiAds in a more detailed way in the future. As the future works we will also optimize the ad retrieval model by utilizing more data sources and advanced models. Meanwhile, there are some open problems in designing the reasonable mechanism for ROI-constrained bidders.

ACKNOWLEDGMENTS

We would sincerely like to thank the advertisers who participated in the early experiments of AiAds. We also thank the anonymous reviewers for their valuable comments and helpful suggestions.

REFERENCES

- [1] Deepak Agarwal, Souvik Ghosh, Kai Wei, and Siyu You. 2014. Budget pacing for targeted online advertisements at linkedin. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1613–1619.

- [2] Tasos Anastasakos, Dustin Hillard, Sanjay Kshetramade, and Hema Raghavan. 2009. A collaborative filtering approach to ad recommendation using the query-ad click graph. In *Proceedings of the 18th ACM conference on Information and knowledge management*. ACM, 1927–1930.
- [3] Yoram Bachrach, Sofia Ceppi, Ian A Kash, Peter Key, and David Kurokawa. 2014. Optimising trade-offs among stakeholders in ad auctions. In *Proceedings of the fifteenth ACM conference on Economics and computation*. ACM, 75–92.
- [4] Baidu. 2019. Baidu ad formats. Retrieved 2019 from http://yingxiao.baidu.com/new/home/product/product/id/50?ly=product_union_author_lists
- [5] Baidu. 2019. Baidu conversion tracking. Retrieved 2019 from <http://ocpc.baidu.com/developer/d/guide>
- [6] Bing. 2019. Bing automated bid strategies. Retrieved 2019 from <https://help.bingads.microsoft.com/apex/index/3/en/56786>
- [7] Andrei Broder, Peter Ciccolo, Evgeniy Gabrilovich, Vanja Josifovski, Donald Metzler, Lance Riedel, and Jeffrey Yuan. 2009. Online expansion of rare queries for sponsored search. In *Proceedings of the 18th international conference on World wide web*. ACM, 511–520.
- [8] Andrei Broder, Evgeniy Gabrilovich, Vanja Josifovski, George Mavromatis, and Alex Smola. 2011. Bid generation for advanced match in sponsored search. In *Proceedings of the fourth ACM international conference on Web search and data mining*. ACM, 515–524.
- [9] Andrei Z Broder, Peter Ciccolo, Marcus Fontoura, Evgeniy Gabrilovich, Vanja Josifovski, and Lance Riedel. 2008. Search advertising using web relevance feedback. In *Proceedings of the 17th ACM conference on Information and knowledge management*. ACM, 1013–1022.
- [10] Matthew Cary, Aparna Das, Ben Edelman, Ioannis Giotis, Kurtis Heimerl, Anna R Karlin, Claire Mathieu, and Michael Schwarz. 2007. Greedy bidding strategies for keyword auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*. ACM, 262–271.
- [11] Ruggiero Cavallo, Prabhakar Krishnamurthy, Maxim Sviridenko, and Christopher A Wilkens. 2017. Sponsored search auctions with rich ads. In *Proceedings of the 26th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 43–51.
- [12] Yejin Choi, Marcus Fontoura, Evgeniy Gabrilovich, Vanja Josifovski, Mauricio Mediano, and Bo Pang. 2010. Using landing pages for sponsored search ad selection. In *Proceedings of the 19th international conference on World wide web*. ACM, 251–260.
- [13] Yuxiao Dong, Nitesh V Chawla, and Ananthram Swami. 2017. metapath2vec: Scalable representation learning for heterogeneous networks. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 135–144.
- [14] Eyal Even Dar, Vahab S Mirrokni, S Muthukrishnan, Yishay Mansour, and Uri Nadav. 2009. Bid optimization for broad match ad auctions. In *Proceedings of the 18th international conference on World wide web*. ACM, 231–240.
- [15] Jon Feldman and S Muthukrishnan. 2008. Algorithmic methods for sponsored search advertising. In *Performance Modeling and Engineering*. Springer, 91–122.
- [16] Negin Golrezaei, Ilan Lobel, and Renato Paes Leme. 2018. Auction Design for ROI-Constrained Buyers. (2018).
- [17] Google. 2019. Google ads extensions. Retrieved 2019 from <https://support.google.com/google-ads/answer/2375499>
- [18] Google. 2019. Google automated bid strategies. Retrieved 2019 from <https://support.google.com/google-ads/answer/2979071?hl=en>
- [19] Mihajlo Grbovic, Nemanja Djuric, Vladan Radosavljevic, Fabrizio Silvestri, Ricardo Baeza-Yates, Andrew Feng, Erik Ordentlich, Lee Yang, and Gavin Owens. 2016. Scalable semantic matching of queries to ads in sponsored search advertising. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. ACM, 375–384.
- [20] Sonal Gupta, Mikhail Bilenko, and Matthew Richardson. 2009. Catching the drift: learning broad matches from clickthrough data. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1165–1174.
- [21] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*. 1024–1034.
- [22] Jason Hartline, Nicole Immorlica, Mohammad Reza Khani, Brendan Lucier, and Rad Niazadeh. 2018. Fast Core Pricing for Rich Advertising Auctions. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. ACM, 111–112.
- [23] Benamin Heymann. 2018. ROI constrained Auctions. *arXiv preprint arXiv:1809.08837* (2018).
