



# Bid keyword suggestion in sponsored search based on competitiveness and relevance <sup>☆</sup>

Ying Zhang <sup>a,\*</sup>, Weinan Zhang <sup>b</sup>, Bin Gao <sup>c</sup>, Xiaojie Yuan <sup>a</sup>, Tie-Yan Liu <sup>c</sup>

<sup>a</sup> College of Software, Nankai University, Tianjin 300071, PR China

<sup>b</sup> Department of Computer Science, University College London, London WC1E 6BT, UK

<sup>c</sup> Microsoft Research, No. 5 Danling Street, Beijing 100080, PR China

## ARTICLE INFO

### Article history:

Received 21 September 2012

Received in revised form 3 January 2014

Accepted 19 February 2014

### Keywords:

Bid keyword suggestion

Bid optimization

Sponsored search

## ABSTRACT

In sponsored search, many advertisers have not achieved their expected performances while the search engine also has a large room to improve their revenue. Specifically, due to the improper keyword bidding, many advertisers cannot survive the competitive ad auctions to get their desired ad impressions; meanwhile, a significant portion of search queries have no ads displayed in their search result pages, even if many of them have commercial values. We propose recommending a group of relevant yet less-competitive keywords to an advertiser. Hence, the advertiser can get the chance to win some (originally empty) ad slots and accumulate a number of impressions. At the same time, the revenue of the search engine can also be boosted since many empty ad shots are filled. Mathematically, we model the problem as a mixed integer programming problem, which maximizes the advertiser revenue and the relevance of the recommended keywords, while minimizing the keyword competitiveness, subject to the bid and budget constraints. By solving the problem, we can offer an optimal group of keywords and their optimal bid prices to an advertiser. Simulation results have shown the proposed method is highly effective in increasing ad impressions, expected clicks, advertiser revenue, and search engine revenue.

© 2014 Elsevier Ltd. All rights reserved.

## 1. Introduction

Sponsored search is one of the major types of online advertising and is also the key source of revenue for search engine companies. In sponsored search, a set of ads are displayed along with organic search results when answering a query. Although displayed simultaneously and in similar forms, sponsored search results are actually generated by a quite different mechanism from that of organic search. While the organic search results are produced according to the relevance of each web page to the query, the sponsored search results are generated according to an auction process (Aggarwal, Goel, & Motwani, 2006; Varian, 2007).

Before the auction happens, each advertiser is asked to participate in a bid process, in which he/she creates a group of ad copies and bids on some keywords for the ad group with their match types. The match type might be exact match or broad match (a.k.a. advanced match). When a query is issued, the search engine will first retrieve some candidate ads whose

<sup>☆</sup> This work was done when the first and the second authors were interns at Microsoft Research Asia.

\* Corresponding author. Address: College of Software, Nankai University, No. 94, Weijin Road, Nankai District, Tianjin 300071, PR China. Tel.: +86 22 23508468; fax: +86 22 23509175.

E-mail address: [zhangying@dbis.nankai.edu.cn](mailto:zhangying@dbis.nankai.edu.cn) (Y. Zhang).

bid keywords match the query. Then the search engine will run an auction on these candidate ads by considering both the ad quality and the bid prices of them (Feng, Bhargava, & Pennock, 2003). Those ads with the highest quality scores and bid prices will win the auction and be displayed on the search result page. If such an ad is clicked by a user, its advertiser will be charged by the search engine. Mainstream search engines adopt the generalized second price (GSP) (Edelman, Ostrovsky, & Schwarz, 2005) auction mechanism, which means that the advertiser's cost of a click depends on the bid price and the relevance score of the next ad in the ranking list of the auction.

As can be seen above, an advertiser should carefully consider which keywords to bid and what price to set for each of these keywords, in order to win the auction. However, data statistics show that not many advertisers are good at dealing with this (see Section 3.1.2). For example, too many advertisers bid on a small number of popular keywords, and thus as a result those advertisers with low bid prices will fail due to the hot competition in the auction process and do not have as many ad impressions as expected. This is bad for these advertisers since they have missed their campaign goals. This is also bad for search engines since the real contributions of these advertisers to the revenue of search engine will be much lower than their budgets. On the other hand, however, data statistics also show that a significant portion of search queries have no ads (or no enough ads) displayed in their search result pages, even if many of them have commercial values (see Section 3.1.1). In other words, many keywords that may potentially lead to ad clicks have been ignored by the advertisers. Again, this is bad for search engines since many potentially valuable ad slots have been wasted. This is also bad for advertisers since they have missed many advertising opportunities.

It is clear that if we can effectively solve the aforementioned problem, we can improve the satisfaction of advertisers and increase the revenue of search engines simultaneously. A straightforward way to solve this problem is to suggest appropriate keywords to the advertisers to bid. Actually, the ad platforms in many search engines already provide this type of services, such as Keyword Group Detection of Microsoft AdCenter,<sup>1</sup> Overture (Yahoo!) Keyword Selector Tool,<sup>2</sup> and the Keyword Tool of Google AdWords.<sup>3</sup> In the research community, there are also many papers on bid keyword optimization (Abhishek & Hosanagar, 2007; Bartz, Murthi, & Sebastian, 2006; Kitts & Leblanc, 2004; Chen, Xue, & Yu, 2008; Fuxman, Tsaparas, Achan, & Agrawal, 2008; Kiritchenko & Jiline, 2008). However, most of these existing works are based on the semantic similarity between keywords and/or the co-occurrence of bid keywords across advertisers. As a consequence, it is easy to understand that these methodologies will not effectively solve the aforementioned problem, and sometimes make the situation even worse: they will tend to suggest popular keywords to the advertisers and the competition on these popular keywords will become tougher and tougher.

In this paper, we propose a novel keyword suggestion technology, which can alleviate the aforementioned problem. In particular, instead of suggesting popular keywords, we recommend a group of relevant yet less-competitive keywords to the advertisers, by optimizing the expected advertiser revenue. Here, less-competitive keywords correspond to the keywords that have not been intensively bid by advertisers. This idea is inspired by the long tail theory (Anderson, 2008). That is, the aggregated popularity of a large numbers of less-competitive (sometimes tail) items can make a large fraction of the total popularity. As pointed out by many previous works, for search engines, the query volume just follows a long tail distribution. To validate this phenomenon in the data analysis experiment, we used a search log dataset that records the submitted queries from a commercial search engine,<sup>4</sup> which was collected in half a month (from 01-April-2011 to 15-April-2011). It contains about 3.3 billion user queries in volume, in which there are 623 million unique queries. The distribution of the query volume is shown in a log–log scale in Fig. 1. From the figure, we can see that the query volume follows a long tail distribution obviously. This suggests that the aggregated volume of less-competitive queries may take a large part of the total query volume. Therefore, if an advertiser bids on a package of less-competitive keywords, he/she may win a large number of ad auctions and accumulate high ad impressions (and potentially high return on investment, i.e., ROI). Meanwhile, the search engine can also get extra revenue since many empty ad slots related to the suggested keywords can be utilized.

Mathematically, we formulate the keyword recommendation problem as a mixed integer optimization problem. Given a target ad group, we first collect a set of candidate keywords whose relevance score to the ad group can exceed the floor relevance score in sponsored search. Then we perform selection among these candidate keywords and try to give an optimal bid price for each selected keyword, by maximizing the revenue of ad group with the budget control constraint. In this process, more relevant and less competitive candidate keywords will have higher possibility to be selected. This constrained optimization problem can be solved by binary integer programming and sequential quadratic programming (SQP) in an alternate manner. Simulation results on the sponsored search log obtained from a commercial search engine,<sup>5</sup> show that our proposed technology can effectively increase the ad impressions of advertisers with a low average cost per click, and it can effectively help advertisers obtain increased revenue. At the same time, our method can significantly reduce the empty ad slots and enlarge the revenue of search engine.

To sum up, the contributions of our work are listed as below. (i) We perform a comprehensive study on keyword bidding in sponsored search, and point out a problem with the current sponsored search services that may make both advertisers and search engines unsatisfied. (ii) We propose a bid keyword suggestion method based on competitiveness and relevance, and

<sup>1</sup> <http://adlab.msn.com/Keyword-Group-Detection>.

<sup>2</sup> <http://inventory.overture.com/d/searchinventory/suggestion>.

<sup>3</sup> <https://adwords.google.com/select/KeywordToolExternal>.

