



Dynamic dual adjustment of daily budgets and bids in sponsored search auctions

Jie Zhang ^a, Yanwu Yang ^{b,*}, Xin Li ^c, Rui Qin ^a, Daniel Zeng ^{a,d}

^a State Key Lab of Management and Control for Complex Systems, CAS, Beijing 100190, China

^b School of Management, Huazhong University of Science and Technology, Wuhan 430074, China

^c Department of Information Systems, City University of Hong Kong, Kowloon Tong, Hong Kong

^d Department of Management Information Systems, The University of Arizona, USA

ARTICLE INFO

Article history:

Received 3 May 2012

Received in revised form 11 July 2013

Accepted 16 August 2013

Available online 29 August 2013

Keywords:

Sponsored search auction

Budget adjustment

Continuous reinforcement learning

Dynamic adjustment

ABSTRACT

As a form of targeted advertising, sponsored search auctions attract advertisers bidding for a limited number of slots in paid online listings. Sponsored search markets usually change rapidly over time, which requires advertisers to adjust their advertising strategies in a timely manner according to market dynamics. In this research, we argue that both the bid price and the advertiser (claimed) daily budget should be dynamically changed at a fine granularity (e.g., within a day) for an effective advertising strategy. By doing so, we can avoid wasting money on early ineffective clicks and seize better advertising opportunities in the future. We formulate the problem of dual adjusting (claimed) daily budget and bid price as a continuous state – discrete action decision process in the continuous reinforcement learning (CRL) framework. We fit the CRL approach to our decision scenarios by considering market dynamics and features of sponsored search auctions. We conduct experiments on a real-world dataset collected from campaigns conducted by an e-commerce advertiser on a major Chinese search engine to evaluate our dual adjustment strategy. Experimental results show that our strategy outperforms two state-of-the-art baseline strategies and illustrate the effect of adjusting either (claimed) daily budget or bid price in advertising.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

With the recent trend of “economics meets search” [16], there is a rapid growth of search engine-based advertisements. Such sponsored search is often managed through a form of auctions, where a bidding contract is triggered once a query of certain keywords is submitted. Advertisers need to carefully manage their advertising strategies through the parameters of the bidding setup in order to compete with other advertisers. The high volume of search demands makes bidding in sponsored search auctions a continuous and dynamic process. Once an advertiser adjusts her advertisements and/or advertising strategies, rankings of the sponsored links and cost-per-click will be changed accordingly. Thus, to achieve an effective campaign, advertisers should continuously monitor the auction market and adjust their strategies in response to market dynamics.

Among all factors to be managed, budget is one critical factor that is endogenous to the auction process [3]. Previous research usually assumes budget to be fixed and take it as a constraint in strategy development. For example, some research [18,24] studied the allocation of budget over keywords. This method cannot effectively deal with the dynamic sponsored search auctions provided by major search engines (e.g., Google). It is necessary to examine budget allocation at finer

granularity, such as real-time adjustment within a day, for advertising. At this level, so far, most studies considered only bid price adjustment [13,22] and ignored budget control.

In the lifecycle of advertising campaigns in sponsored search, budget decisions occur at three levels [35]: allocate budget across search markets [34], distribute budget over a series of temporal slots (e.g. day) [36,37], and adjust the claimed budget within a temporal slot (to simplify the wording, this paper uses ‘day’ to replace ‘temporal slot’ hereafter). Note that after deciding each day’s budget during a given promotional period, an advertiser can claim a daily budget (named daily budget in short) which is different from their actual internal budget (named daily budget limit in short). By changing the (claimed) daily budget, the advertiser can restrict the amount of clicks to be directed to their sponsored links by the search engine in a unit time. Thus, in a sponsored search auction, after an advertiser determines her daily budget limit for a campaign (or each keyword of the campaign), she should monitor the advertising performance and adjust the daily budget together with bid prices corresponding to market dynamics, until the end of the daily advertising schedule or when the day’s remaining budget is zero. Thus she can avoid wasting money on ineffective clicks and save money for future opportunities.

In a previous work [35] we developed a hierarchical budget optimization framework (BOF), which considers all three levels of budget allocation in the entire lifecycle of advertising campaigns in sponsored search auctions. Following that, this work aims to tackle the advertising strategy optimization problem by dynamically adjusting daily budget and bid price within a day. To the best of our knowledge, our work might be the first effort in this direction.

* Corresponding author. Tel.: +86 27 87542254; fax: +86 27 87542254.

E-mail addresses: zhangjie.isec@gmail.com (J. Zhang), yangyanwu.isec@gmail.com (Y. Yang), Xin.Li.PhD@gmail.com (X. Li), qinrui.isec@gmail.com (R. Qin), zeng@email.arizona.edu (D. Zeng).

We address the dual adjustment problem as a continuous state – discrete action decision process in the continuous reinforcement learning (CRL) framework, since it can be viewed as a special multi-stage dynamic decision problem. Specifically, we employ continuous-time and continuous-state reinforcement learning [7] to capture the continuous state variables in the problem (such as the day's remaining budget, click dynamics, etc.). We adapt the CRL framework from two perspectives to meet the requirements of the advertising decision problem in sponsored search auctions. First, in classic CRL, behaviors of the environment were usually given as prior knowledge. However, we don't have information due to the complicated nature of sponsored search. Thus, we employ a Back Propagation Neural Network (BPNN) to approximate the parameters of the environment using historical data. Second, since some search engines (e.g., Google) only permit making a limited number of changes on (claimed) daily budget during a day, we employ step functions to represent the multiple discrete actions, which can occur at any time of the day. We provide an algorithm to solve our CRL-based dual adjustment model that iteratively optimizes the actions along the decision temporal points. In order to evaluate the effectiveness of our proposed approach, we conduct some experiments using a real-world dataset collected from search advertising campaigns conducted by an e-commerce advertiser on a major Chinese search engine. Experimental results illustrate that our strategy outperforms two state-of-the-art baseline strategies. They also show that budget adjustment plays a more important role in the dual adjustment strategy.

The contributions of this work can be summarized as follows:

- (1) We frame the dual adjustment problem (on changing daily budget and bid price within a day) for advertising strategy development in sponsored search auctions.
- (2) We build a continuous-time, continuous-state, and discrete action model under a CRL framework that fits the decision scenarios of the dual adjustment problem in sponsored search auctions.
- (3) We develop an iterative numerical approach that can efficiently solve our proposed model for online applications.

The rest of this paper is organized as follows. The next section briefly reviews some relevant literature. In Section 3, we describe budget decision scenarios in sponsored search auctions and state the research problem under consideration. Section 4 presents our dual adjustment model based on CRL. Section 5 provides a numerical solution to our proposed model. Sections 6 and 7 report the experiments to evaluate our dual adjustment model and discuss the implications of results. Section 8 concludes this work.

2. Literature review

Effective advertising strategy is an important problem in marketing. In the middle of the last century, [33] proposed the concept of advertising effectiveness and equations on sales response dynamics. They provided an optimal solution for the allocation of limited budgets considering sales dynamics. [23] introduced the concept of advertising goodwill to reflect the flow of current advertising expenditures. They considered that aggregated advertising effectiveness would influence future budget decisions and built a dynamic adjustment framework to optimize advertising strategies and price policies. The framework was later generalized by [28] into the case with limited budgets.

In recent years, many researchers have attention on the impact of market dynamics on advertisers' decision, such as budget allocation and adjustment [12,19,26,27,30]. Some studies recognized that the shape of the advertising response function plays an important role in advertising strategies [20,21,29,32]. Particularly, the S-shaped response function was carefully examined [10,14] since its convexity at low expenditure levels makes it easy to obtain the periodic optima in practice. Krishnan and Jain [17] investigated the optimal advertising policy for new products considering the influence of information diffusion

and concluded that optimal advertising strategies are determined by the advertising effectiveness, discount rate, and the ratio of advertisement to profits.

The increase in sponsored search auctions has led to advertising competitions in this unique auction market environment [15]. In such environments, various mathematic programming algorithms have been developed to improve advertising strategies. Integer programming and nonlinear programming were used to find optimal solutions for budget allocation over sponsored keywords [18,24]. OZluk and Cholette [24] showed that price elasticities of the click-through rate and response functions were key factors for budget decisions, and investing in more keywords under a certain threshold could help improve advertisers' profits. Fruchter and Dou [11] established an optimal control model to study the optimal budget allocation problem among web portals, and used dynamic programming to derive the analytical solution to the problem.

