

# Adaptive Risk-Aware Bidding with Budget Constraint in Display Advertising

Zhimeng Jiang<sup>†</sup>, Kaixiong Zhou<sup>‡</sup>, Mi Zhang<sup>§</sup>, Rui Chen<sup>§</sup>, Xia Hu<sup>‡</sup>, Soo-Hyun Choi<sup>§</sup>

Department of Computer Science and Engineering, Texas A&M University, TX, USA

Department of Computer Science, Rice University, TX, USA

Samsung Electronics America

<sup>†</sup>zhimengj@tamu.edu, <sup>‡</sup>{Kaixiong.Zhou,xia.hu}@rice.edu,

<sup>§</sup>{mi.zhang,rui.chen1,sh9.choi}@samsung.com

## ABSTRACT

Real-time bidding (RTB) has become a major paradigm of display advertising. Each ad impression generated from a user visit is auctioned in real time, where demand-side platform (DSP) automatically provides bid price usually relying on the ad impression value estimation and the optimal bid price determination. However, the current bid strategy overlooks the randomness of the user behaviors (e.g., click) and the cost uncertainty caused by the auction competition. In this work, we propose a novel adaptive risk-aware bidding algorithm with budget constraint via reinforcement learning, which is the first to simultaneously consider estimation uncertainty and the dynamic risk tendency of a DSP. Specifically, we explicitly factor in the *uncertainty of estimated ad impression values* and model the *risk preference* of a DSP under a *specific state and market environment* via a sequential decision process. Additionally, we theoretically unveil the intrinsic relation between the uncertainty and the risk tendency based on value at risk (VaR). Consequently, we propose two instantiations to model risk tendency, including an expert knowledge-based formulation embracing three essential properties and an adaptive learning method based on self-supervised reinforcement learning. We conduct experiments on public datasets and show that the proposed framework achieves better performance in terms of the number of clicks under different budget constraints<sup>1</sup>.

## Keywords

Risk-Aware Bidding Strategy; Budget Constraint; Display Advertising

## 1. INTRODUCTION

In the past few years, real-time bidding (RTB) has quickly become tens of billions of markets in the globe [9; 34]. In RTB, a demand-side platform (DSP) buys ad impressions in a programmatic manner on behalf of advertisers. The success of a DSP heavily relies on its *bid optimization* (i.e., the process of identifying the optimal bid price for each bid request) capability [32], whose goal is to maximize the key performance indicator (KPI) agreed upon with advertisers

(e.g., the total number of clicks or return on ad spend). In practice, bid optimization normally involves two steps, namely user response prediction and bid price determination [30]. User response prediction is performed to estimate the true value of a potential ad impression. Taking the estimated value as an input, the bid price determination step aims to generate the optimal bid price for a bid request in sequential decision making.

RTB is a highly competitive and dynamic marketplace. The prerequisite for a DSP to stay competitive is the capability of accurately predicting user responses [4; 8], e.g., click-through rate (CTR) or conversion rate (CVR), and advanced bid price determination algorithm. A large number of prediction models [38; 27; 12; 33; 41; 6] have been proposed in the literature, and return a *point* estimation for a bid request. Subsequently, such estimation is used to approximate the true value of the corresponding ad impression. Despite the substantial progress that has been made, their accuracy is still far from perfect. One major reason is due to the large randomness of the user behaviors (e.g., click) and the cost uncertainty caused by the auction competition [1], which result in inherent *uncertainties* of estimated values. Motivated by work [37], explicitly factoring such uncertainties plays a critical role in optimizing a campaign's performance.

On the other hand, the dynamic nature of RTB requires modeling the correlations of bid requests under a given budget constraint in view of varying market competition [7]. The latest research considers a DSP's bidding process as a sequential decision process and proposes model-based or model-free reinforcement learning based bidding strategies [3; 31]. While it lays a solid groundwork for bid optimization, these studies are based on the fundamental assumption that the estimations of ad impressions' values are accurate. Unfortunately, as explained before, this assumption can hardly hold in practice. We point out that a reasonable bid optimization solution needs to consider three inherently correlated components, including the uncertainties of estimations of ad impression values, the state of a DSP (e.g., remaining budget and future auction number), and market competition. The latter two components determine the DSP's *risk tendency* (e.g., take more risk by bidding more aggressively or reduce risks by bidding more conservatively).