<sup>4</sup> <http://www.bing.com>.

<sup>5</sup> <http://www.bing.com>.

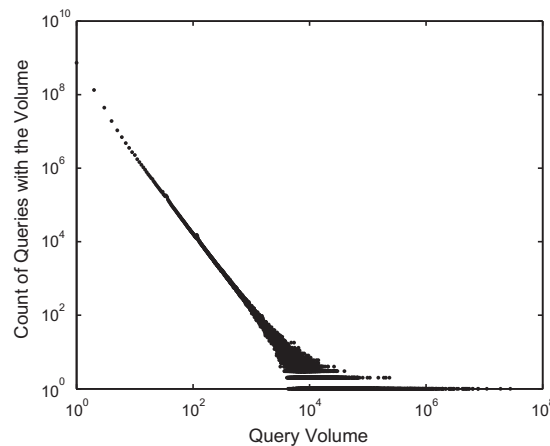


Fig. 1. The long tail distribution of the query volume.

verify the effectiveness of the method. (iii) We model bid keyword suggestion as a mixed integer optimization problem, and solve it using binary integer programming and sequential quadratic programming alternatively.

The rest of the paper is organized as follows. In Section 2, we review the state of the art on bid keyword suggestion and make discussions regarding the difference between our work and previous work. In Section 3, we report our statistical study on the sponsored search log from a commercial search engine to show the possibility of using a package of relevant yet less-competitive keywords to improve ad impressions and search engine revenue. The problem definition and the optimization algorithm are described in Section 4. The experimental results are presented and discussed in Section 5. Conclusions and future work are given in the last section.

## 2. Related work

As mentioned in the introduction, bid keywords play a critical role in sponsored search. In the literature, a lot of research has been done regarding topics related to bid keywords. For example, Broder et al. (Broder et al., 2007; Broder et al., 2008; Broder et al., 2009), Even-Dar et al. (Even-Dar, Mirrokni, Muthukrishnan, Mansour, & Nadav, 2009) and Zhang et al. (Zhang, He, Rey, & Jones, 2007) studied how to improve the broad match of bid keywords for a given query. Gupta et al. (Gupta, Bilenko, & Richardson, 2009) proposed an adaptive algorithm which could utilize arbitrary similarity functions and catch the dynamics in the broad match. Ravi et al. (Ravi et al., 2010) discussed how to generate bid keywords for some given landing pages of the advertisers. Pandey et al. (Pandey, Punera, Fontoura, & Josifovski, 2010) studied the advertisability of tail queries in sponsored search system. Zhang et al. (Zhang et al., 2012) studied the hierarchical structure of sponsored search advertisers and proposed a joint optimization of keyword bid price and campaign budget allocation under a multi-campaign sponsored search account.

In addition, several dedicated works have been proposed for bid keyword recommendation. Bid keyword recommendation aims at finding a group of keywords for an advertiser to bid based on his/her original bid keywords and/or ad copies. Representative work on bid keyword recommendation include (Abhishek & Hosanagar, 2007; Chen et al., 2008; Fuxman et al., 2008; Joshi & Motwani, 2006; Bartz et al., 2006; Kitts & Leblanc, 2004; Kiritchenko & Jilinc, 2008; Zainal-Abidin & Wang, 2010; Broder, Gabrilovich, Josifovski, Mavromatis, & Smola, 2011; Berg, Greenwald, Naroditskiy, & Sodomka, 2010; Sodomka, Lahaie, & Hillard, 2011). Most of these works have considered keyword relevance as a key factor in their algorithms. For example, Chen et al. (Chen et al., 2008) built a hierarchy of concepts based on a web directory. Given a keyword, they first matched it to some relevant concepts, and then considered these concepts and their parent concepts in the hierarchy for keyword recommendation. Fuxman et al. (Fuxman et al., 2008) proposed using the click graph extracted from search logs to compute the keyword similarity for recommendation. Joshi et al. (Joshi & Motwani, 2006) proposed a graph model named *TermsNet*, which regards terms as the vertices and the similarities between terms as the weights of the edges in a graph. The similarity score was calculated based on the snippets of the top search results when submitting the terms as queries to a search engine. The keyword recommendation was based on similarity propagation on the graph. Abhishek et al. (Abhishek & Hosanagar, 2007) proposed a system called *Wordy* for bid keyword recommendation, in which a term graph similar to *TermsNet* was used. The difference lies in that in *Wordy* the whole retrieved documents were used to compute the similarity score rather than only using the snippets. In addition, the authors adopted a broad search algorithm to find the less frequent terms within a certain distance to the original bid keywords in the graph as the recommended keywords. To sum up, the above works are very similar in their nature, i.e., they recommend bid keywords based on relevance, and the relevance scores are calculated from semantic relationship structures such as query similarity graph or concept hierarchy. The problem with these approaches is that semantic similarity is not enough to improve advertising effectiveness. The end goal of advertisers is to increase their expected revenue or ROI, while recommending similar but highly-competitive keywords would do little help.

In the industry of online advertising, there are also some tools for bid keyword recommendation, such as Keyword Group Detection provided by Microsoft AdCenter, Overture Keyword Selector Tool provided by Yahoo!, and the AdWords Keyword Tool provided by Google. Usually, both relevance and popularity of the keywords are considered in these tools. That is, these tools tend to recommend those keywords that are relevant to the original bid keywords, and have been bid by many other advertisers. However, as discussed in the introduction, such a keyword recommendation mechanism may increase the competition of the auctions, and as a result many advertisers without high enough budgets and bids will lose the opportunities to show their ads.

To solve the problems with the previous work, in this paper, we investigate the problem of keyword recommendation by looking at the factors that are more related to advertising effectiveness, and by avoiding the hot competition among advertisers. Our goal is not only to help advertisers optimize their campaigns or ad groups but also to help search engines improve their revenues. This is significantly different from the attempts made by the aforementioned previous work and by search engine optimization (SEO) companies.<sup>6</sup>

### 3. Data analysis on sponsored search

Our proposal of bid keyword suggestion is motivated by an intensive study on the sponsored search log of a commercial search engine. We used three kinds of data in our study: the search log that records the submitted queries and the corresponding ad impressions and clicks, the auction log that records the detailed auction processes, and the advertiser database that includes the bid keywords, bid prices, and the budget for each ad group. The data were collected in half a month (from 01-April-2011 to 15-April-2011), which contains about 623 million unique queries and about 31 million of active ad groups. The findings of the study are reported in this section.

#### 3.1. Existing problems in sponsored search

The key mechanism in sponsored search is the auction for the ad slots in the search result page of each query. When the match type is *exact match*, the ad groups will be involved in the auction if they exactly bid the query as their keyword. When the match type is *broad match* or *advanced match*, the original query will be expanded to several related keywords, and then all the ad groups that bid any of these keywords will participate in the auction. As can be seen above, bid keywords play a critical role in the auction process. Improper keyword bidding will affect the effectiveness of the auction and thus influence the performance of both the advertisers and the search engine. Regarding keyword bidding, we found the following problems through our data analysis.

##### 3.1.1. Wasted ad slots

According to our study, 34.5% of the unique queries have no ad impression on their search result pages, and these queries occupy 41.0% search traffic. Another 52.4% unique queries (corresponding to 46.6% search traffic) have ad impressions but the number of ad groups winning in the corresponding auction is smaller than eight.<sup>7</sup> The rest 13.1% unique queries (corresponding to only 12.4% search traffic) attract eight or more ad groups in their auctions. In other words, 87.6% of the search traffic has not been fully utilized by the advertisers, and nearly half of the traffic has not been used at all. Therefore, a lot of ad slots have actually been wasted, even if some of the corresponding queries have commercial values.<sup>8</sup>

Our explanations to the above observations are as follows. On one aspect, advertisers usually prefer popular keywords and ignore the potential of rare yet relevant keywords. On another aspect, the existing keyword recommendation tools provided by search engines also tend to suggest popular keywords to advertisers, because most of these tools are based on the semantic similarity or co-occurrences of keywords across ad groups. As a result, many tail keywords are not bid and their corresponding ad slots are left empty.

##### 3.1.2. Loss in highly-competitive auctions

According to our study (see Table 1), 55.3% ad groups have no ad impression, 92.3% ad groups have no ad click. The owners of all these related ad groups should be regarded as “unsuccessful” in sponsored search. The reason is that, their ads are not shown in the search result pages or not clicked by the users, so the information in the landing pages cannot be seen by the users, and there will be no conversion behaviors (registration, add something to the shopping cart, purchase something, etc.) from the users. The major problems with these ad groups are listed as below, i.e., low relevance<sup>9</sup> and low bid price.