Considering the dynamic nature of sponsored search auction markets [36], it is natural to formulate the budget optimization problem as a Markov decision process [8] or an optimal control model [1,11]. Archak et al. [1] showed that under a reasonable assumption, online advertising has positive carryover effects on the propensity and form of user interactions with the same advertiser in the future. Based on the Nerlove–Arrow advertising framework [23], Rutz and Bucklin [25] proposed a dynamic linear model to capture the potential spillover from generic to branded paid search. The budget optimization problem can also be established as an online (multiple-choice) knapsack problem [4,5], from which advertisers can achieve a provably optimal competitive ratio. By considering bid dynamics and rankings of advertisers, a cyclical bid adjustment model in a two-player competition game was studied by Zhang and Feng [38], where an equilibrium bidding price for two advertisers can be obtained. Some researches have explored periodical budget allocation strategies by considering temporal features, such as weekday, weekends, and months of a year [9]. However, this exploration is not detailed enough for real-time strategy adjustment.

3. Budget decisions in sponsored search auctions

3.1. Budget decision scenarios

Three different budget decision scenarios occur during the lifecycle of sponsored search advertising campaigns. Fig. 1 describes budget decision scenarios according to temporal granularity, with two dimensions: over time (the horizontal axis) and across markets (the vertical axis). The interested reader is referred to see article [35] for more details of budget decision scenarios in sponsored search auctions.

Suppose the advertiser intends to conduct a sponsored search marketing campaign. First, at a long-term level, an advertiser needs to decide how to allocate her search advertising budget across multiple markets, given a predetermined total budget. Second, at a medium-term level, the advertiser needs to distribute her advertising budget across a series of temporal slots/days (named daily budget limit). Here, if necessary, the advertiser can optimize one day's budget limit based on historical advertising performance. However, once determined, the advertiser would expect to use up all budget in that temporal slot. Third, in each temporal slot of an ongoing advertising campaigns, an advertiser can also adjust some factors, such as (claimed) daily budget and bid price, in real-time to affect the ranking of the advertiser's sponsored links in response to advertising dynamics.

Budget strategies at these three levels complement each other and form an integrated budget optimization problem chain. That is, results of higher-level decisions constrain lower-level decisions, and operational results at lower levels create feedback for decisions at higher levels. Moreover, budget operations at these three levels all interact with outside environments, which contain great uncertainties ranging from advertising resources to advertising performance.

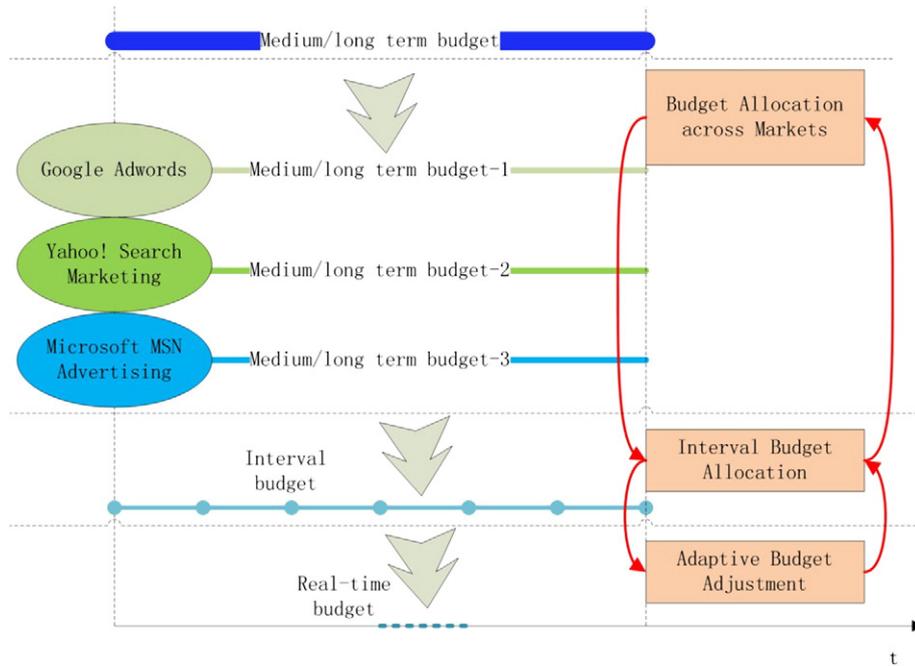


Fig. 1. A multi-level framework of budget decisions in search advertisements.

3.2. Problem statement

As explained above, in sponsored search auctions an advertiser can tweak advertising strategies at multiple levels of temporal granularity. However, in terms of real-time adjustment of advertising strategies, previous literature often relied on adjusting bid prices. Nevertheless, from a budget perspective, advertisers can set a claimed daily budget (named daily budget) in real-time to affect advertising performance. The search engine would use the (claimed) daily budget to decide whether to stop directing traffic towards the sponsored links. If the (claimed) daily budget is reached, the advertisements will not appear on the search engine. Note that the (claimed) daily budget may not equal the actual remaining budget of the day (which is a non-increasing function due to the cumulative consumption of budgets by received clicks). Thus, by changing (claimed) daily budget, an advertiser can assign her budget to specific temporal periods. For instance, when search demands are relatively high but conversion rates are relatively low, it would be wise to lower the daily budget and save budget for later traffic

with higher potential to purchase. Both bid price and daily budget can be used in managing sponsored search auctions.

This study focuses on the use of real-time budget adjustment in advertising strategy management. Specifically, suppose the budget for a temporal slot (i.e., a day) is specified, we want to know how to dynamically adjust daily budget and bid price to improve advertising performance.

The notations used in this paper are listed in Table 1. In this paper, we assume that a daily budget limit y_{ij} can be determined using the BOF framework given the overall budget B and requirements of multiple search engines and the length of the campaign. Thus, an advertiser cannot spend over y_{ij} in a day. Under this constraint, at any time t the advertiser can change bid price, β , and (claimed) daily budget, d , considering the change of market environments, specifically Cost-Per-Click (CPC) c , effective click-through rate (CTR), ρ , and remaining budget of a day, b . Here, CPC is affected by the bidding of all advertisers on the market. Defining clicks on sponsored links that can generate value to the advertiser as effective clicks, effective CTR measures the percentage of effective clicks among all clicks:

Table 1
List of notations.

Notation	Definition
B	The overall budget for search advertisements
n_1	The number of search markets
x_i	The allocated budget for the i th search market
n_2	The number of temporal slots/days during a promotional period
y_{ij}	The allocated budget for the j th temporal slot/day in the i th search market
t	Current time (within a temporal slot/day)
c	Cost-Per-Click (CPC)
ρ	Effective CTR
b	Remaining budget (within the temporal slot/day)
\mathbf{m}	State of environment, and $\mathbf{m} = (c, \rho, b)$
\mathbf{X}	A set of environment states, $\mathbf{X} = (\mathbf{m}_1, \mathbf{m}_2, \dots)$
N	The number of adjustments during each temporal slot/day allowable by search engine
d	(Claimed) daily budgets
β	Bid price
\mathbf{u}	An action, $\mathbf{u} = (d, \beta)$
\mathbf{U}	A set of actions, $\mathbf{U} = \{\mathbf{u}_1, \mathbf{u}_2, \dots\}$
μ	A policy, which is a map from state \mathbf{m} to action \mathbf{u}

$$\text{Effective CTR} = \frac{\text{effective clicks}}{\text{total clicks}} \tag{1}$$

(Note that the effective CTR is equivalent to the conversion rate if these kinds of user behaviors are defined as conversion actions by advertisers.) CPC and effective CTR determine the cost and reward for the advertiser on the sponsored search campaigns. Given the constraints of search engines, the advertiser can only change daily budget N time a day. In this problem, an advertiser aims to find the adjustment policy that will bring the highest reward.

4. The dual adjustment of budgets and bids within a day

The strategic decisions on advertising budgets and bids can be viewed as a special multi-stage dynamic decision problem with Markov properties [1]. Decisions at time t depend on both the current marketing state and decisions at time $t-1$. This makes Reinforcement Learning (RL) an appropriate technique to model budget and bid adjustments in

sponsored search auctions. Note that sponsored search auctions are a continuous process due to the high volume of search queries, which demand a CRL approach with flexible components suitable to encode the various states and actions in sponsored search auctions. In a previous study [35], a BOF framework was built that incorporates interactions between the different levels of decisions in advertising management. This research extends this framework to the dual adjustment setting. We first describe the formulas for the first two levels of decisions and then discuss our extension of the BOF framework on real-time dual adjustment. For the sake of simplicity, we do not differentiate the keywords contained in each campaign.