Based on the above observations, we propose an adaptive risk-aware bidding algorithm via reinforcement learning. To the best of our knowledge, it is the first work that simultaneously considers *prediction uncertainty* and the *dynamic risk tendency* of a DSP. We first theoretically unveil the intrinsic

<sup>1</sup>Our code is available at: <https://github.com/zhimengj0326/ekRLB>

Table 1: Notations and descriptions

Notation	Description
$\mathbf{x}$	The features of bid request.
$r_{mean}(\mathbf{x})$	The mean value of predicted CTR.
$r_{std}(\mathbf{x})$	The standard deviation of predicted CTR.
$(t, b)$	The remaining auction numbers and budget.
$\theta$	The ad impression value.
$m(\delta)$	The probability density distribution of market price.
$V(t, b)$	The expected total reward with starting state $(t, b)$ taking the optimal policy.
$a(t, b, \mathbf{x})$	The bid price for the request features $\mathbf{x}$ .
$\beta(t, b)$	The risk tendency for the resource state $(t, b)$ .
$U(t, b)$	The expected bid price depleting remaining budget.

relation between prediction uncertainty and risk tendency, which helps generate a modified value of an ad impression. With this formulation of an ad impression value, it is critical to properly optimize a DSP’s risk tendency. To this end, we propose two instantiations to model risk tendency, including an expert knowledge based formulation embracing three essential properties and an adaptive learning method based on self-supervised reinforcement learning. We summarize our key contributions as follows.

- We present an adaptive risk-aware bidding algorithm, which, for the first time, considers both prediction uncertainty and dynamic risk tendency to optimize bidding performance. This framework is based on a new formulation of an ad impression value by revealing the intrinsic relation between prediction uncertainty and risk tendency. We theoretically prove that this formulation allows achieving the optimal bid price based on VaR analysis.
- We propose two ways to determine the risk tendency of a DSP. We identify three basic properties of risk tendency, which lead to an expert knowledge based instantiation. To mitigate the extensive manual tuning efforts, we also design a self-supervised reinforcement learning method to learn the risk tendency based on experience.
- We conduct extensive experiments on two public datasets to validate the superiority of our adaptive risk-aware bidding algorithm and demonstrate the benefits of considering both prediction uncertainty and risk tendency.

## 2. PROBLEM FORMULATION

In the RTB system, each bidder of a DSP acts on behalf of an advertiser and competes for the advertisement auction every time a bid request is generated from a user visit. Given each auction opportunity, the bidder estimates the ad impression value and uncertainty, and then determine the bid price to maximize the cumulative ad impression value<sup>2</sup>. We aim to obtain better bidding strategy under the second-price auction, i.e., the bidder with the highest bid price wins the auction with the second-highest price payment. Related notations are summarized in Table 1.

### 2.1 Problem Definition

Considering budget constraints in real-time bidding, we formulate the bid optimization problem as a Markov decision

<sup>2</sup>In this paper, we use bidder and DSP interchangeably, and adopt CTR to estimate the ad impression value, while other metrics, such as CVR, can be adapted similarly.

process (MDP) in the *episode* level, where each episode consists of  $T$  sequential bid auctions accompanied with a budget of  $B$ . For each auction, we consider three pieces of critical information: (i) the remaining auction number  $t \in \{0, \dots, T\}$ ; (ii) the remaining budget  $b \in \{0, \dots, B\}$  and (iii) the mean value of predicted CTR (pCTR)  $r_{mean}(\mathbf{x}_t)$  and the corresponding standard deviation  $r_{std}(\mathbf{x}_t)$  for a bid request with feature vector  $\mathbf{x}_t$ . Hence, the bidder’s state  $s$  is defined as  $s \triangleq (t, b, \mathbf{x}_t)$ . We define  $V(t, b, \mathbf{x})$  as the long-run expected ad impression value accumulated from the current state  $s = (t, b, \mathbf{x})$ . Our target problem is that, given the remaining action number  $t$ , remaining budget  $b$ , the mean value of pCTR  $r_{mean}(\mathbf{x}_t)$  and corresponding standard deviation  $r_{std}(\mathbf{x}_t)$  for current bid request features  $\mathbf{x}_t$ , how can we determine the optimal bid price  $a(t, b, \mathbf{x}_t)$  to optimal cumulative ad impression value in a sequential decision-making process?

### 2.2 MDP Formulation

Reinforcement learning can be represented by tuple

$(\mathcal{S}, \mathcal{A}_s, \mathcal{P}_{sa}^s, \mathcal{R}_{sa}^s)$ , where  $\mathcal{S}$  denotes the state space,  $\mathcal{A}_s$  denotes the action (i.e., bid price) space for state  $s$ ,  $\mathcal{P}_{sa}^s$  and  $\mathcal{R}_{sa}^s$  represent the state transition probability and the immediate reward (i.e., pCTR) for the transition from state  $s$  to  $s'$  under action  $a$ . Note that  $t = 0$  and  $b = 0$  represent the end of an episode and the state with the depleted budget in sequential decision-making, respectively.