<sup>6</sup> <http://www.keywordperformance.com/>.

<sup>7</sup> The maximum number of ad slots in each search result page is eight for the commercial search engine.

<sup>8</sup> According to a commercial query classifier provided by the search engine (Dai et al., 2006), 68.5% of the 87.6% search traffic are commercial queries.

<sup>9</sup> When we talk about relevance in the paper, we are actually referring the relevance score calculated by the commercial search engine. We take query-keyword relevance score as example and other relevance scores are calculated in the similar method. First, we extract some features between query and keyword, such as query-keyword similarities (cosine similarity, Jaccard similarity, etc.), semantic similarities (calculated by knowledge base like Freebase (<http://www.freebase.com/>)), and taxonomy information (obtained from a hierarchical text classifier). Then we ask human judges to label the relevance degrees (very relevant, relevant, neutral, irrelevant, and very irrelevant) of some selected query-keyword pairs. After that, we use a learning to rank method (Ranking SVM (Joachims & Thorsten, 2002)) to learn a relevance scoring function based on the features and the labeled data. Thus, for any given query-keyword pair, we can calculate their relevance score.

**Table 1**  
The statistics of the ad groups.

Ad groups	Count	Percentage
Total	31,047,416	100
Participated auctions	17,815,780	57.4
Had ad impressions	13,883,463	44.7
Had ad clicks	2,399,234	7.7

- **Low relevance.** Among the 55.3% ad groups having no ad impression, 42.6% ad groups were never involved in any auction. This is because the keywords that they bid are irrelevant to any query issued during the period of our study. As a result, either they were not triggered or they were filtered out from the auction process due to their low relevance.
- **Low bid price.** Among the 55.3% ad groups, 12.7% ad groups participated in some auctions but did not get any impression, although the keywords they bid are relevant. This is because their bid prices are too low as compared to those of other advertisers, especially when the competition is hot. For those ad groups with impressions but no clicks (34.0%), the major reason might also be the low bid price. That is, their bid prices are not high enough to make them ranked on the top and therefore their corresponding click-through rate (CTR) are low.

It is clear that these unsuccessful advertisers will not be satisfied with the sponsored search system. They might choose to switch to another publisher if the problem cannot be solved after a certain period of time.


### 3.2. Potential of advertisability

Intuitively, we can improve the aforementioned situations by bidding on rare but relevant keywords. In this way, we can fill some empty ad slots and can also ensure certain impressions of those ad groups with relatively low bid prices. However, the question is whether the improvement can be significant. In other words, is there enough advertisability in these keywords given that they are rare? In this subsection, we report our findings regarding this question.

#### 3.2.1. Distribution of search queries

To understand the potential advertisability of rare queries, we sort the queries according to their decreasing search volumes and evenly divide them into 50 buckets, i.e., each bucket contains 2% unique queries. Then we count the total numbers of ad impressions and ad clicks associated with the queries in each bucket (See Fig. 2).

According to Fig. 2, up to 40.9% ad impressions and 27.8% ad clicks are associated with the 98% least frequent queries. These 98% unique queries are rare queries. According to our statistics, only 16% of these rare queries have triggered auctions and only 14% of them eventually have ad impressions. Furthermore, 98.9% of these auctions were triggered by broad match or advanced match, indicating that very few advertisers actually bid on these queries as keywords. If one bids on some of these keywords that are relevant to his/her ad group, it is very likely that he/she will obtain impressions since there is little competition in the corresponding auctions. Furthermore, bidding on a group of such keywords will probably accumulate a number of ad impressions for him/her as a result.



**Fig. 2.** The percentage of query volume, ad impression, ad click and the value of CTR for each bucket of queries.

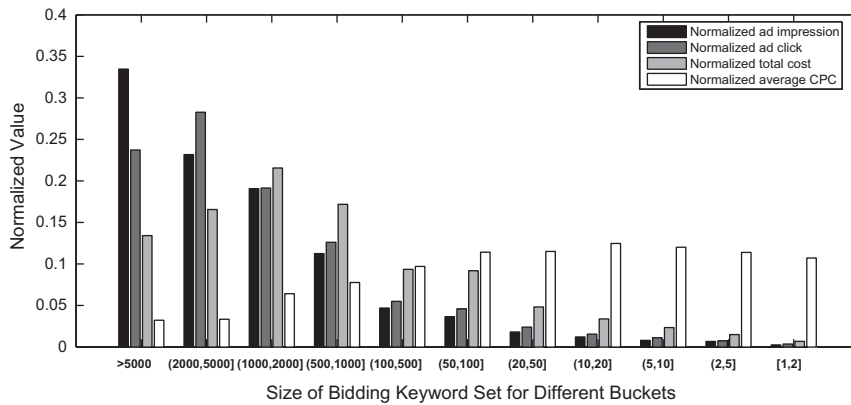


Fig. 3. The percentage of normalized ad impression, ad click, total cost, and average CPC of the ad groups among buckets.

In addition, as shown in the figure, the average click-through rate (CTR) is stable with respect to query popularity. That is, even though the queries are rare, as long as we have a large number of total impressions (by collectively considering a large number of such queries), we can still expect a large click number. Therefore, the potential return for these rare queries is considerable.

### 3.2.2. Cost for large number of bid keywords

We investigate the effects of bidding a group of keywords. We sort the ad groups according to the decreasing number of their bid keywords and divide them into buckets. Then we look at ad impressions, ad clicks, total cost,<sup>10</sup> and average cost per click (CPC) with respect to the buckets. The distributions of these metrics<sup>11</sup> are shown in Fig. 3. We have the following observations from it. (i) The number of ad impressions/clicks increases as the number of bid keywords in an ad group increases. This is quite intuitive and consistent with our discussions in the previous subsection. (ii) The average CPC decreases as the number of bid keywords in an ad group increases. This accords with the investment principal known as *averaging down* which refers to the average price an investor paid per share reduces by buying additional shares after the price has dropped. Each time the price goes down he buys more and his average price per share keeps going down. Therefore, it is good for advertisers that bidding a group of keywords since they can get a lot more clicks by paying less average cost.

In summary, the above study reveals that it is an effective strategy to bid a group of relevant yet less-competitive keywords with low prices, since these cheap keywords may collectively attract accumulated ad impressions and clicks for the advertiser.

## 4. Bid keyword suggestion

In this section, we describe our proposal to solve the problem as mentioned in Section 3.1. The key idea is to recommend a group of relevant but less-competitive keywords with optimal low prices for the advertisers to bid. To better illustrate this idea, we first give some necessary notations and definitions including winning score, winning price, impression probability based on bid price, and impression confidence based on competitiveness. After that we formulate bid keyword suggestion as a mixed integer optimization problem and discuss how to solve it efficiently.

### 4.1. Notations and definitions

The problem setting of bid keyword suggestion is as follows. The input is a target ad group  $\omega$  with its ad copies, original bid keywords,<sup>12</sup> and bid prices. In addition, each ad group is associated with a budget  $g_\omega$  by the advertiser.

#### 4.1.1. Winning score and winning price

In sponsored search, an issued query might trigger an auction. If one of the associated keywords of an ad matches the query by the corresponding match functions according to the match types of the keywords, the ad (together with the matched bid keyword) will be involved in the auction. The search engines may use the product of the bid price and the relevance score as the rank score in their auctions (Feng et al., 2003). The ads in the auction are ranked by their rank scores. The winning score of an auction is defined as the following.

<sup>10</sup> Total cost is the total amount of money that search engine charges the advertiser on an ad group by its clicks.

<sup>11</sup> To protect the business secrete of the commercial search engine, here we only report the normalized values.

<sup>12</sup> A keyword may have different match types and different bid prices accordingly. For simplicity, we regard them as different keywords.



**Definition 1** (Winning Score). Given an auction  $\theta$ , its winning score  $\mu_\theta$  is the lowest rank score to display an ad on the corresponding search result page.

In an auction, if the number of participating ad groups is bigger than the number of ad slots, some ad groups will lose the auction. In this case, the winning score is the rank score of the ad that wins the last ad slot. If the number of participating ad groups is smaller than or equal to the number of ad slots, the winning score will be  $s_\epsilon$ , which is the floor rank score for ads to be showed in the result pages in sponsored search system.