The budget for each of the n_1 search engines is determined by maximizing the total reward, as in:

$$\begin{aligned} \max \quad & \sum_{i=1}^{n_1} g_i^{(1)}(x_i) \\ \text{s.t.} \quad & \sum_{i=1}^{n_1} x_i - B \leq 0 \\ & x_i \geq 0, i = 1, 2, \dots, n_1, \end{aligned} \quad (2)$$

where $g_i^{(1)}$ is the maximized campaign reward in the i th search engine, given as:

$$\begin{aligned} g_i^{(1)}(x_i) := \max \quad & \sum_{j=1}^{n_2} g_{i,j}^{(2)}(y_{i,j}) \\ \text{s.t.} \quad & \sum_{j=1}^{n_2} y_{i,j} - x_i \leq 0 \\ & y_{i,j} \geq 0, j = 1, 2, \dots, n_2, \end{aligned} \quad (3)$$

where $g_{i,j}^{(2)}(y_{i,j})$ is the maximized campaign reward in the j th temporal slot of the i th search engine.

Within the scope of a temporal slot (i.e., a day), the adjustment advertising strategy depends on the dynamics of the market, namely, the CPC c , the effective CTR ρ , and the remaining budget of the day b . We thus take them as state variables in $\mathbf{m} = (c, \rho, b)$ and compose a set of environment states $\mathbf{X} = \{\mathbf{m}_1, \mathbf{m}_2, \dots\}$ to represent the sponsored search auction environment. Furthermore, an advertiser can only control the daily budget d and bid price β . Thus, we take them as action variables in $\mathbf{u} = (d, \beta)$ and compose a set of actions $U = \{\mathbf{u}_1, \mathbf{u}_2, \dots\}$. A policy μ is a mapping from the set of states to the set of actions, which are optimal advertising actions for each corresponding state.

The environment state in sponsored search auctions varies with time. Thus we represent it with differential equations as:

$$\begin{aligned} \dot{c}(t) &= f_1(\mathbf{m}(t), \mathbf{u}(t)) \\ \dot{\rho}(t) &= f_2(\mathbf{m}(t), \mathbf{u}(t)) \\ \dot{b}(t) &= f_3(\mathbf{m}(t), \mathbf{u}(t)) \end{aligned}$$

which depict the changing rate of the environment state. Note that we don't have access to $\mathbf{f} = (f_1, f_2, f_3)$ in advance. However it's possible to estimate the variation tendency of \mathbf{f} from historical data.

Assume action $\mathbf{u}(t)$ is taken at time t , when the system is in state $\mathbf{m}(t)$. After a short temporal period Δt , the system state transits to $\mathbf{m}(t + \Delta t)$. Within this period, the day's remaining budget reduces from $b(t)$ to $b(t + \Delta t)$. Thus the cost from t to $t + \Delta t$ is $b(t) - b(t + \Delta t)$. Note that, b is non-increasing as fees are deducted by the search engine. Given this cost value, the advertiser gets $\frac{b(t) - b(t + \Delta t)}{c(t)}$ clicks, among which $\rho(t) \frac{b(t) - b(t + \Delta t)}{c(t)}$ are effective clicks. We use $r(t) = r(\mathbf{m}(t), \mathbf{u}(t))$ to denote these effective clicks, which is a measure of the instant reward at time t . Based on it, we can derive the cumulative reward from time t to the end of temporal slot, T , as reward function $V(t)$:

$$V(t) = \int_t^T e^{-\frac{s-t}{\tau}} r(\mathbf{m}(s), \mathbf{u}(s)) ds,$$

in which we apply discount factor $\tau \in (0, 1]$ to reflect the time value of money.

In practice, the number of allowed budget adjustments is limited by search engine auction providers. In this research, we assume that an

advertiser adjusts the daily budget and bid prices simultaneously, which are both N times during a day. Thus, \mathbf{u} can be represented as a step function with N pieces:

$$\mathbf{u}(t) = \begin{cases} \mathbf{u}_1, & \text{if } 0 \leq t < t_1 \\ \mathbf{u}_2, & \text{if } t_1 \leq t < t_2 \\ \dots & \\ \mathbf{u}_N, & \text{if } t_{N-1} \leq t < T. \end{cases}$$

The advertiser thus needs to make N decisions. At time t , an advertiser needs to choose the optimal control variable $u(t)$ to maximize $V(t)$.

With the above derived state variables and cost/reward values, we formulate the budget and bids adjustment problem within a day as:

$$\begin{aligned} g_{i,j}^{(2)}(y_{i,j}) := \max_{\mathbf{u} \in U} \quad & V \quad (0) \\ \text{s.t.} \quad & \dot{c}(t) = f_1(\mathbf{m}(t), \mathbf{u}(t)) \\ & \dot{\rho}(t) = f_2(\mathbf{m}(t), \mathbf{u}(t)) \\ & \dot{b}(t) = f_3(\mathbf{m}(t), \mathbf{u}(t)) \\ & V(t) = \int_t^T e^{-\frac{s-t}{\tau}} r(\mathbf{m}(s), \mathbf{u}(s)) ds \\ & c_0 = c_{i,j}, \rho_0 = \rho_{i,j}, b_0 = y_{i,j} \end{aligned} \quad (4)$$

where the advertiser's purpose is to find out the optimal action $\mathbf{u}(t)^*$ characterizing a policy of budget and bids adjustments with the maximized profits in terms of effective clicks. Models (2)–(4) reflect advertisers' three levels of decisions during the lifecycle of a campaign.

Model (4) has two unique features that are different classic CRL models due to the characteristics of sponsored search auctions. First, due to the complexity of sponsored search auctions, it is impossible to get an explicit dynamic system in the CRL, i.e., \mathbf{f} is not a determined function given in advance. We need to train it from historical data. Secondly, constraints on allowed budget changes make the decision space discontinuous, which hinders the derivation of optimal actions. Here we design \mathbf{u} as a step function with limited steps. The changes of the step function can happen at any time during a day though. These features raise difficulties in reaching a closed form optimal solution for the model.

5. Numerical solution

To solve our proposed dual adjustment model, in this section, we discuss a numeric solution algorithm. Note that since an advertising schedule is time-bounded, our numerical approach got to be efficient.

5.1. Estimation of function \mathbf{f}

In order to estimate the function of $\mathbf{f}(\mathbf{m}, \mathbf{u})$, we use a back-propagation neural network as the universal approximation $\bar{\mathbf{f}}(\mathbf{m}, \mathbf{u}; \nu)$, where ν is a parameter of the function approximator, and use a reward-based mechanism to improve the approximation. We initialize the parameters of $\bar{\mathbf{f}}$ based on the domain knowledge, then at time t the parameters can be improved based on the difference of system states from the function estimation by minimizing

$$E_1(t) = \frac{1}{2} \left\| \dot{\mathbf{m}}(t) - \bar{\mathbf{f}}(\mathbf{m}(t), \mathbf{u}; \nu) \right\|^2. \quad (5)$$

After iterations, the approximation function $\bar{\mathbf{f}}$ will approach to the exact \mathbf{f} , following the Cybenko Theorem [6,31].

Theorem 1. Cybenko Theorem

Given any $\mathbf{f} \in \mathbf{C}^1(\mathbb{R}^n, \mathbb{R}^m)$, and any $\epsilon > 0$, we can always find a neural network $\bar{\mathbf{f}}(\mathbf{x}; \nu)$ that approximates $\mathbf{f}(\mathbf{x})$, i.e.

$$\max_{\mathbf{x}} |\bar{\mathbf{f}}(\mathbf{x}; \nu) - \mathbf{f}(\mathbf{x})| \leq \epsilon. \quad (6)$$

5.2. Finding the optimal action \mathbf{u}^*

In sponsored search auctions, the control variable \mathbf{u} can only be a step function with limited number of adjustments, which makes it difficult to find the optimal actions. Given the N allowed changes, the advertiser faces N decisions. At time t , an advertiser needs to choose the optimal control variable $u(t)$, considering the remaining adjustment times n and the system status vector \mathbf{m} , to maximize $V(t)$. If an advertiser has a different number of remaining decision points, she may have different choices and expected payoffs. Here we denote V_n as the value with n remaining adjustment chances and propose a recursion method to find the optimal action series.