In the episode level bidding process, the state space  $\mathcal{S} = \{0, \dots, T\} \times \{0, \dots, B\} \times \mathbf{X}$ , where  $\mathbf{X}$  denotes the set of bid request features. Given state  $s = (t, b, \mathbf{x}_t)$ , the action space  $\mathcal{A}_s$  consists of all possible bid prices in set  $\{0, \dots, b\}$ , since possible bid prices are constrained by the remaining budget  $b$ . Let  $p\mathbf{x}(\mathbf{x}_t)$  denote the probability of the bid request feature  $\mathbf{x}_t$  for a potential ad impression and  $m(\delta|\mathbf{x}_t)$  denote the probability of market price  $\delta$  given feature  $\mathbf{x}_t$ . As there is no obvious dependency between winning price distribution and bid request features [40] in iPinYou dataset<sup>3</sup>, we assume that market environment  $m(\delta) \simeq m(\delta|\mathbf{x}_t)$ , i.e., the market price distribution is independent of the bid request feature. Such independent distribution assumption can be justified by empirical evaluation via comparing winning bid distribution against different features in the real-world iPinYou dataset [40].

For the state transition, if the bid price  $a$  is larger than the market price  $\delta$  (highest bid price among other competitors), then the bidder wins the ad auction, and state  $(t, b, \mathbf{x}_t)$  will transit to  $(t - 1, b - \delta, \mathbf{x}_{t-1})$  with probability  $p\mathbf{x}(\mathbf{x}_{t-1}) \sum_{\delta=0}^a m(\delta)$ . Otherwise, if  $a < \delta$ , the bidder will lose the auction and transit to state  $(t - 1, b, \mathbf{x}_{t-1})$  with probability  $p\mathbf{x}(\mathbf{x}_{t-1}) \sum_{\delta=a+1}^{+\infty} m(\delta)$ . The immediate reward is given by  $r_{mean}(\mathbf{x}_t)$  for  $t$ -th auction if the bidder wins the auction; otherwise, it is 0. Mathematically, the state transition probability and reward function are expressed as follows:

$$\mathcal{P}_{(t,b,\mathbf{x}_t),a}^{(t-1,b-\delta,\mathbf{x}_{t-1})} = \begin{cases} p\mathbf{x}(\mathbf{x}_{t-1})m(\delta), & \text{if } \delta \leq a; \\ p\mathbf{x}(\mathbf{x}_{t-1}) \sum_{\delta=a+1}^{+\infty} m(\delta), & \text{if } \delta > a. \end{cases}$$

<sup>3</sup>The similar box plot of winning price distribution against the features, such as hour, weekday, user browser, operation system and location regions of bid requests, can be observed in many campaigns of iPinYou [40].




Figure 1: An overview of the adaptive risk-aware bidding framework. Both pCTR and prediction uncertainty come from the auction environment. Considering the market environment and bidder state, the framework combines prediction uncertainty and risk tendency to adjust ad impression values and decide bid prices.

$$\mathcal{R}_{(t,b,\mathbf{x}_t),a}^{(t-1,b-\delta,\mathbf{x}_{t-1})} = \begin{cases} \theta(t,b,\mathbf{x}_t), & \text{if } \delta \leq a; \\ 0, & \text{if } \delta > a. \end{cases}$$

### 3. METHODOLOGY

In this section, we introduce the **R**isk-aware **R**einforcement **L**earning **B**idding (RRLB) framework that effectively integrates prediction uncertainty and a bidder’s risk tendency into a reinforcement learning framework. An overview of RRLB is illustrated in Figure 1. In the following sections, we first explain the uncertainties of CTR prediction with Bayesian logistic regression and then describe the risk-aware bid optimization framework and two proposed instantiations to determine the *risk tendency*. Finally, we describe the model-based reinforcement learning method mapping the adjusted ad impression value to the final bid price.

#### 3.1 Uncertainty of CTR Prediction

Similar to [37], we adopt Bayesian logistic regression to explicitly measure the uncertainties of predicted CTR (pCTR) values. In Bayesian logistic regression, each weight  $\mathbf{w}$  is treated as a random variable instead of a parameter, and the variance of the random variable represents the uncertainty of the corresponding feature. The likelihood of observing the correct binary click label  $y$  given features  $\mathbf{x}$  and weights of logistic regression  $\mathbf{w}$  is

$$p(y|\mathbf{x}, \mathbf{w}) = \sigma(\mathbf{w}^T \mathbf{x})^y (1 - \sigma(\mathbf{w}^T \mathbf{x}))^{1-y}, \quad (1)$$