Suppose the relevance score of an ad group  $\omega$  in an auction  $\theta$  is  $r_{\omega,\theta}$ , which can be calculated based on the features like query-ad similarity, semantic similarity, taxonomy, and user query time (Graepel, Candela, Borchert, & Herbrich, 2010; Hillard, Schroedl, Manavoglu, Raghavan, & Leggetter, 2010; Radlinski et al., 2008).<sup>13</sup> The relevance score  $r_{\omega,\theta}$  of an ad group can be different in different auctions, due to some contextual information on time, location, and user for the triggering query (Graepel et al., 2010). Such a relevance score indicates the probability that an ad will be clicked after it is shown in the search result page. Usually, there is a floor relevance score  $r_\epsilon$  in the sponsored search system. The ad group with  $r_{\omega,\theta} < r_\epsilon$  is not eligible to participate in the auction. Based on the relevance score, we introduce the following definition of winning price of an ad group in an auction.

**Definition 2** (Winning Price). Given an ad group  $\omega$  in an auction  $\theta$  with its relevance score  $r_{\omega,\theta}$ , the winning price  $\beta_{\omega,\theta}$  is calculated as  $\frac{\mu_\theta}{r_{\omega,\theta}}$ , which is the least bid price that can make  $\omega$  be shown on the corresponding search result page.

#### 4.1.2. Impression probability based on bid price

Intuitively, the more the rank score of an ad group exceeds the winning score of the auction, the higher the probability for the ad group to have an impression in the search result page is. Considering the relationship between the bid price and the rank score, the bid price will directly affect the impression probability of an ad group. We define the impression probability based on bid price as below.

**Definition 3** (IMPRESSION PROBABILITY BASED ON BID PRICE). Given the winning score  $\mu_\theta$  of an auction  $\theta$  and the relevance score  $r_{\omega,\theta}$  of an ad group  $\omega$  in auction  $\theta$ , the impression probability of ad group  $\omega$  based on bid price  $b_\omega$ , denoted by  $p_{\omega,\theta}$ , is defined as

$$p_{\omega,\theta} = \Phi \left[ \frac{b_\omega r_{\omega,\theta} - \mu_\theta + \varepsilon}{\lambda} \right].$$

Here  $\Phi[x] = \int_{-\infty}^x \mathcal{N}(z; 0, 1) dz$  is the standard cumulative Gaussian density.<sup>14</sup> It maps the rank score  $b_\omega r_{\omega,\theta}$  and the winning score  $\mu_\theta$  from interval  $(-\infty, \infty)$  to probabilities in interval  $(0, 1)$ . Parameter  $\varepsilon$  adjusts the base impression probability and parameter  $\lambda (\lambda > 0)$  scales the steepness of the mapping function. From the definition, we can see that higher bid price correspond to higher impression possibility.

#### 4.1.3. Impression confidence based on competitiveness

Suppose the sponsored search system maintains a universal set of keywords, which for example can be the union of the queries submitted by users in the historical search logs and the bid keywords given by advertisers. We denote this set as  $W = \{w_1, w_2, \dots, w_m\}$ , where  $m$  is the number of keywords.

Every keyword  $w_i$  ( $i = 1, 2, \dots, m$ ) in  $W$  might trigger its auctions  $\Theta_i$ . Each auction  $\theta \in \Theta_i$  will have its winning score  $\mu_\theta$ . When an ad group  $\omega$  participates in an auction  $\theta$ , we may compute the relevance score  $r_{\omega,\theta}$  and the impression probability based on bid price  $p_{\omega,\theta}$ . Note that  $\Theta_i$  might be empty if  $w_i$  is never matched by any types of matching functions.

Given a keyword, its competitiveness will also affect the possibility that the related ad groups will have an ad impression after bidding the keyword. We define the following impression confidence based on competitiveness to reflect the above effect.

**Definition 4** (IMPRESSION CONFIDENCE BASED ON COMPETITIVENESS). Let  $n$  be the average number of ad groups that are involved in the auctions triggered by the keyword  $w$ , the impression confidence based on competitiveness of an ad group in the auctions triggered by  $w$ , denoted by  $c$ , is defined as

$$c = 1 - \frac{1}{1 + e^{-\tau l}} + \frac{1}{1 + e^{\tau(n-l)}},$$

where  $l$  is the maximum number of ad slots in each search result page of a sponsored search system and  $\tau > 0$  is a coefficient.

From the definition, we can see that impression confidence based on competitiveness is a value in the interval  $[0, 1)$ . Fig. 4 shows some example curves of it. The coefficient  $\tau$  controls the steepness of the curve around  $l$ .

The intuitive explanation of impression confidence based on competitiveness is as follows. In the auction for a query keyword  $w$ , the number of ad slot  $l$  is a critical value to describe the competitiveness.

<sup>13</sup> Please refer to FootNote 9 for the detailed calculation process.

<sup>14</sup> The idea of using the standard cumulative Gaussian density in the definition is borrowed from (Graepel et al., 2010).

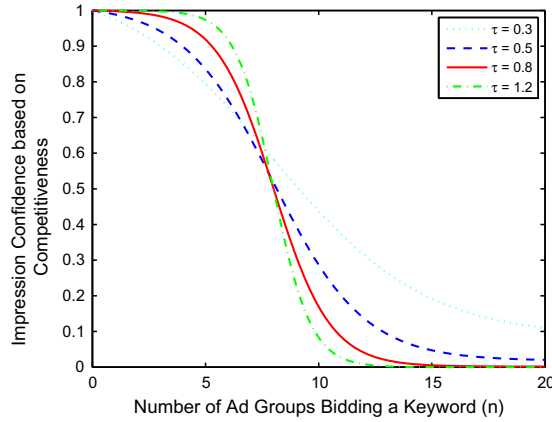


Fig. 4. The curves of impression confidence based on competitiveness.

- If there are no more than  $l$  ad groups involved in the auction, the winning score will be the floor rank score  $s_\epsilon$  and the auction is not competitive. The fewer auction participants we have, the easier a participant will achieve a high-ranked ad slot as long as his/her rank score exceeds  $s_\epsilon$ . Intuitively, in this case, the impression confidence based on competitiveness decrease from 1 very slowly as  $n$  increases from zero to  $l$ .
- If there are more than  $l$  ad groups participating in the auction, the winning score will be determined by the rank score of the ad group that wins the last ad slot. The more participants we have, the more competitive the auction will be. However, the increasing rate of the impression confidence based on competitiveness might almost “linearly” decrease when  $n$  increases in a certain extent around  $l$ . When  $n$  increases over this short range, the decreasing rate of impression confidence based on competitiveness should become slower with the increase of  $n$ . This is because when  $n$  is very large, adding one more new participant will not affect the already highly competitive auction, and the impression confidence based on competitiveness in the auction is almost zero.

It is clear that the formula for calculating impression confidence based on competitiveness (which is actually a translation of the sigmoid function) can well reflect the above two intuitions.

Actually, impression confidence based on competitiveness has a strong relation with rare keywords. Rare keywords are usually not bid by any advertisers or only bid by a small number of advertisers, and therefore their impression confidence based on competitiveness will be high. According to our statistics, the impression confidence of 97% rare keywords are almost 1. Therefore, if we can attach importance to keywords with high impression confidence, many of these keywords will be rare keywords.

#### 4.2. Optimization problem

The output of bid keyword suggestion is a subset of keywords  $W' \subseteq W$  with their optimal bid prices, i.e.,

$$\{(w'_1, b'_1), (w'_2, b'_2), \dots, (w'_k, b'_k)\} \subseteq W \times \mathcal{R},$$

where  $\mathcal{R}$  is the set of non-negative real numbers,  $b'_1, b'_2, \dots, b'_k$  are suggested bid prices, and  $k$  is the number of recommended keywords.

Given a target ad group  $\omega$ , keywords whose relevance scores with ad group  $\omega$  are less than the floor relevance score will be filtered out. Therefore, a subset  $W_\omega \subseteq W$  of much relevant candidate keywords are generated.