According to Bellman's Principle of Optimality [2], the optimal payoff for Model (4) is:

$$V^*(\mathbf{m}(t)) = \max_{\mathbf{u}(t+\Delta t)} \left[\int_t^{t+\Delta t} e^{-\frac{s-t}{\tau}} r(\mathbf{m}(t), \mathbf{u}(t)) ds + e^{-\frac{\Delta t}{\tau}} V^*(\mathbf{m}(t+\Delta t)) \right]. \quad (7)$$

Let $\Delta t \rightarrow 0$, we have:

$$V^*(\mathbf{m}(t)) = \tau \max_{\mathbf{u}(t)} \left[r(\mathbf{m}(t), \mathbf{u}(t)) + \frac{\partial V^*}{\partial \mathbf{m}} \mathbf{f}(\mathbf{m}(t), \mathbf{u}(t)) \right]. \quad (8)$$

For the decision related to V_n , considering the actions in two subsequent temporal points t and $t-$, we have

If $\mathbf{u}(t) = \mathbf{u}(t-)$:

$$V_n^*(\mathbf{m}(t)) = \tau \left[r(\mathbf{m}(t), \mathbf{u}(t-)) + \frac{\partial V_n^*}{\partial \mathbf{m}} \mathbf{f}(\mathbf{m}(t), \mathbf{u}(t-)) \right]. \quad (9)$$

If $\mathbf{u}(t) \neq \mathbf{u}(t-)$:

$$V_n^*(\mathbf{m}(t)) = \tau \max_{\mathbf{u}(t)} \left[r(\mathbf{m}(t), \mathbf{u}(t)) + \frac{\partial V_{n-1}^*}{\partial \mathbf{m}} \mathbf{f}(\mathbf{m}(t), \mathbf{u}(t)) \right]. \quad (10)$$

Hence, for the last decision point, $n = 0$:

$$V_0^*(\mathbf{m}(t)) = \tau \left[r(\mathbf{m}(t), \mathbf{u}(t-)) + \frac{\partial V_0^*}{\partial \mathbf{m}} \mathbf{f}(\mathbf{m}(t), \mathbf{u}(t-)) \right], \quad (11)$$

and for earlier decision points with $n > 0$:

$$V_n^*(\mathbf{m}(t)) = \tau \max \left\{ \max_{\mathbf{u}(t)} \left[r(\mathbf{m}(t), \mathbf{u}(t)) + \frac{\partial V_{n-1}^*}{\partial \mathbf{m}} \mathbf{f}(\mathbf{u}(t)) \right], \right. \\ \left. \left[r(\mathbf{m}(t), \mathbf{u}(t-)) + \frac{\partial V_n^*}{\partial \mathbf{m}} \mathbf{f}(\mathbf{m}(t), \mathbf{u}(t-)) \right] \right\}. \quad (12)$$

For the scenario where $n > 0$, if we define 'action indicator' h as

$$h = \max \left[r(\mathbf{m}(t), \mathbf{u}(t)) + \frac{\partial V_{n-1}^*}{\partial \mathbf{m}} \mathbf{f}(\mathbf{m}(t), \mathbf{u}(t)) \right] \\ - \left(r(\mathbf{m}(t), \mathbf{u}(t-)) + \frac{\partial V_n^*}{\partial \mathbf{m}} \mathbf{f}(\mathbf{m}(t), \mathbf{u}(t-)) \right), \quad (13)$$

then we get:

$$\mathbf{u}^*(t) = \operatorname{argmax}_{\mathbf{u}(t)} \left\{ r(\mathbf{m}(t), \mathbf{u}(t)) + \frac{\partial V_{n-1}^*}{\partial \mathbf{m}} \mathbf{f}(\mathbf{m}(t), \mathbf{u}(t)) \right\}, \text{ if } h > 0, \quad (14)$$

$$\mathbf{u}^*(t) = \mathbf{u}^*(t-), \text{ if } h < 0. \quad (15)$$

Let $V_n(\mathbf{m}; \mathbf{w}_n)$ be the universal approximation of the total discounted reward function V_n^* , where \mathbf{w}_n is the parameter vector generated from the back-propagation neural network. We can compute the mean

squared TD error, which should be minimized during the reinforcement learning process:

$$E_2^0(t) = \frac{1}{2} \left| r(\mathbf{m}(t), \mathbf{u}(t)) + \frac{\partial V_0^*}{\partial \mathbf{m}} \mathbf{f}(\mathbf{m}(t), \mathbf{u}(t)) - V_0(\mathbf{m}(t)) \right|^2, \quad (16)$$

$$E_2^n(t) = \begin{cases} \frac{1}{2} \left| r(\mathbf{m}(t), \mathbf{u}(t)) + \frac{\partial V_{n-1}^*}{\partial \mathbf{m}} \mathbf{f}(\mathbf{m}(t), \mathbf{u}(t)) - V_{n-1}(\mathbf{m}(t)) \right|^2, & \text{if } d \leq 0, \\ \frac{1}{2} \left| r(\mathbf{m}(t), \mathbf{u}(t)) + \frac{\partial V_n^*}{\partial \mathbf{m}} \mathbf{f}(\mathbf{m}(t), \mathbf{u}(t)) - V_n(\mathbf{m}(t)) \right|^2, & \text{otherwise.} \end{cases} \quad (17)$$

5.3. The solution algorithm

Based on the above discussion, we propose a solution algorithm to find the optimal action \mathbf{u}^* of the model. Algorithm 5.1 aims to iteratively find a near-optimal action at time t . In each round, we also updated the BPNN's estimation of function f and minimize the TD error for reinforcement learning. The computing procedure is illustrated in Fig. 2. The convergence of this algorithm is guaranteed by Theorem 1.

Algorithm 5.1.

```

1: procedure REALTIME ADJUSTMENT( $\mathbf{m}$ ,  $r$ )
2:    $n = N$ ,  $t = 0$ 
3:   repeat
4:      $\nu \leftarrow \operatorname{argmin} E_1(t)$  (Eq. 5)
5:     for  $k \leftarrow 0, n$  do
6:        $\omega_k \leftarrow \operatorname{argmin} E_2^k$  (Eq. 16 and 17)
7:     end for
8:     if  $n = 0$  then
9:        $u(t) = u(t-)$ 
10:    else
11:       $h \leftarrow \max[r(\mathbf{m}(t), \mathbf{u}(t)) + \frac{\partial V_{n-1}^*}{\partial \mathbf{m}} \mathbf{f}(\mathbf{m}(t), \mathbf{u}(t))$ 
12:         $- (r(\mathbf{m}(t), \mathbf{u}(t-)) + \frac{\partial V_n^*}{\partial \mathbf{m}} \mathbf{f}(\mathbf{m}(t), \mathbf{u}(t-)))]$ 
13:      if  $h > 0$  then
14:         $u(t) \leftarrow u^*(t)$  (Eq. 14)
15:       $n \leftarrow n - 1$ 
16:    else
17:       $u(t) \leftarrow u(t-)$ 
18:    end if
19:  end if
20:  Read  $t$  (from the outside system).
21:  until  $t = T$ 
22: end procedure

```

Next, we discuss the computational load of Algorithm 5.1. The algorithm can be divided into two cases. If $n = 0$, then it simply records the current state generated by the system and updates parameters of the two neural networks. The computational complexity for this step is polynomial. If $n > 0$, then it needs to compute 2 local optima using the

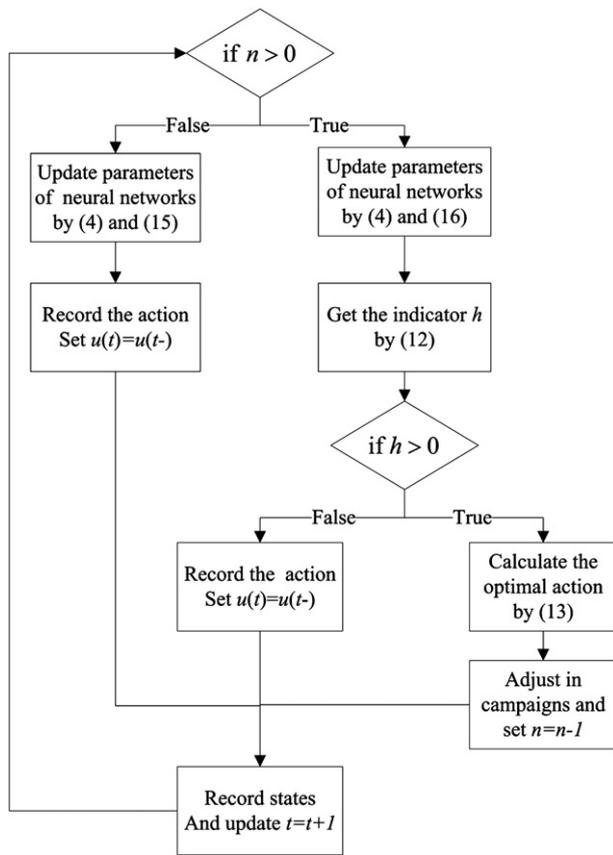


Fig. 2. The flow chart of computing procedure.

interior point method in addition to updating the two neural networks. The computational complexity for this step is also polynomial. Note that if n is large, this algorithm might need a large memory to store the neural networks. Nevertheless, major search engines (e.g., Google) restrict the number of changes on daily budget, which eliminates this concern. In this sense, Algorithm 5.1 is an efficient solution for the dual adjustment problem for online deployment.