where the sigmoid function is defined as  $\sigma(x) = \frac{1}{1+e^{-x}}$ . Note that  $\mathbf{w}$  is modeled as a random variable with a p.d.f  $p(\mathbf{w})$  in the Bayesian version of logistic regression. The marginal conditional probability  $p(y|\mathbf{x})$  is  $p(y|\mathbf{x}) = \int p(y|\mathbf{x}, \mathbf{w})p(\mathbf{w})d\mathbf{w}$ . The Gaussian prior  $N(\mu_0, q_0^{-1}\mathbf{I})$  and Laplace approximation of posterior on  $\mathbf{w}$  are adopted for weight distribution update, where  $q_0$  is initial weight precision. The goal is to maximize the posterior weight distribution  $p(\mathbf{w}|\mathbf{x}, y)$  given a bid request’s feature vector  $\mathbf{x}$  and label  $y$ . The model output is a probability estimation of the occurrence of a click event, defined as  $\hat{y} = P(y = 1|\mathbf{x})$ . The variance of weight is lower when the associated feature emerges more frequently, which means that such a model can measure the data completeness for each feature. Via updating the mean and covariance matrix of the weight  $\mathbf{w}$ , the distribution of CTR  $p_{\hat{y}|\mathbf{x}}(\hat{y})$  can be obtained. Subsequently, we define the mean and standard deviation of CTR as  $r_{mean}(\mathbf{x}) = \mathbb{E}_{p_{\hat{y}|\mathbf{x}}}[\hat{y}]$  and  $r_{std}(\mathbf{x}) = \sqrt{\mathbb{D}_{p_{\hat{y}|\mathbf{x}}}[\hat{y}]}$ , where  $\mathbb{E}_{p_{\hat{y}|\mathbf{x}}}[\cdot]$  and  $\mathbb{D}_{p_{\hat{y}|\mathbf{x}}}[\cdot]$  denote the

expectation and variance over the CTR distribution  $p_{\hat{y}|\mathbf{x}}$ . The standard deviation of CTR reflects the uncertainties of pCTR values. How to obtain uncertainties of other prediction models is beyond the scope of this paper.

#### 3.2 Theoretical Relation Between Uncertainty and Risk Tendency

The key intuition of RRLB is to decompose the value of an ad impression  $\theta(t, b, \mathbf{x}_t)$  as the weighted sum of two parts: the mean pCTR and a compound term that simultaneously reflects prediction uncertainty and a bidder’s risk tendency. Formally, the ad impression value is defined as follows:

$$\theta(t, b, \mathbf{x}_t) = r_{mean}(\mathbf{x}_t) + \beta(t, b)r_{std}(\mathbf{x}_t), \quad (2)$$

where  $\beta(t, b)$  denotes the bidder’s risk tendency at resource state  $(t, b)$ , which is a subset of bidding state  $(t, b, \mathbf{x}_t)$ , which represents the intrinsic status of a bidder in terms of remaining auction number  $t$  and remaining budget  $b$  in an episode. Next we provide a theoretical motivation for the formulation of Eq. (2). The core idea is based on the value at risk (VaR) theory borrowed from finance [26; 19], where VaR estimates how much the predicted CTR under/over-estimates with a given probability. Note that the goal of a bidding strategy is to improve the cumulative ad impression value. Given the current bid request feature vector  $\mathbf{x}_t$ , remaining budget  $b$ , and remaining auctions number  $t$ , let  $V^a(t, b, \mathbf{x}_t)$  and  $V_{std}^a(t, b, \mathbf{x}_t)$  be the cumulative estimated impression value and uncertainty of the winning ads with the bidding strategy  $a(t, b, \mathbf{x}_t)$ . Then we define VaR of the cumulative ad impression value with the bidding strategy  $a(t, b, \mathbf{x}_t)$  as follows:

$$V_\lambda^a(t, b, \mathbf{x}_t) \triangleq V^a(t, b, \mathbf{x}_t) + \lambda(t, b)V_{std}^a(t, b, \mathbf{x}_t), \quad (3)$$

where state-associated coefficient  $\lambda(t, b)$  is the risk preference that balances the cumulative estimated impression value and uncertainty. Note that  $\beta(t, b)$  and  $\lambda(t, b)$  balance the estimated value and uncertainty for the current auction and all the remaining auctions, respectively. The optimal VaR bid price  $a_{VaR}$  maximizes the VaR of the cumulative ad impression value  $V_\lambda^a(t, b, \mathbf{x}_t)$  as follows,

$$a_{VaR}(t, b, \mathbf{x}_t) = \arg \max_{0 \leq a \leq b} V_\lambda^a(t, b, \mathbf{x}_t). \quad (4)$$

The theorem below shows that the linear combination in Eq. (2) can achieve the optimal VaR bid price  $a_{VaR}$ .