Selecting a subset  $W'$  of bid keywords among candidate keywords subset  $W_\omega \subseteq W$  is equivalent to determining an indicator vector  $\eta = (\eta_1, \eta_2, \dots, \eta_m)^T$ , where

$$\eta_i = \begin{cases} 1, & w_i \in W' \\ 0, & w_i \notin W' \end{cases} (i = 1, 2, \dots, m). \quad (1)$$

In order to find the optimal  $\eta$  and recommend reasonable bid prices  $b_i$  for each keyword  $w_i$ , we consider maximizing the revenue of the ad group which is highly related to the final profit the ad group can make from the service of sponsored search.

The constraints are the lower bound and the upper bound of the bid prices, and the budget of the ad group.

To ease the illustration, we further introduce two notations as below. (i)  $\rho_\omega$  – the average CTR of an ad group  $\omega$ , calculated from the historical log. Considering the definition of the impression probability of ad group  $\omega$  based on bid price  $b_\omega$ , denoted by  $p_{\omega, \theta}$ , we can approximately estimate the average probability of click on ad group  $\omega$  given bid price  $b_\omega$  as  $p_{\omega, \theta} \rho_\omega$ . (ii)  $v_\omega$  – the average true value of a click on the ads of an ad group  $\omega$ . It will be the upper bound of the bid prices as to make a profit.



The maximum of the winning price and the floor price  $b_\epsilon$  is the lower bound of the bid prices. The winning price for an ad group  $\omega$  in an auction  $\theta$  is calculated as  $\frac{\mu_\theta}{r_{\omega,\theta}}$ , where  $\mu_\theta$  is the winning score of the auction  $\theta$ , and  $r_{\omega,\theta}$  is the relevance score of ad group  $\omega$  in an auction  $\theta$ .

The budget constraint serves as the upper bound of the advertiser's cost. As the bid price is a natural upper bound of the cost in GSP, the upper bound of the advertiser's cost on an ad group  $\omega$  with bidding keyword  $w_i$  can be written as  $\sum_{\theta \in \Theta_i} b_{\omega,i} p_{\omega,\theta} \rho_\omega$ , in which  $b_{\omega,i}$  is the bid price for keyword  $w_i$  in the ad group  $\omega$ , and  $\Theta_i$  denotes all auctions triggered by the keyword  $w_i$ . For all candidate keywords of an ad group  $\omega$ , the upper bound of the advertiser's total cost should be

$$\sum_{w_i \in W_\omega} \left\{ \sum_{\theta \in \Theta_i} b_{\omega,i} p_{\omega,\theta} \rho_\omega \eta_i \right\}. \quad (2)$$

The objective is to maximize the expected advertiser revenue. Given bidding keyword  $w_i$ , the expected advertiser revenue of ad group  $\omega$  on this keyword can be approximately calculated as  $\sum_{\theta \in \Theta_i} (v_\omega - b_{\omega,i}) p_{\omega,\theta} \rho_\omega$ , where  $v_\omega$  is the average true value of the ad group  $\omega$ . Considering the impression confidence  $c_i$  based on competitiveness of all candidate keywords to an target ad group, we can write the total expected revenue as,

$$\sum_{w_i \in W_\omega} \left\{ c_i \sum_{\theta \in \Theta_i} (v_\omega - b_{\omega,i}) p_{\omega,\theta} \rho_\omega \eta_i \right\}. \quad (3)$$

To sum up the above constraints and the objective, we can obtain the following mixed integer optimization problem for an ad group  $\omega$ .

$$\begin{aligned} \max_{\eta_i, b_{\omega,i}} \quad & \sum_{w_i \in W_\omega} \left\{ c_i \sum_{\theta \in \Theta_i} (v_\omega - b_{\omega,i}) p_{\omega,\theta} \rho_\omega \eta_i \right\}, \\ \text{s.t.} \quad & \sum_{w_i \in W_\omega} \left\{ \sum_{\theta \in \Theta_i} b_{\omega,i} p_{\omega,\theta} \rho_\omega \eta_i \right\} \leq g_\omega, \\ & \max_{\theta \in \Theta_i} \left\{ b_\epsilon, \frac{\mu_\theta}{r_{\omega,\theta}} \right\} \leq b_{\omega,i} \leq v_\omega, (i = 1, 2, \dots, m). \end{aligned} \quad (4)$$

### 4.3. Efficient solution

The above problem is a mixed integer programming problem, which is NP-hard. To efficiently solve it, we employ an alternative optimization method. There are two variable vectors in the above problem, i.e.,  $\eta = \{\eta_i\}$  and  $b_\omega = \{b_{\omega,i}\}$ . We first fix  $b_\omega$  to the corresponding winning prices, and then the problem is converted to a binary integer optimization problem, which can be solved like a knapsack problem (Kellerer, Pferschy, & Pisinger, 2004). After that, we fix  $\eta$  to the optimal result of the binary integer optimization problem, and thus the problem is converted to a constraint optimization problem on  $b_\omega$ , which can be solved by means of sequential quadratic programming (SQP) (Cottle, Pang, & Stone, 2009). The above alternative process can be conducted iteratively, and we can prove that the process can converge in a finite number of steps (Tong & Koller, 2002).

Specifically, for the binary integer optimization problem, we solve it using the standard knapsack method (Kellerer et al., 2004). For the constraint optimization problem on  $b_\omega$ , we first rewrite the constraint optimization problem on  $b_\omega$  as the following form.

$$\begin{aligned} \min_{b_\omega} \quad & f(b_\omega), \\ \text{s.t.} \quad & h_j(b_\omega) \geq 0 (j = 1, 2, 3). \end{aligned} \quad (5)$$

Here we have,

$$\begin{aligned} f(b_\omega) &= - \sum_{w_i \in W_\omega} \left\{ c_i \sum_{\theta \in \Theta_i} (v_\omega - b_{\omega,i}) p_{\omega,\theta} \rho_\omega \eta_i \right\}, \\ h_1(b_\omega) &= g_\omega - \sum_{w_i \in W_\omega} \left\{ \sum_{\theta \in \Theta_i} b_{\omega,i} p_{\omega,\theta} \rho_\omega \eta_i \right\} \geq 0, \\ h_2(b_\omega) &= v_\omega e - b_\omega \geq 0, \\ h_3(b_\omega) &= b_\omega - d_\omega \geq 0, \end{aligned} \quad (6)$$

where  $e$  is a vector with all its elements equal to 1, and  $d_\omega$  is a vector with the values  $\max_{\theta \in \Theta_i} \left\{ b_\epsilon, \frac{\mu_\theta}{r_{\omega,\theta}} \right\}$  in the corresponding elements.

The above problem can be approximately solved by SQP, in which a quadratic programming (QP) subproblem is solved in each iteration by linearizing the constraints and approximating the following Lagrangian function  $L(b_\omega, \eta)$  quadratically.

$$L(b_\omega, \eta) = f(b_\omega) - \sum_{i=1}^3 \lambda_i h_i(b_\omega). \quad (7)$$

Here  $\lambda = \{\lambda_i\}$  is Lagrangian multiplier. Starting from any initial  $b_\omega^{(0)}$ , suppose  $b_\omega^{(j)}$  is the solution in the  $j$ -th iteration,  $\lambda^{(j)}$  is the corresponding Lagrangian multipliers, and  $H^{(j)} = H(b_\omega^{(j)}, \lambda^{(j)})$  is the Hessian of the Lagrangian function. Then the following QP subproblem should be solved in the  $(j+1)$ -th iteration.

$$\begin{aligned} \min_z \quad & \frac{1}{2} z^T H^{(j)} z + \nabla L(b_\omega, \eta)^T z + L(b_\omega, \eta), \\ \text{s.t.} \quad & \nabla h_i(b_\omega^{(j)})^T z + h_i(b_\omega^{(j)}) \geq 0 \quad (i = 1, 2, 3). \end{aligned} \quad (8)$$

Here  $\nabla$  denotes gradient calculus. If  $z^{(j)}$  is the solution of the above QP subproblem and  $u^{(j)}$  is the corresponding multiplier of this subproblem, then we use the following formulas to update the solution of the SQP problem.

$$\begin{aligned} b_\omega^{(j+1)} &= b_\omega^{(j)} + z^{(j)}, \\ \lambda^{(j+1)} &= u^{(j)}. \end{aligned} \quad (9)$$

The above iterations lead to a local approximation of the solution. The algorithm can be stabilized by line search.

To sum up, by solving the binary integer optimization problem and the constraint optimization problem on  $b_\omega$  alternatively, we can get the optimal values  $\eta^*$  and  $b_\omega^*$ . Then we can select the keywords corresponding to the top-scored elements in  $\eta^*$  to compose the recommended keyword set, and set the suggested bid prices according to  $b_\omega^*$ . The algorithm flowchart is depicted in Fig. 5.