6. Experimental validation

In this section, we conduct two sets of experiments to evaluate the proposed model and optimization algorithm. First, we assess the performance of our dual adjustment strategy by comparing with two baseline strategies. Second, we assess the necessity of adjusting both daily budgets and bids by comparing with strategies adjusting only one of them.

6.1. Dataset

We collected a dataset about campaigns conducted by an e-commerce advertiser on one of the largest search engines in China for the experiment. The dataset includes field reports and logs of search advertising campaigns from September 2008 to August 2010. The dataset contains 4 campaigns covering 170 keywords. In the data collection period, 476,832 sponsored links to the e-commerce website were triggered.

6.2. Evaluation metrics

In this research, we employ the reward advertisers received from the campaigns to evaluate the effectiveness of optimization strategies.

Search engine users may click on sponsored links and initiate a visit to the target website. The advertiser will be charged for all clicks going through her sponsored links. (In practice, advertisers generally would use up all their budget on the received clicks.) However, among these clicks, only the *effective clicks* bring the advertiser a reward. In this research, we consider that a click generates value if it causes purchase, registration, staying on (i.e., viewing) the landing page for more than 5 s, surfing the target website for more than 2 links, or bookmarking/downloading of relevant pages. (In our dataset, 1402 clicks were counted as effective clicks.) We assume that the reward associated with each effective click is the same. Thus, the reward can be represented by the number of effective clicks received. A higher reward means that the advertiser can earn more money from investing the same amount in a campaign.

6.3. Baseline methods

The literature review uncovered only a few studies on dynamic adjustment of daily budgets and bids in sponsored search auctions. For comparison purposes, we implement two baseline strategies. The first baseline strategy (denoted as Fixed-strategy) fixes daily budgets and bids during the scope of a campaign. It is a simple but widely used method when advertisers don't have sufficient resources and knowledge to keep track of advertising performance. The second baseline strategy (denoted as TwoBid-strategy) randomizes bid prices between two values q_1 and q_2 on all keywords in the campaign until the actual daily budget is exhausted [13].

Moreover, we tweak our dual adjustment strategy (denoted as DynAdjustment) to derive two other baseline methods to evaluate the effect of adjusting only daily budgets or bids in the campaigns. Specifically, the Budget-Adjustment strategy is a special case of DynAdjustment that fixes the bid as a constant (i.e., 2.0) and employs the daily budget as the only control factor. The Bid-Adjustment strategy is a special case of DynAdjustment that fixes the daily budget as a constant (i.e., 60.0) and employs the bid as the only control factor.

6.4. Experimental procedure

In practice, advertisers (especially for small enterprises) don't have sufficient budget, so it is necessary to deal with the budget constraint. In the experiments, we assume that the (insufficient) daily budget limit is $B = 100$. Thus, the advertisers' links can only be shown among the sponsored links during some temporal periods of a day. By conducting simulation on adjustment strategies and comparing with the actual clicks and effective clicks in the dataset, we can calculate the reward for each strategy.

In the experiments, the parameters of our dual adjustment model need to be specified based on the data. Particularly, we obtain the average values of cost-per-click (CPC), daily budget limits, bid prices, and total clicks directly from reports provided by search engines. We extract the dynamics of cost-per-click (CPC) from reports of the campaigns, and the dynamics of effective click-through-rate (CTR) and clicks from log files on the advertisers' websites. We extract effective clicks from weblogs using regular expression matching and calculate corresponding effective CTR according to Eq. (1). Fig. 3 illustrates changes of CPC and effective CTR over time. This information is used to initialize and train function f by following formula (5). The baselines on adjusting only daily budgets or bids can be trained in a similar fashion. The Fixed-strategy and TwoBid-strategy do not require parameter tuning.

Since our model is based on strategy adjustments within a day, we mainly compare the dynamics of the strategies' performances within a one-day scope. We also provide a monthly performance curve to illustrate how our strategy may affect the advertising performance at a coarser temporal granularity level. For each of the experiments, we conduct simulation 30 times and report their averaged performance.

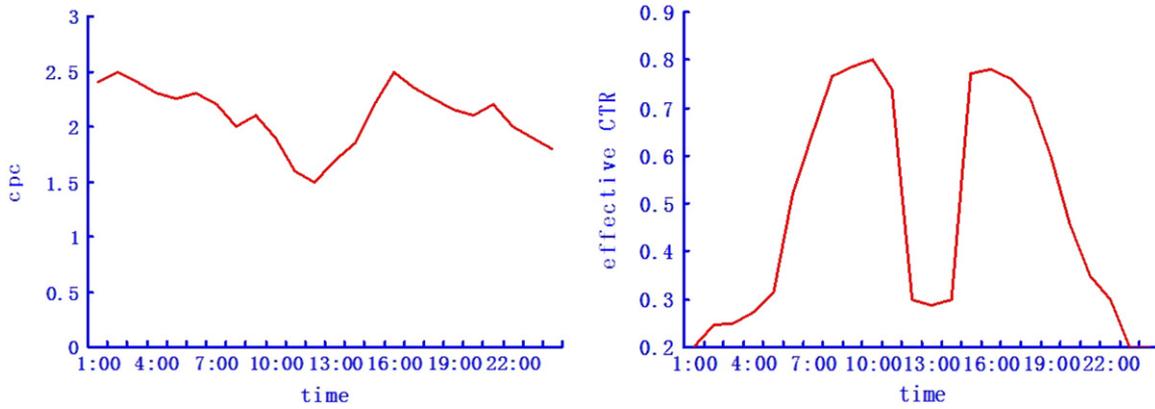


Fig. 3. CPC and effective CTR during a day.

6.5. Results

6.5.1. Performance as compared with existing strategies

The first experiment is intended to evaluate the performance of our dual adjustment strategy by comparing it with two baseline strategies.

We first assess the results of various strategies within a typical day. Fig. 4 reports the change of a) daily budgets, b) bid prices, and c) the day's remaining budget (caused by changes of bids and budgets) after implementing our method and two baseline methods. As we can see from the figure, the Fixed-strategy also takes both a fixed bid price and a fixed budget on the entire day. The TwoBid-strategy takes a fixed budget but tweaks the bid price a little bit. As a result, both methods' remaining budgets reduce from the very beginning of the day and are used up by about 3 pm. Our dual adjustment strategy, however, shows a quite different behavior. In general, the strategy gradually increases the budget limit along with time, so that the budget won't be used up at the beginning of the day. The strategy also increases the bid prices at two peak hours of the day. As we can see in Fig. 3, the effective CTR is relatively higher during these times. The increased price thus attracts more clicks and brings more reward. The remaining budget, under our strategy, starts to decrease at about 6 am and lasts until 9 pm.

Fig. 5 reports the rewards accumulated by executing these three strategies during a day. As a direct result of manipulating bids and budgets to the temporal slots with higher effective CTRs, our dual adjustment strategy obtains the highest reward (19,681 effective clicks) among the three. The TwoBid-strategy receives 16,000 effective clicks and the Fixed-strategy receives 14,634 effective clicks. Within a day, our strategy outperforms the two baseline strategies by about 23% and 34.5%, respectively. We also notice that the accumulation of rewards starts later with our strategy. However, it experiences a much faster increase from 7 am to 12 pm and from 2 pm to 9 pm, which leads to its high reward.

Fig. 6 further reports experiment results for the three strategies over a one-month temporal period (where daily traffic varies). Since each strategy generally receives the same amount of rewards per day, the increase of the total reward is roughly linear. In a longer temporal scope, the TwoBid-strategy performances are relatively better than before, which is still worse than our proposed strategy. Within a month, our strategy obtains 607,926 effective clicks. The TwoBid-strategy receives 550,064 effective clicks. The Fixed-strategy receives 520,833 effective clicks. Our strategy outperforms the two baseline strategies by 10.5% and 16.7%, respectively.