- [24] Jukka Jylänki. 2010. A thousand ways to pack the bin—a practical approach to two-dimensional rectangle bin packing. *retrived from http://clb.demon.fi/files/RectangleBinPack.pdf* (2010).
- [25] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. Lightgbm: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems*. 3146–3154.
- [26] Kuang-Chih Lee, Ali Jalali, and Ali Dasdan. 2013. Real time bid optimization with smooth budget delivery in online advertising. In *Proceedings of the Seventh International Workshop on Data Mining for Online Advertising*. ACM, 1.
- [27] Vahab Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. 2016. Dynamic auctions with bank accounts. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. AAAI Press, 387–393.
- [28] Vahab Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. 2017. Non-clairvoyant dynamic mechanism design. (2017).
- [29] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
- [30] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529.
- [31] Sandeep Pandey, Kunal Punera, Marcus Fontoura, and Vanja Josifovski. 2010. Estimating advertisibility of tail queries for sponsored search. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*. ACM, 563–570.
- [32] Tao Qin, Wei Chen, and Tie-Yan Liu. 2015. Sponsored search auctions: Recent advances and future directions. *ACM Transactions on Intelligent Systems and Technology (TIST)* 5, 4 (2015), 60.
- [33] Sujith Ravi, Andrei Broder, Evgeniy Gabrilovich, Vanja Josifovski, Sandeep Pandey, and Bo Pang. 2010. Automatic generation of bid phrases for online advertising. In *Proceedings of the third ACM international conference on Web search and data mining*. ACM, 341–350.
- [34] Benjamin Rey and Ashvin Kannan. 2010. Conversion rate based bid adjustment for sponsored search. In *Proceedings of the 19th international conference on World wide web*. ACM, 1173–1174.
- [35] Yelong Shen, Xiaodong He, Jianfeng Gao, Li Deng, and Grégoire Mesnil. 2014. A latent semantic model with convolutional-pooling structure for information retrieval. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*. ACM, 101–110.
- [36] Yelong Shen, Xiaodong He, Jianfeng Gao, Li Deng, and Grégoire Mesnil. 2014. Learning semantic representations using convolutional neural networks for web search. In *Proceedings of the 23rd International Conference on World Wide Web*. ACM, 373–374.
- [37] Hyun-Je Song, A Kim, Seong-Bae Park, et al. 2017. Translation of Natural Language Query Into Keyword Query Using a RNN Encoder-Decoder. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 965–968.
- [38] Statista. 2019. Search Advertising - worldwide | Statista Market Forecast. Retrieved 2019 from <https://www.statista.com/outlook/219/100/search-advertising/worldwide>
- [39] Yizhou Sun, Jiawei Han, Xifeng Yan, Philip S Yu, and Tianyi Wu. 2011. PathsIm: Meta path-based top-k similarity search in heterogeneous information networks. *Proceedings of the VLDB Endowment* 4, 11 (2011), 992–1003.
- [40] Christopher A Wilkens, Ruggiero Cavallo, and Rad Niazadeh. 2016. Mechanism design for value maximizers. *arXiv preprint arXiv:1607.04362* (2016).
- [41] Jian Xu, Kuang-chih Lee, Wentong Li, Hang Qi, and Quan Lu. 2015. Smart pacing for effective online ad campaign optimization. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2217–2226.
- [42] Su Yan, Wei Lin, Tianshu Wu, Daorui Xiao, Xu Zheng, Bo Wu, and Kaipeng Liu. 2017. Beyond Keywords and Relevance: A Personalized Ad Retrieval Framework in E-Commerce Sponsored Search. *arXiv preprint arXiv:1712.10110* (2017).
- [43] Bianca Zadrozny and Charles Elkan. 2002. Transforming classifier scores into accurate multiclass probability estimates. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 694–699.
- [44] Weinan Zhang, Dingquan Wang, Gui-Rong Xue, and Hongyuan Zha. 2012. Advertising keywords recommendation for short-text web pages using Wikipedia. *ACM Transactions on Intelligent Systems and Technology (TIST)* 3, 2 (2012), 36.
- [45] Weinan Zhang, Ying Zhang, Bin Gao, Yong Yu, Xiaojie Yuan, and Tie-Yan Liu. 2012. Joint optimization of bid and budget allocation in sponsored search. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1177–1185.
- [46] Wei Vivian Zhang, Xiaofei He, Benjamin Rey, and Rosie Jones. 2007. Query rewriting using active learning for sponsored search. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 853–854.
- [47] Ying Zhang, Weinan Zhang, Bin Gao, Xiaojie Yuan, and Tie-Yan Liu. 2014. Bid keyword suggestion in sponsored search based on competitiveness and relevance. *Information Processing & Management* 50, 4 (2014), 508–523.
- [48] Jun Zhao, Guang Qiu, Ziyu Guan, Wei Zhao, and Xiaofei He. 2018. Deep Reinforcement Learning for Sponsored Search Real-time Bidding. *arXiv preprint arXiv:1803.00259* (2018).
- [49] Han Zhu, Junqi Jin, Chang Tan, Fei Pan, Yifan Zeng, Han Li, and Kun Gai. 2017. Optimized cost per click in taobao display advertising. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2191–2200.