## 5. Experimental evaluation

In this section, we first introduce our experimental settings, including the datasets, ad group sampling method, baseline algorithms, and evaluation measures. Then we report the experimental results on our proposed algorithm and make analysis and discussion.

### 5.1. Experimental settings

#### 5.1.1. Datasets

The data in our experiments mainly came from three sources provided by a commercial search engine: search log, auction log, and advertiser database. The log data was collected in one month (April 2011) and is split into two periods: the first half of the month (01-April-2011 to 15-April-2011) for training and the second half (16-April-2011 to 30-April-2011) for testing. The database is a snapshot at the end of the month, which contains 31 million ad groups.

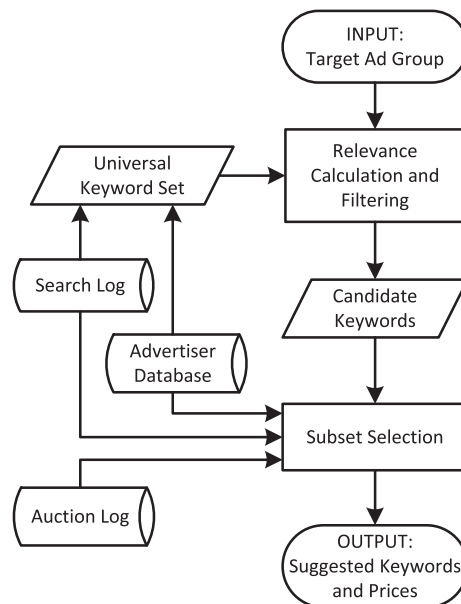


Fig. 5. The algorithm flowchart of bid keyword suggestion.

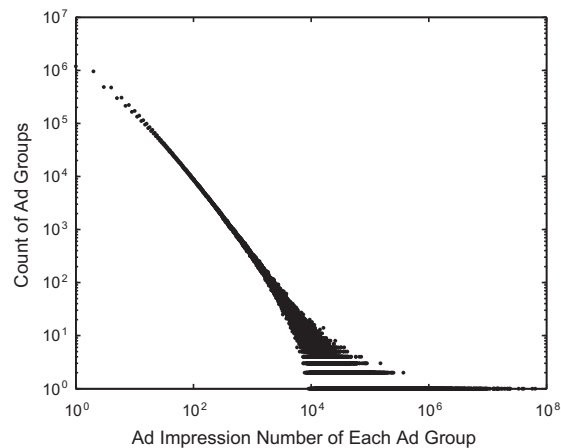


Fig. 6. The power law distribution of the ad impression number for each ad group.

We used the search log data in the training period to build the candidate keyword set. We extracted over 1.5 billion raw queries and reduced them to about 1 billion using query normalization. There are 623 million unique queries in the training period and 639 million unique queries in the test period. The information of ad impressions and clicks is also available in the search log, from which we can compute the average CTR for each ad group. We used the auction log in the training period to compute the winning scores and the winning prices. We used the advertiser database to extract budget, ad copies, bid keywords, bid prices, and bid numbers, and then compute the average true values,<sup>15</sup> the impression confidence scores based on competitiveness, and the relevance scores.

#### 5.1.2. Ad group sampling

As mentioned in the introduction, the motivation of our proposal is to help those advertisers who failed in their auctions due to limited budget and improper keyword bidding. Intuitively, it is not necessary to perform such recommendation to those big advertisers, who tend to set high bid prices on their preferred keywords, and obtain a lot of ad impressions and clicks without much concern about the cost. In this regard, it will be informative if we can investigate the performance of our proposed technology on different types of advertisers. In our experiments, we extracted two sets of ad groups as below.

- **Small advertiser set.** The set of ad groups with fewer than 500 ad impressions in the training period. We assume these advertisers are not very successful, and are desired for the help from bid keyword suggestion.
- **Big advertiser set.** The set of ad groups with more than 5,000 ad impressions in the training period. We assume that they have been successful enough and might not really care about the recommendation of keywords.

Fig. 6 shows the distribution of the ad impressions for the ad groups. From our statistics, we know that the majority (91.9%) of the ad groups belong to the small advertiser set and majority (82.2%) ad impression volume belong to the big advertiser set. This somehow shows both sets are much important in sponsored search and our proposed technology should not only work on the small advertiser set, but also not hurt the performance of the big advertiser set of ad groups. For each of the two sets, we randomly sampled 800 ad groups as the target ad groups in our experiments.

#### 5.1.3. Comparison algorithms

We compared the proposed algorithm with the following methods for bid keyword suggestion on the two target sets. These methods are representative approaches on keyword recommendation as mentioned in Section 2. After we get the suggested keywords and their bid prices, we run simulations (Acharya, Krishnamurthy, Deshpande, Yan, & Chang, 2007) on the auction log in the testing period see the performance.

1. **No Recommendation (NR).** We simply used the original bid keywords of the target ad group to get the performance in the testing period. This is the baseline for all the other algorithms.
2. **Popular and Relevant (PR).** This algorithm is implemented in accordance with the common idea of existing bid keyword recommendation tools in commercial search engines, such as AdWords Keyword Tool.<sup>16</sup> It tends to recommend those keywords that are relevant to the ad group and are popular in web users and many other advertisers.

<sup>15</sup> Note that we approximately calculate the true value according to the idea in (Athey & Nekipelov, 2010), i.e., we ran a simulation on the auction log to get the incremental cost per click (ICPC) (Athey & Nekipelov, 2010) for each keyword and use it as the estimated true value for the keyword.

<sup>16</sup> <http://www.googlekeywordtool.com/>.

Specifically, we obtain the search frequency of a keyword and we calculate the relevance score between a keyword and an ad group. The relevance calculation method is similar to the process explained in FootNote 9. The only difference is that we regard all the text content (ad titles, ad copies, and keywords) in the ad group as a bag-of-word *document* and then extract the features used for the relevance scoring function like the similarities between the keyword and the *document*. The keywords that have the higher product values of the relevance scores and the logarithmic search frequencies will be prioritized for recommendation. In addition, the keywords whose relevance scores or logarithmic search frequencies are below some given thresholds are not qualified for recommendation.

3. **Amnesiac Averaged Perceptron (AAP)**. We implement the learning based approach proposed in (Gupta et al., 2009). The proposed keyword similarity model is adaptive to combine arbitrary existing similarity functions, such as the syntactic, taxonomy, and search results click similarity. It can be the state of the art algorithm for keyword matching. In addition, considering the dynamics in sponsored search, the weights of the voting perceptrons are updated by the time.
4. **Bid Keyword Suggestion (BKS)**. This is our proposed algorithm, which generates the recommended keywords by solving a subset selection problem.
5. **Naive Bid Keyword Suggestion (N-BKS)**. This is a simplified version of our proposed algorithm, which generates the recommended keywords by solving the binary integer optimization problem and use the winning prices as the suggested bid prices.

For NR, we used the original bid price on each keyword. For PR, N-BKS, and AAP, we set the bid price of each recommended keyword to be its mean of winning prices in all auctions. For BKS and N-BKS, the relevance score  $r_{w,\theta}$  is normalized to the interval  $[0, 10000]$  and the parameter setting is  $\lambda = 500$ ,  $\epsilon = 500$ . For the impression confidence, we find  $\tau = 0.8$  can best reflect the intuitions in the definition by empirical study. For all the algorithms, the upper bound on the number of the recommended keywords is 2000. We assume that the target ad groups will bid on the recommended keywords with the recommended bid prices. Then we perform simulation on the data in the testing period. The target ad group will participate in the auctions in the test data with the constraint of their budgets. We then compute the performances of all the comparison algorithms in terms of the evaluation measures as described in the next subsection.

#### 5.1.4. Evaluation measures

We used the following five measures in our experiments.

1. **Ad impressions**. This measure reflects the return of an ad group on impressions. We simulated the auctions for the recommended keywords for each target ad group in the testing period and calculated the corresponding ad impression.
2. **Expected clicks**. This measure reflects the return of an ad group on clicks. We use the average CTR and ad impression to calculate the expected click.
3. **Average cost per click**. This measure reflects the average cost per click of an ad group in the auctions according to GSP mechanism.
4. **Expected advertiser revenue**. This measure reflects the expected revenue of the advertiser as calculated in Eq. 3.
5. **Advertiser total cost (a.k.a., search engine revenue)**. This measure reflects the total cost of the advertiser as calculated in Eq. 2, which equals to the search engine revenue.