6.5.2. Effects of adjusting budget and bid in our strategy

The second experiment explores which part of our dual adjustment strategy is more important, daily budget adjustment or bid adjustment. For comparison purposes, the figures include our dual adjustment strategies' performances, which are actually the same as reported in the previous subsection.

Fig. 7 reports the change of a) daily budgets, b) bid prices, and their caused c) remaining budgets within a day by adjusting only daily budget or bid prices. When we take a Bid-Adjustment strategy, the daily budget remains the same while bid price varies (peaks from 6 am to 12 pm) during the day. Using the Budget-adjustment strategy, the bid price remains the same while daily budget varies (gradually increases from 7 am to 4 pm) during the day. Due to the model's learning abilities, none of the strategies uses their budget at the beginning of a day. The Bid-adjustment strategy starts using budget at about 5 am and the Budget-adjustment strategy starts 2 h later. Note that both strategies spend their budget faster than our dual-adjustment strategy.

As a result of their abilities to tweak bids and budgets, the two strategies receive different rewards. Fig. 8 shows that the budget adjustment strategy and the bid adjustment strategy have steeper reward increases but less total rewards. The Bid adjustment strategy has a quick reward increase after 5 o'clock and ends up with 16,759 effective clicks. The budget adjustment strategy shows a two-stage rapid increase and ends up with 17,951 effective clicks.

When viewing the three strategies in the scope of one month, their performance differences are even clearer (Fig. 9). Among the three strategies, the bid adjustment strategy has the lowest performance with 516,667 effective clicks per month. However, the budget adjustment strategy has a performance (591,667 effective clicks per month) that is much closer to the dual adjustment strategy. Obviously, the budget adjustment part plays a more important role in our dual adjustment strategy in directing budgets to the most appropriate temporal slots for effective clicks.

7. Discussion and managerial insights

From the experiments, we can see that our proposed dual adjustment strategy is a feasible solution to the advertising strategy optimization problem in sponsored search auctions. It outperforms two state-of-the-art baseline strategies in terms of daily level rewards and monthly level rewards. From the behavior of this strategy in the experiment, we notice that it prevents the budget from being exhausted too early in the day, reserving it for later intervals with higher possible rewards. The two baseline strategies, however, tend to consume the budget at a relatively stable rate from the very beginning of each day. In general, all strategies took advantage of all chances provided by the search engine to tweak budgets, which means that our proposed optimization strategy is effective in capturing market dynamics and altering advertising strategies in a timely manner.

The experiments also illustrate the effects of adjusting daily budget vs. adjusting bid price. The effect of budget adjustment is slightly stronger than the effect of bid adjustment on the accumulated reward. This phenomenon provides another lens on the keyword bidding problem currently under investigation by researchers: instead of solely relying on change of bid prices, it may be more effective (and less

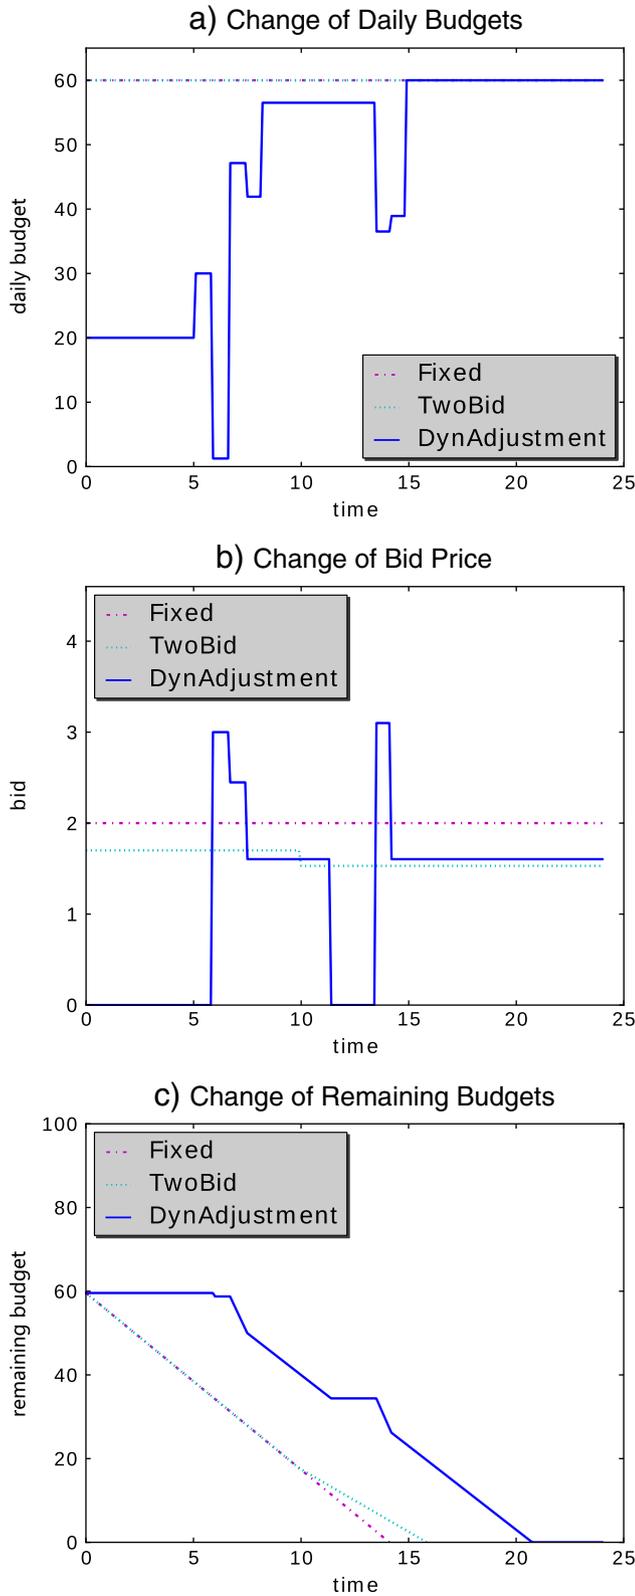


Fig. 4. The variations of daily budgets and bid prices of three strategies and their caused remaining budget during a day.

computationally intensive) to inspect the alteration of daily budget as a control variable in designing advertising strategies.

In addition to the algorithm contributions, this research provides important managerial insights for advertisers in sponsored search

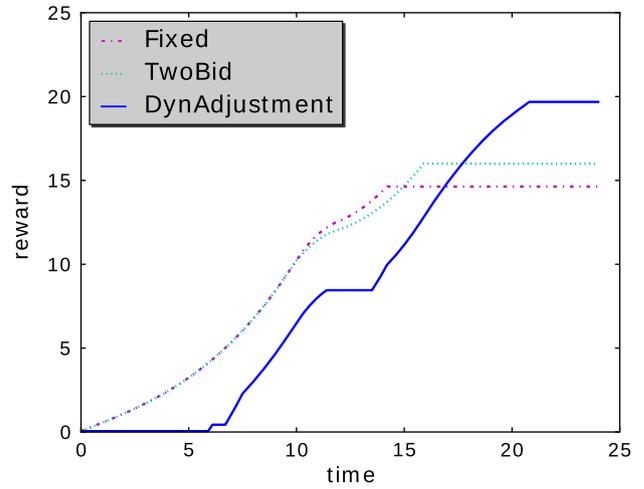


Fig. 5. The change of the rewards during a day.

auctions. First, advertisers should not solely conduct long-term and medium-term budget optimization and leave budget and bid price fixed within a day. In response to the inherent dynamic nature of sponsored search auctions, advertisers should carefully track the advertising performance, and adjust the daily budget and bids correspondingly. Our work indicates that a simple strategy of daily budget adjustment can, to some degree, improve the advertising performance in terms of effective clicks. For example, a manager should tune down budget when the traffic is high but the effective CTR is low. Second, most advertisers pay more attention to bidding strategies in order to maximize advertising performance. In this research, we provide opportunities for advertisers to adjust both the daily budget and bid prices for management. The improved performance of the dual-adjustment method indicates the interaction of the two factors and illustrates the possibilities of including other types of factors to control advertising. It is worthwhile for advertisers to empirically explore such opportunities and for researchers to investigate such problems. Third, since more changes of budgets/bids enable better tracking and response to market status, advertisers should take advantage of all chances provided by search engine to reinspect their advertising strategy, especially in competitive markets and during hours with rapid change of effective CTRs. In our experiments, our proposed approach used up all change opportunities allowed by search engines.