**Theorem 1 (Risk Tendency Optimality).** *The RRLB framework adopting the linear formulation in Eq. (2) can achieve the optimal VaR bid price  $a_{VaR}(t, b, \mathbf{x}_t)$  if the same risk tendency  $\beta(t, b) = \lambda(t, b)$  is used.*

**PROOF.** *Note that the bidder wins an ad impression if the bid price  $a$  is larger than the market price  $\delta$ . The cumulative estimated ad impression value and the corresponding uncertainty should satisfy*

$$V^a(t, b, \mathbf{x}_t) = \sum_{\delta=0}^{a(t,b,\mathbf{x}_t)} m(\delta)r_{mean}(\mathbf{x}_t) + \sum_{\delta=0}^{a(t,b,\mathbf{x}_t)} m(\delta)V(t-1, b-\delta) + \sum_{\delta=a(t,b,\mathbf{x}_t)+1}^{\infty} m(\delta)V^a(t-1, b),$$

$$V_{std}^a(t, b, \mathbf{x}_t) = \sum_{\delta=0}^{a(t, b, \mathbf{x}_t)} m(\delta) r_{std}(\mathbf{x}_t) + \sum_{\delta=0}^{a(t, b, \mathbf{x}_t)} m(\delta) V_{std}(t-1, b-\delta) + \sum_{\delta=a(t, b, \mathbf{x}_t)+1}^{\infty} m(\delta) V_{std}(t-1, b),$$

Combining above two equations, we have

$$V_{\lambda}^a(t, b, \mathbf{x}_t) = \sum_{\delta=0}^{a(t, b, \mathbf{x}_t)} m(\delta) (r_{mean}(\mathbf{x}_t) + \lambda(t, b) r_{std}(\mathbf{x}_t)) + \sum_{\delta=0}^{a(t, b, \mathbf{x}_t)} m(\delta) V_{\lambda}^a(t-1, b-\delta) + \sum_{\delta=a(t, b, \mathbf{x}_t)+1}^{\infty} m(\delta) V_{\lambda}^a(t-1, b),$$

By firstly setting  $\theta(t, b, \mathbf{x}_t) = r_{mean}(\mathbf{x}_t) + \beta(t, b) r_{std}(\mathbf{x}_t)$  and adopting the model-based method [3], it can be seen that the linear combination in  $\theta(t, b, \mathbf{x}_t)$  can achieve the optimal bid price  $a_{VaR}(t, b, \mathbf{x}_t)$ , which maximizes the VaR of the cumulative ad impression value.

Intuitively, a rational risk tendency should be a function of the resource state, defined by  $(t, b)$ , of a bidder. A budget-restrained/abundant bidder would act very differently in taking risks during the bid. However, we deem that the risk tendency is independent of a random bid request  $\mathbf{x}$ .

### 3.3 Expert Knowledge Based Risk Tendency

We leverage expert knowledge on RTB to design the first instantiation of risk tendency  $\beta(t, b)$  to reveal the intrinsic risk preference of a rational bidder. We distill three key rules as follows. (i) The sufficiency of the remaining budget  $b$  determines the sign of risk tendency  $\beta(t, b)$ , where a positive risk tendency indicates a strong preference to win auctions. (ii) The partial derivative of risk tendency  $\beta(t, b)$  w.r.t. remaining budget  $b$  (remaining auction number  $t$ ) should be positive (negative) because more budget naturally allows the bidder to take the risk of bidding more ad impressions. (iii) When the remaining budget  $b$  and remaining auction number  $t$  are relatively large (e.g., the beginning of an ad campaign), risk tendency  $\beta(t, b)$  depends on the ratio of  $b$  to  $t$  and the extent of market competition. Formally, we express the three rules as follows.

(i) Sign of risk tendency:

$$\beta(t, b) \begin{cases} \geq 0, & b \text{ is sufficient at current } t; \\ < 0, & \text{otherwise.} \end{cases} \quad (5)$$

(ii) Monotonicity of risk tendency:  $\frac{\partial \beta(t, b)}{\partial t} < 0$ ,  $\frac{\partial \beta(t, b)}{\partial b} > 0$ .

(iii) Approximation for the scenarios of large remaining budget and auction number:  $\beta(t, b) \simeq \beta(t', b')$  if  $\frac{b}{t} = \frac{b'}{t'}$ .

Before elaborating the exact formulation of  $\beta(t, b)$ , we have to quantify the sufficiency of the budget as pointed out in the first rule. In practice, the budget richness is highly related to the level of market competition. Given a certain market price distribution  $m(\delta)$  and resource state  $(t, b)$ , let  $U(t, b)$  denote the expected bid price for an auction such that the remaining budget will be depleted in the remaining future auctions. Supposing that budget  $b$  is evenly allocated to the remaining  $t$  auctions, we can calculate  $U(t, b)$  through the following formula:

$$\sum_{\delta=0}^{U(t, b)} \delta m(\delta) = \frac{b}{t}. \quad (6)$$

---

**Algorithm 1:** Expert knowledge based reinforcement learning bidding strategy

---

**Input:** market price probability  $m(\delta)$ , episode length  $T$ , budget  $B$ , average pCTR and standard deviation of CTR  $\bar{r}_{mean}$  and  $\bar{r}_{std}$