## 5.2. Experimental results

We present the performances of the algorithms NR, PR, AAP, BKS and N-BKS with respect to ad impressions, expected clicks, average cost per click, expected advertiser revenue and advertiser total cost (a.k.a, search engine revenue). Note that we normalized the values under each evaluation measure by dividing them by the maximum value in the corresponding results, to protect the business secrets of the search engine.

### 5.2.1. Ad impressions

Fig. 7 shows the average ad impressions of each algorithm on the two sets of ad groups. From the figure we have the following observations.

On the small advertiser set, (i) PR, AAP, N-BKS, and BKS successfully achieve a lot more impressions as compared to NR, showing that the four recommendation methods can all help the small advertisers improve their performance. (ii) AAP slightly performs better than PR. The reason is that taking more algorithms such as search result click similarity will recommend more practical keywords and thus improve the impressions. (iii) N-BKS and BKS significantly outperform PR and AAP. The reason is that N-BKS and BKS takes competitiveness into consideration besides relevance and keyword frequency. They tend to recommend a group of relevant keywords with low competitiveness, which can help the advertisers accumulate satisfying ad impressions. (iv) BKS outperforms N-BKS. The reason is that BKS seeks the optimal bid prices compared with N-BKS. Overall speaking, BKS is the best choice for the small advertisers.

On the big advertiser set, (i) The order of the performance of NR, AAP, N-BKS, and BKS are the same as that of the small advertiser set. The reasons are similar to what we have explained for the small advertiser set. (ii) Differently, PR performs the




Fig. 7. The average ad impressions.

best among the algorithms. The big advertisers can afford the high bid prices on hot keywords and they often have big budgets, so the keywords suggested by PR can be well utilized by them. (iii) BKS performs better than NR, indicating that the proposed method does not hurt the performance of the big advertisers.

#### 5.2.2. Expected clicks

The expected ad clicks for the five algorithms on the two sets of ad groups are shown in Fig. 8. From the figure we have the following observations. For the small advertisers, we can see that the trend is very similar to that of the ad impressions. All the algorithms improve the click number as compared to NR, and BKS still achieves the best performance by recommending less competitive keywords with optimal bid prices. For the big advertisers, (i) BKS still outperforms NR, showing that the proposed method does not hurt the performance of the big advertisers. (ii) PR performs the best among the algorithms. The reason is similar to what happens to ad impressions. (iii) AAP relies a lot on the keyword similarity and popularity, so it gets a second highest expected clicks for the big advertisers.

#### 5.2.3. Average cost per click

Fig. 9 shows the average CPC of each algorithm on the two sets of ad groups. From the figure, we have the following observations. (i) On both sets, BKS and N-BKS lead to the lowest average CPC. This is because the keywords they recommend are mostly less-competitive keywords, whose winning prices are just above the floor price. (ii) On both sets, the keywords suggested by PR are much more expensive than those suggested by AAP, N-BKS, and BKS. The reason is that PR tends to recommend popular keywords which usually are highly competitive and thus with high CPC.

#### 5.2.4. Expected advertiser revenue

Fig. 10 shows the expected advertiser revenue of each algorithm on the two sets of ad groups. From the figure, we have the following observations. (i) For the small advertisers, BKS and N-BKS can help them get the biggest reward compared with NR, PR, and AAP. The keywords recommended by PR and AAP are often very popular and expensive, so the corresponding revenue for the advertisers will be small. (ii) For the big advertisers, PR and AAP can bring more reward than BKS and N-BKS, for these advertisers have big budgets and they can afford the keywords with high bid prices.

#### 5.2.5. Advertiser total cost (search engine revenue)

Fig. 11 shows the average total cost for different algorithms on the two sets of ad groups. From the figure, we have the following observations. (i) On the small advertiser set, PR, AAP, N-BKS, and BKS can all enlarge their total cost. However,

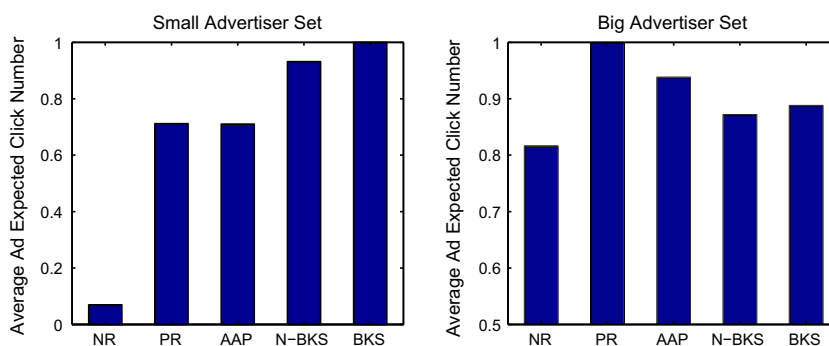


Fig. 8. The average expected clicks.

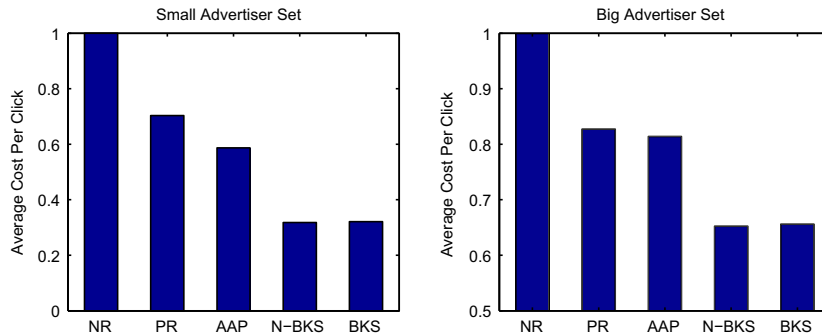


Fig. 9. The average cost per click.




Fig. 10. The average ad group expected revenue.

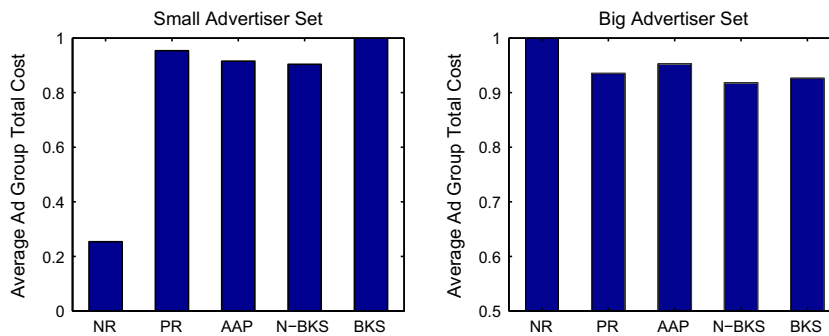


Fig. 11. The advertiser total cost (search engine revenue).

considering the comparison on expected click and average CPC, we find that the increased cost by BKS is the most valuable to the advertisers, for they can get more clicks with less CPC by this method. (ii) On the big advertiser set, the total cost among the five algorithms does not have much difference. (iii) The advertiser total cost actually equals to the search engine revenue. The above results show that BKS can help bring more revenue to the search engine from the small advertisers. As explained in the introduction, the extra revenue comes from utilizing the empty ad slots related to the less-competitive keywords.

#### 5.2.6. Summary and discussions

To sum up all the aforementioned experimental results, we can get the following conclusions.

On the small advertiser set, (i) the proposed BKS algorithm can help the advertisers get largely improved performance on impression, click, CPC, and advertiser revenue compared with several popularly-used keyword suggestion methods. (ii) Although all compared methods can help increase the search engine revenue substantially, from the perspective of the advertisers, the increased cost by the proposed BKS algorithm is much more rewarding as BKS can have more expected clicks and lower average CPC. This clearly demonstrates the effectiveness of our proposal for both the advertisers and the search engine.



On the big advertiser set, (i) the BKS algorithm can still improve their performance on each evaluation metric, especially the average CPC. (ii) PR can help the big advertisers achieve more, because the big advertisers often have enough budgets and can bid popular keywords with much higher price, and then the philosophy of PR is more favorite to them.

In total, our proposed technology can not only work well on the small advertisers that the majority of the ad groups belong to, but also not affect the performance of the big advertisers.