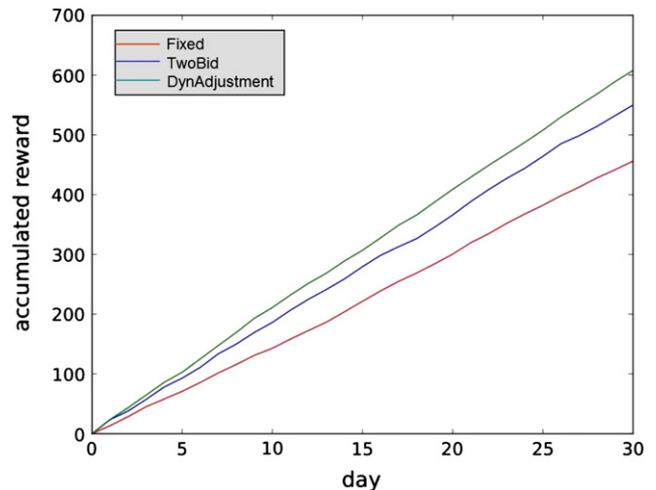


Fig. 6. The rewards of these strategies during a month.

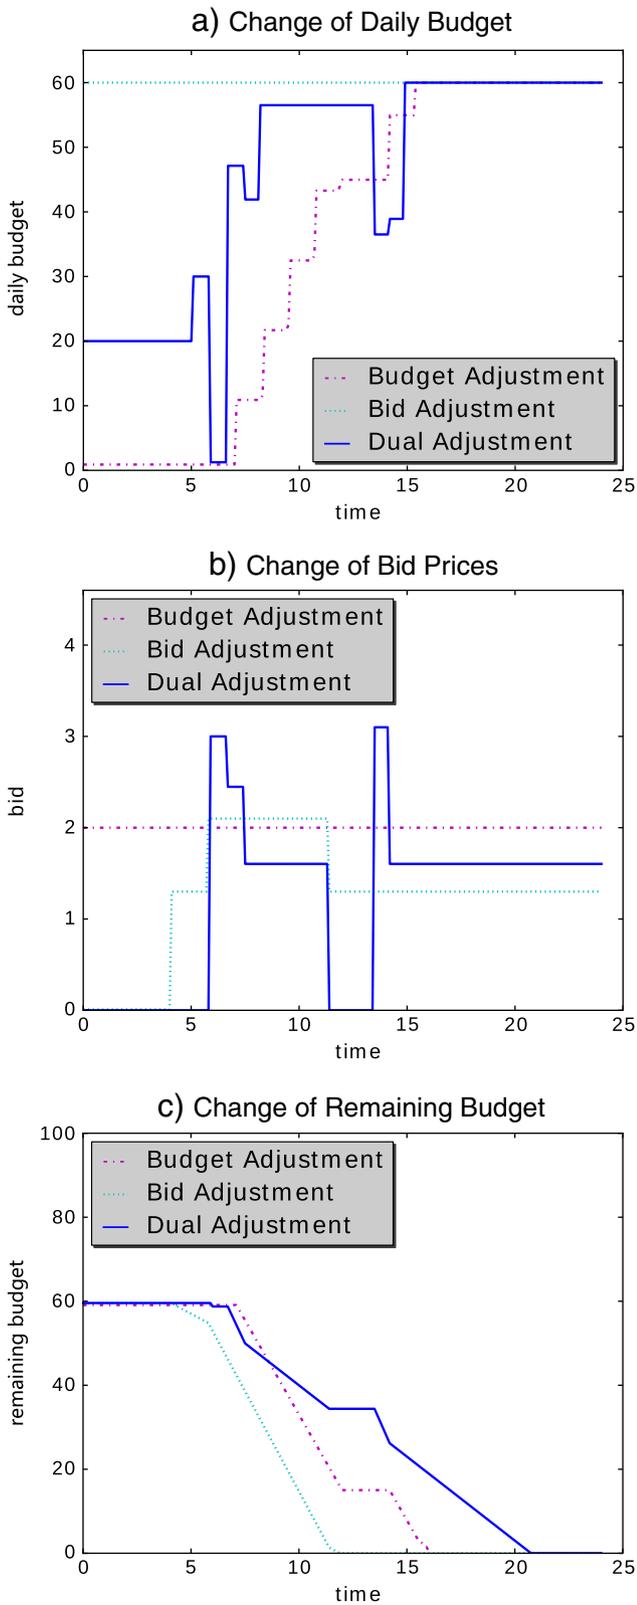


Fig. 7. Effects of the daily budget and bids adjustment on the remaining budget for these strategies during a day.

8. Conclusions and future work

In this paper, we propose a reinforcement learning approach to deal with the problem of adjusting daily budget and bid price for better

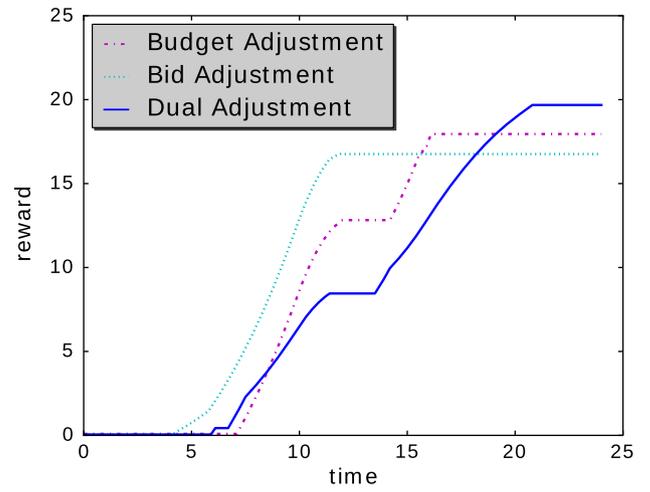


Fig. 8. Effects of the daily budget and bids adjustment on the rewards of these strategies during a day.

advertising output. Since the sponsored search auction market changes rapidly, we propose a continuous-time, continuous-state model for dual adjustment of daily budget and bid price. In addition, since budgets can only be changed a limited times during a day, we incorporate discrete-time-action with step functions in the model, where the actions can be taken at any time of a day. We also provide a numerical solution to our model, considering the difficulty of deriving close form solutions. We conducted some experiments to validate our proposed approach with the real-world data about campaigns conducted by an e-commerce advertiser on a major Chinese search engine. Experimental results illustrate that our strategy outperforms the two baseline strategies. It is also found that the daily budget adjustment plays a more important role than bid adjustment in affecting advertising effectiveness in our model.

This work opens a new venue for advertising strategy development. In the future we intend to explore (a) the theoretical basis and empirical evidence for co-optimization of daily budget and bids; (b) the inter-operations of dynamic budget adjustment across search markets; and (c) more efficient computational algorithms to facilitate online applications of advertising strategy optimization.

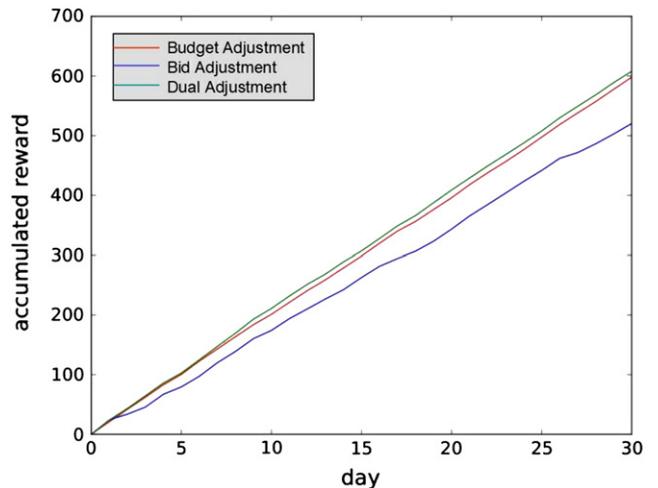


Fig. 9. Effects of the daily budget and bids adjustment in a month.

Acknowledgments

We are thankful to anonymous reviewers who provided valuable suggestions that led to a considerable improvement in the organization and presentation of this manuscript. This work is partially supported by NSFC (71071152, 71272236, 71332001, 71202169, 71025001, 71231002) and Beijing Natural Science Foundation (4132072).