**Output:** Optimal bid price  $a$  for current state  
Parameter tuning for the manually designed risk tendency;

Update Value function  $V(t, b)$  based on risk tendency ;

**for**  $\delta = 0, 1, \dots, \min(\delta_{max}, b)$  **do**

**if**  $\theta(t, b, \mathbf{x}) + V(t-1, b-\delta) - V(t-1, b) < 0$  **then**  
         $a(t, b, \mathbf{x}) \leftarrow \delta$ ;  
        **break**;

**end**

---

Since the bidder wins an auction only if bid price  $U(t, b)$  is higher than market price  $\delta$ , the left side of Eq. (6) represents the expected actual cost, which should be the same as  $\frac{b}{t}$  based on the assumption of even budget allocation.

Based on the above intuition of budget richness, we formally define risk tendency as follows:

$$\beta(t, b) = \tanh\left(\alpha \frac{U(t, b) - \hat{U}}{\hat{U}}\right), \quad (7)$$

where  $\alpha$  is a positive hyperparameter that controls the slope of risk tendency,  $\hat{U}$  is the budget richness threshold tuned from historical data, and function  $\tanh(\cdot)$  confines risk tendency within the range  $(-1, 1)$ . It can be observed that the proposed risk tendency formulation satisfies all three expert knowledge based rules. First, the expected bid price  $U(t, b)$ , which measures the amount of budget richness, determines the sign of risk tendency  $\beta(t, b)$ , which is non-negative only if  $U(t, b) \geq \hat{U}$  as required in rule (i). Second, both  $U(t, b)$  and  $\beta(t, b)$  increase with budget  $b$ , and decrease with  $t$  as required in rule (ii). Third, the expected cost is proportional to the ratio  $\frac{b}{t}$  as shown in Eq. (6), which helps the subsequent design of risk tendency to meet rule (iii).

By using the expert knowledge based risk tendency, we can obtain the adjusted ad impression value using  $\theta(t, b, \mathbf{x}) = r_{mean}(\mathbf{x}) + \beta(t, b) r_{std}(\mathbf{x})$ . Subsequently, the model-based reinforcement learning method is applied to obtain the final bid price as will be introduced in Section 3.5.

Compared with the previous efforts [3], the bid price not only reveals the estimated mean value of an ad impression but also takes into account the estimation risk and risk tendency based on the bidding state. The expert knowledge based Reinforcement Learning Bidding (ekRLB) algorithm is summarized in Algorithm 1, where the value function update could be calculated using dynamic programming.

### 3.4 Self-Supervised Risk Tendency

Although the above explicit design of risk tendency may capture the intrinsic risk preference well, it requires extensive human expert knowledge and efforts on tuning Equation (7). To be specific, a careful selection of hyperparameters  $\alpha$  and  $\hat{U}$  is required through many trials. With the variation of the market environment, these presetting hyperparameters tend to obtain sub-optimal performance and need to be updated periodically. Furthermore, the computation cost of calculating budget richness  $U(t, b)$  is prohibitive as it requires considering all combinations of  $t$  and  $b$ . To avoid such

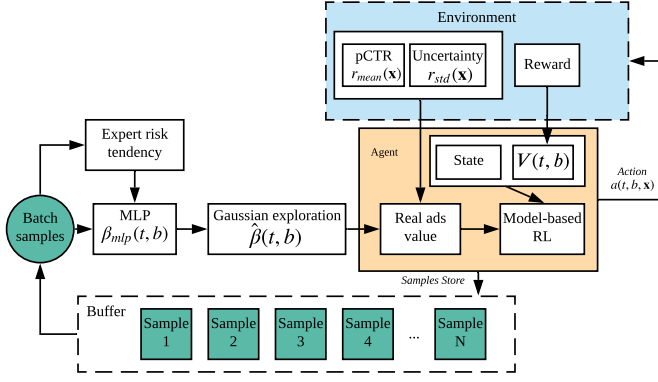


Figure 2: An overview of the self-supervised reinforcement learning. The risk tendency is generated from an MLP trained by experience buffer and batch sampling in a self-supervision way. The bidder decides bid prices based on risk tendency and interacts with the RTB environment.

manual efforts, we propose a self-supervised reinforcement learning bidding (ssRLB) method as shown in Figure 2 to automatically generate risk tendency via a multi-layer perceptron (MLP). Self supervised by the bidding history, we update the mapping function  $\beta_{mlp}(t, b) = MLP(t, b; \mathbf{W}_{mlp})$  generated from an MLP to approximate the risk tendency at resource state  $(t, b)$  with trainable weight  $\mathbf{W}_{mlp}$ . The ssRLB framework consists of a Gaussian exploration block, an experience buffer, an MLP mapping function, and batch sampling. We explain the details of each component as follows:

**Gaussian exploration.** We adopt exploration on risk tendency by adding Gaussian noise to  $\beta_{mlp}(t, b)$  as  $\hat{\beta}(t, b) = \beta_{mlp}(t, b) + \epsilon$ , where noise  $\epsilon \sim \mathcal{N}(0, \sigma^2)$ . The noise variance  $\sigma^2$  is adjustable to provide a trade-off between exploitation and exploration in reinforcement learning.