## 6. Conclusion and future work

In this paper, we conducted an intensive study on the sponsored search data of a commercial search engine, and identified two problems with the current sponsored search systems: (i) a large number of advertisers cannot get satisfying ad impressions and clicks due to improper keyword bidding; (ii) a significant portion of search queries have no ads displayed in their search result pages. To tackle these problems, we proposed a novel bid keyword suggestion technology. The key component of the technology is a subset selection model which can be efficiently solved by alternatively utilizing a knapsack problem solver and an SQP solver. The objective function for subset selection ensures that the proposed technology can simultaneously help advertisers and improve search engine revenue. Our experimental results verified the effectiveness of the proposed algorithm, in terms of several measures such as ad impression, expected clicks, average CPC, search engine revenue, and advertiser revenue.

For future work, we plan to work on the following directions. The first direction is to consider user experience in the evaluation. As we all know, sponsored search is a three-player game in which user is one of the key players. Therefore, we should also consider user experience besides relevance in order to obtain a comprehensive evaluation of the entire ecosystem. The second direction is to further refine the proposed bid keyword suggestion model to perform optimization for multiple ad groups simultaneously. In particular, we will consider the inter-competition within the ad groups that consume our keyword recommendation service.

## Acknowledgments

This work is partially supported by National Natural Science Foundation of China under Grant No. 61170184, the National 863 Project of China under Grant No. 2013AA013204, and the Tianjin Municipal Science and Technology Commission under Grant No. 13ZCZDZX02200 and 13ZCZDZX01098. This work is also partially supported by Microsoft Research Asia under Grant No. FY12-RES-OPP-094.

## References

- Abhishek, V., & Hosanagar, K. (2007). Keyword generation for search engine advertising using semantic similarity between terms. In *Proceedings of the ninth international conference on electronic commerce* (pp. 89–94). ACM.
- Acharya, S., Krishnamurthy, P., Deshpande, K., Yan, T., & Chang, C. (2007). A simulation framework for evaluating designs for sponsored search markets. In *WWW 2007 workshop on sponsored search auctions*.
- Aggarwal, G., Goel, A., & Motwani, R. (2006). Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM conference on electronic commerce* (pp. 1–7). ACM.
- Anderson, C. (2008). The long tail: Why the future of business is selling less of more. *Hyperion*.
- Athey, S., & Nekipelov, D. (2010). A structural model of sponsored search advertising auctions. In *Sixth ad auctions workshop*.
- Bartz, K., Murthi, V., & Sebastian, S. (2006). Logistic regression and collaborative filtering for sponsored search term recommendation. In *Second workshop on sponsored search auctions* (pp. 61–70). Citeseer.
- Berg, J., Greenwald, A., Naroditskiy, V., & Sodomka, E. (2010). A knapsack-based approach to bidding in ad auctions. In *The 19th European conference on artificial intelligence (ECAI)*.
- Broder, A., Ciccolo, P., Fontoura, M., Gabrilovich, E., Josifovski, V., & Riedel, L. (2008). Search advertising using web relevance feedback. In *Proceedings of the 17th ACM conference on information and knowledge management* (pp. 1013–1022). ACM.
- Broder, A., Ciccolo, P., Gabrilovich, E., Josifovski, V., Metzler, D., Riedel, L., et al. (2009). Online expansion of rare queries for sponsored search. In *Proceedings of the 18th international conference on World Wide Web* (pp. 511–520). ACM.
- Broder, A., Fontoura, M., Gabrilovich, E., Joshi, A., Josifovski, V., & Zhang, T. (2007). Robust classification of rare queries using web knowledge. In *Proceedings of the 30th annual international ACM SIGIR conference on research and development in information retrieval* (pp. 231–238). ACM.
- Broder, A., Gabrilovich, E., Josifovski, V., Mavromatis, G., & Smola, A. (2011). Bid generation for advanced match in sponsored search. In *Proceedings of the fourth ACM international conference on Web search and data mining* (pp. 515–524). ACM.
- Chen, Y., Xue, G., & Yu, Y. (2008). Advertising keyword suggestion based on concept hierarchy. In *Proceedings of the international conference on Web search and Web data mining* (pp. 251–260). ACM.
- Cottle, R., Pang, J., & Stone, R. (2009). The linear complementarity problem. 60, Society for Industrial Mathematics.
- Dai, H., Zhao, L., Nie, Z., Wen, J., Wang, L., & Li, Y. (2006). Detecting online commercial intention (oci). In *Proceedings of the 15th international conference on World Wide Web* (pp. 829–837).
- Edelman, B., Ostrovsky, M., & Schwarz, M. (2005). Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. Technical report, National Bureau of Economic Research.
- Even-Dar, E., Mirrokni, V., Muthukrishnan, S., Mansour, Y., & Nadav, U. (2009). Bid optimization for broad match ad auctions. In *Proceedings of the 18th international conference on World Wide Web* (pp. 231–240). ACM.
- Feng, J., Bhargava, H., & Pennock, D. (2003). Comparison of allocation rules for paid placement advertising in search engines. In *Proceedings of the 5th international conference on electronic commerce* (pp. 294–299). ACM.
- Fuxman, A., Tsaparas, P., Achan, K., & Agrawal, R. (2008). Using the wisdom of the crowds for keyword generation. In *Proceedings of the 17th international conference on World Wide Web* (pp. 61–70). ACM.
- Graepel, T., Candela, J., Borchert, T., & Herbrich, R. (2010). Web-scale bayesian click-through rate prediction for sponsored search advertising in microsofts bing search engine. In *Proceedings of the 27th international conference on machine learning*. San Francisco, CA: Morgan Kaufmann.

- Gupta, S., Bilenko, M., & Richardson, M. (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* (pp. 1165–1174). ACM.
- Hillard, D., Schroedl, S., Manavoglu, E., Raghavan, H., & Leggetter, C. (2010). Improving ad relevance in sponsored search. In *Proceedings of the third ACM international conference on Web search and data mining* (pp. 361–370). ACM.
- Joachims, & Thorsten, (2002). Optimizing search engines using clickthrough data. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 133–142). ACM.
- Joshi, A., & Motwani, R. (2006). Keyword generation for search engine advertising. In *Sixth IEEE international conference on data mining workshops, 2006 (ICDM workshops 2006)* (pp. 490–496). IEEE.
- Kellerer, H., Pferschy, U., & Pisinger, D. (2004). *Knapsack problems*. Springer.
- Kiritchenko, S., & Jilinc, M. (2008). Keyword optimization in sponsored search via feature selection. In *JMLR: Workshop and conference proceedings 4* (pp. 122–134).
- Kitts, B., & Leblanc, B. (2004). Optimal bidding on keyword auctions. *Electronic Markets*, 14, 186–201.
- Pandey, S., Punera, K., Fontoura, M., & Josifovski, V. (2010). Estimating advertisibility of tail queries for sponsored search. In *Proceeding of the 33rd international ACM SIGIR conference on research and development in information retrieval* (pp. 563–570). ACM.
- Radlinski, F., Broder, A., Ciccolo, P., Gabrilovich, E., Josifovski, V., & Riedel, L. (2008). Optimizing relevance and revenue in ad search: A query substitution approach. In *Proceedings of the 31st annual international ACM SIGIR conference on research and development in information retrieval* (pp. 403–410). ACM.
- Ravi, S., Broder, A., Gabrilovich, E., Josifovski, V., Pandey, S., & Pang, B. (2010). Automatic generation of bid phrases for online advertising. In *Proceedings of the third ACM international conference on Web search and data mining, WSDM*.
- Sodomka, E., Lahaie, S., & Hillard, D. (2011). A hierarchical model for value estimation in sponsored search. In *EC 2011 seventh ad auctions workshop*.
- Tong, S., & Koller, D. (2002). Support vector machine active learning with applications to text classification. *The Journal of Machine Learning Research*, 2, 45–66.
- Varian, H. (2007). Position auctions. *International Journal of Industrial Organization*, 25, 1163–1178.
- Zainal-Abidin, A., & Wang, J. (2010). Maximizing clicks of sponsored search by integer programming. In *Proceedings of the ADKDD*.
- Zhang, W., He, X., Rey, B., & Jones, R. (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* (pp. 853–854). ACM.
- Zhang, W., Zhang, Y., Gao, B., Yu, Y., Yuan, X., & Liu, T. Y. (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* (pp. 1177–1185). Springer.