References

- [1] N. Archak, V.S. Mirrokni, S. Muthukrishnan, Budget optimization for online advertising campaigns with carryover effects, The Eleventh ACM SIGECOM International Conference on Electronic Commerce, Harvard, 2010.
- [2] R.E. Bellman, *Dynamic Programming*, Princeton University Press, Princeton, NJ, 1957.
- [3] J.P. Benoit, V. Krishna, Multiple-object auctions with budget constrained bidders, *Review of Economic Studies* 68 (1) (2001) 155–179.
- [4] M. Babaioff, N. Immorlica, D. Kempe, R. Kleinberg, A knapsack secretary problem with applications, *Lecture Notes in Computer Science* 4627 (2007) 16–28.
- [5] D. Chakrabarty, Y. Zhou, R. Lukose, Budget constrained bidding in keyword auctions and online knapsack problems, *Proceedings of the 16th International World Wide Web Conference*, 2007.
- [6] G. Cybenko, Approximations by superpositions of sigmoidal functions, *Mathematics of Control, Signals, and Systems* 2 (4) (1989) 303–314.
- [7] K. Doya, Reinforcement learning in continuous time and space, *Neural Computation* 12 (2000) 215–245.
- [8] R. Du, Q. Hu, S. Ai, Stochastic optimal budget decision for advertising considering uncertain sales responses, *European Journal of Operational Research* 183 (2007) 1042–1054.
- [9] B. DasFuata, S. Muthukrishnan, Stochastic budget optimization in Internet advertising, *Algorithmica* (2010) 1–28.
- [10] F.M. Feinberg, On continuous-time optimal advertising under S-shaped response, *Management Science* 47 (2001) 1476–1487.
- [11] G.E. Fruchter, W. Dou, Optimal budget allocation over time for keyword ads in web portals, *Journal of Optimization Theory and Applications* 124 (1) (2005) 157–174.
- [12] G. Feichtinger, R.F. Hartl, S.P. Sethi, Dynamic optimal control models in advertising: recent developments, *Management Science* (1994) 195–226.
- [13] J. Feldman, S. Muthukrishnan, et al., Budget optimization in search-based advertising auctions, *Proceedings of the 8th ACM conference on Electronic Commerce*, ACM Press, California USA, 2007.
- [14] J.K. Johansson, Advertising and the s-curve: a new approach, *Journal of Marketing Research* (1979) 346–354.
- [15] E.J. Johnson, W.W. Moe, P.S. Fader, S. Bellman, G.L. Lohse, On the depth and dynamics of online search behavior, *Management Science* 50 (3) (March 2004) 299–308.
- [16] B.J. Jansen, A. Spink, Sponsored search: is money a motivator for providing relevant results? *IEEE Computer* 40 (8) (2007) 50–55.
- [17] T.V. Krishnan, D.C. Jain, Optimal dynamic advertising policy for new products, *Management Science* 52 (12) (2006) 1957–1969.
- [18] B. Kitts, B.J. LeBlanc, A trading agent and simulator for keyword auctions, *Proceedings of the Third International Joint Conference on Autonomous Agents & Multi Agent Systems*, IEEE Computer Society, Washington, DC, USA, 2004, pp. 228–235, (New York).
- [19] B. Mueller, *Dynamics of International Advertising: Theoretical and Practical Perspectives*, Peter Lang, 2010.
- [20] H.I. Mesak, A.F. Darrat, On comparing alternative advertising policies of pulsation, *Decision Sciences* 32 (1992) 541–564.
- [21] V. Mahajan, E. Muller, Advertising pulsing policies for generating awareness for new products, *Marketing Science* 5 (1986) 89–111.
- [22] S. Muthukrishnan, M. Pál, Z. Svitkina, Stochastic models for budget optimization in search-based advertising, *Algorithmica* 58 (2010) 1022–1044.
- [23] M. Nerlove, K.J. Arrow, Optimal advertising policy under dynamic conditions, *Economica* 29 (114) (May, 1962) 129–142.
- [24] O. Ozluk, S. Cholette, Allocating expenditures across keywords in search advertising, *Journal of Revenue and Pricing Management* 6 (4) (2007) 347–356.
- [25] O.J. Rutz, R.E. Bucklin, From generic to branded: a model of spillover dynamics in paid search advertising, *Journal of Marketing Research* 48 (1) (2011) 87–102.
- [26] B. Richards, J. Botterill, I. MacRury, *The Dynamics of Advertising*, Routledge, 2000.
- [27] S.P. Sethi, Dynamic optimal control models in advertising: a survey, *SIAM Review* 19 (4) (Oct. 1977) 685–725.
- [28] S.P. Sethi, Optimal advertising for the Nerlove–Arrow model under a budget constraint, *Operations Research Quarterly* 28 (2) (1977) 683–693.
- [29] M.W. Sasieni, Optimal advertising strategies, *Marketing Science* 8 (1989) 358–370.
- [30] D. Simester, Y.J. Hu, E. Brynjolfsson, E.T. Anderson, Dynamics of retail advertising: evidence from a field experiment, *Economic Inquiry* 47 (3) (2009) 482–499.
- [31] D. Tikk, L.T. Kóczy, T.D. Gedeon, A survey on universal approximation and its limits in soft computing techniques, *International Journal of Approximate Reasoning* 33 (2) (2003) 185–202.
- [32] D. Vakratsas, F.M. Feinberg, F.M. Bass, G. Kalyanaram, The shape of advertising response functions revisited: a model of dynamic probabilistic thresholds, *Marketing Science* 23 (2004) 109–119.
- [33] H.L. Vidale, H.B. Wolfe, An operations–research study of sales response to advertising, *Operations Research* 5 (1957) 370–381.
- [34] Y. Yang, J. Zhang, B. Liu, D. Zeng, Optimal budget allocation across search advertising markets, *Proceedings of the 21st Workshop on Information Technologies and Systems*, 2011, p. 97102.
- [35] Y. Yang, J. Zhang, R. Qin, J. Li, F. Wang, W. Qi, A budget optimization framework for search advertisements across markets, *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans* 42 (5) (2011) 1141–1151.
- [36] Y. Yang, J. Zhang, R. Qin, J. Li, B. Liu, Z. Liu, Budget optimization strategies in uncertain environments of search auctions: a preliminary investigation, *IEEE Transactions on Services Computing* 6 (2) (2013) 168–176.
- [37] Y. Yang, R. Qin, B.J. Jansen, J. Zhang, D. Zeng, Budget planning for coupled campaigns in sponsored search auctions, *International Journal of Electronic Commerce* (2013) (Accepted for publication).
- [38] X. Zhang, J. Feng, Cyclical bid adjustments in search-engine advertising, *Management Science* 57 (9) (2011) 1703–1719.

Jie Zhang received his MSc. degree in operation research from Renmin University of China in 2009. He is with the State Key Lab of Management and Control for Complex Systems, Institute of Automation, Chinese Academy of Sciences. His current research interests include online auctions and optimal control. Contact him at zhangjie.isec@gmail.com.

Yanwu Yang received his PhD degree in computer science from the graduate school of Ecole Nationale Supérieure de Arts et Métiers (ENSAM), France in 2006. He is with School of Management, Huazhong University of Science and Technology (HUST). He is currently a professor and the head of the ISEC research group (Internet Sciences and Economic Computing). His current research interests include user modeling, web personalization and online auctions. Contact him at yangyanwu.isec@gmail.com.

Xin Li received his PhD degree in management information systems from the University of Arizona, Tucson in 2009. He is currently an Assistant Professor with the Department of Information Systems, City University of Hong Kong, Kowloon, Hong Kong. His research interests include business intelligence, social network analysis, data mining, and text mining. He is a member of the Association for Computing Machinery and Association for Information Systems. Contact him at xin.li@cityu.edu.hk.

Rui Qin received her MSc. degree in operation research from Hebei University, China in 2010. She is with the State Key Lab of Management and Control for Complex Systems, Institute of Automation, Chinese Academy of Sciences. Her current research interests include online auctions. Contact her at qinrui.isec@gmail.com.

Daniel Zeng received his PhD degrees in industrial administration from Carnegie Mellon University, Pittsburgh, PA in 1998. He is currently a Professor and a Honeywell Fellow with the Department of Management Information Systems, University of Arizona, Tucson, and a Research Professor with the State Key Lab of Management and Control for Complex Systems, Institute of Automation, Chinese Academy of Sciences. His research interests include software agents and their applications, security informatics, social computing, computational game theory, recommender systems, and spatiotemporal data analysis. Contact him at zeng@email.arizona.edu.