**Experience buffer.** Motivated by experience replay optimization in reinforcement learning [35], we adopt the experience buffer to store good experiences represented by a quaternary set  $\mathcal{B} = (t, b, \hat{\beta}(t, b), V_{episode})$  from the bidding history, where  $V_{episode}$  denotes the cumulative reward for the entire episode. A “good” experience means that a larger reward  $V_{episode}$  is obtained by using risk tendency  $\hat{\beta}(t, b)$ . Let  $N$  be the buffer’s total length. The samples with the lowest reward will be removed if the buffer is full. In this way, the experience buffer could always provide the best samples explored so far for training the MLP.

**Batch sampling.** Batch sampling is responsible for sampling batch  $\mathcal{B}_{batch}$  from the buffer. We apply a simple uniform sampling to generate  $\mathcal{B}_{batch}$ .

**Training MLP mapping function.** Given the experience batch  $\mathcal{B}_{batch}$ , we update the weight in the MLP mapping function by minimizing the mean square loss function:

$$\mathcal{L} = \sum_{(t, b, \hat{\beta}(t, b), \cdot) \in \mathcal{B}_{batch}} \|MLP(t, b; \mathbf{W}_{mlp}) - \hat{\beta}(t, b)\|^2.$$

Since we only preserve the experiences with larger rewards in the buffer, the mapping function will be updated under supervision toward learning a good risk tendency.

In a nutshell, experience buffer stores “good” experienced risk tendency while the sampled experienced risk tendency is a supervised signal for MLP mapping function training.

---

### Algorithm 2: Self-Supervised Risk Tendency Learning Algorithm

---

**Input:** The historical data sample with pCTR, risk, market price, and click labels, episode length  $T$ , budget  $B$ .

**Output:** Optimal bid price

Initialize the risk tendency, and uniform replay policy ;  
Update Value function  $V(t, b)$  based on Equation (8);

**for each episode do**

**for each ad auction in the current episode do**

        Provide the bid price based on Algorithm 1;  
        Execute auction and observe  $(t + 1, b)$  and cumulative reward and risk;

**end**

    Calculate the cumulative reward for an entire episode;

**if the cumulative reward  $V_{episode}$  is larger than that in Buffer then**

$\mathcal{B} \leftarrow (t, b, \hat{\beta}(t, b), V_{episode})$ ;

        Uniformly sample a batch  $\mathcal{B}_s$  from  $\mathcal{B}$  ;

        Train the MLP based on the batch sample ;

        Update risk tendency  $\hat{\beta}(t, b)$ ;

        Update Value function  $V(t, b)$  based on Eq. (8).

**end**

---

The algorithm for self-supervised Reinforcement Learning Bidding (ssRLB) is shown in Algorithm 2.

### 3.5 Bid Price Determination

Previous sections introduce how to modify the ad impression value with prediction uncertainty and risk tendency. Next, we explain how to calculate the final bid price. We adopt the model-based reinforcement learning bidding strategy [3] to maximize the cumulative reward. Specifically, we regard the pCTR  $r_{mean}(\mathbf{x}_t)$  as the immediate reward for  $t$ -th auction, and cumulative reward  $V(t, b, \mathbf{x}_t)$  is defined as the expected cumulative reward starting from state  $(t, b, \mathbf{x}_t)$  with the optimal bid price. By definition, we have  $V(0, b, \mathbf{x}_t) = V(0, b) = 0$  since there is no available auction. Similar to [3], the updated policy for the cumulative reward is given by:

$$V(t, b, \mathbf{x}_t) = \max_{0 \leq a \leq b} \left\{ \sum_{\delta=0}^a \int_{\mathbf{X}} m(\delta) p_{\mathbf{x}}(\mathbf{x}_{t-1}) \cdot (r_{mean}(\mathbf{x}_t) + V(t-1, b-\delta, \mathbf{x}_{t-1})) d\mathbf{x}_{t-1} + \sum_{\delta=a+1}^{+\infty} \int_{\mathbf{X}} m(\delta) p_{\mathbf{x}}(\mathbf{x}_{t-1}) V(t-1, b, \mathbf{x}_{t-1}) d\mathbf{x}_{t-1} \right\},$$

where  $\mathbf{X}$  represents the entire feature vector space, and the two integrations represent the immediate reward for winning and losing cases. Furthermore, the cumulative reward without observation on  $\mathbf{x}$  can be obtained by integrating the over bid request feature vector  $\mathbf{x}$ . Formally, the cumulative reward  $V(t, b)$  is:

$$V(t, b) \approx \max_{0 \leq a \leq b} \left\{ \sum_{\delta=0}^a m(\delta) r_{avg} + \sum_{\delta=0}^a m(\delta) V(t-1, b-\delta) + \sum_{\delta=a+1}^{\infty} m(\delta) V(t-1, b) \right\},$$

where  $r_{avg} = \int_{\mathbf{X}} p_{\mathbf{x}}(\mathbf{x}_{t-1}) r_{mean}(\mathbf{x}_{t-1}) d\mathbf{x}_{t-1}$  is the average ad impression value over the entire feature vector space. The cumulative reward  $V(t, b)$  can be iteratively updated given the average ad impression value and market price distribution. Note that  $\sum_{\delta=0}^{\infty} m(\delta) = 1$ . The bid price at state  $(t, b, \mathbf{x}_t)$  is calculated by:

$$\begin{aligned} a(t, b, \mathbf{x}_t) &\triangleq \arg \max_{0 \leq a \leq b} V(t, b, \mathbf{x}_t) \\ &= \arg \max_{0 \leq a \leq b} \left\{ \sum_{\delta=0}^a m(\delta) \left( \theta(t, b, \mathbf{x}_t) \right. \right. \\ &\quad \left. \left. + V(t-1, b-\delta) - V(t-1, b) \right) \right\} \\ &= \arg \max_{0 \leq a \leq b} \left\{ \sum_{\delta=0}^a m(\delta) g(\delta) \right\}, \end{aligned} \quad (8)$$

where  $g(\delta) \triangleq \theta(t, b, \mathbf{x}_t) + V(t-1, b-\delta) - V(t-1, b)$ . The cumulative reward  $V(t, b)$  monotonically increases w.r.t. the remaining budget  $b$ , and  $g(\delta)$  monotonically decreases. If  $g(b) > 0$  (e.g., the impression value is extremely high), the optimal bid price is all the remaining budget. If  $g(b) < 0$  (e.g., the impression value is moderate), there must exist a *unique* integer price, defined as  $A$ , satisfying the following three conditions: (1)  $0 \leq A \leq b$ ; (2)  $g(A) \geq 0$  and (3)  $g(A+1) < 0$ . According to Eq. (8), the optimal price is thus determined by this *unique* integer price  $A$ . To summarize, the optimal bid price is ultimately dictated by:

$$a(t, b, \mathbf{x}_t) = \begin{cases} b, & \text{if } g(b) \geq 0; \\ A, & \text{if } g(b) < 0. \end{cases} \quad (9)$$

## 4. EXPERIMENTS

In this section, we conduct experiments to evaluate our RRLB framework with the two instantiations of risk tendency, ekRLB, and ssRLB, and answer the following three research questions:

- **Q1:** How does our risk-aware bid optimization solution compare with several baselines in terms of the total number of clicks?
- **Q2:** How do prediction uncertainty and risk tendency affect the total number of clicks?
- **Q3:** How do the hyperparameters affect the long-run performance of the risk-aware bidding strategy?

### 4.1 Experimental Setting

#### 4.1.1 Datasets.

Our experiments are conducted on two real-world datasets, iPinYou<sup>4</sup> and YOYI<sup>5</sup>. We follow the data preprocessing in [37] to split training/test sets and obtain the estimations and uncertainties of ad impression values. The dataset description is as follows:

- **iPinYou** dataset contains 19.5M ad impressions, 14.79K clicks, and 16.0K spend (in CNY) over 9 campaigns during 10 days in the year 2013. We follow the data preprocessing configuration in [37] to split it into training/test sets and obtain the estimations and risk of ad impression values.

<sup>4</sup><https://contest.ipinyou.com/>

<sup>5</sup><http://apex.sjtu.edu.cn/datasets/7>

Table 2: The comparison of total click number ( $c_0 = 1/2$ ).

iPinYou	Lin	RLB	ekRLB	CRTRLB	CURLB	ssRLB
1458	401	428	428	416	423	422
2259	19	73	75	59	70	70
2261	9	44	51	36	64	61
2821	109	209	211	198	214	206
2997	295	376	382	384	360	276
3358	208	233	233	231	222	217
3386	157	293	294	290	293	290
3427	263	290	292	295	286	302
3476	206	230	232	241	233	204
Average	185.2	241.8	<b>244.2</b>	238.9	240.6	226.9
YOYI	725	890	894	873	840	<b>914</b>

- **YOYI** dataset includes 402M ad impressions, 500K clicks, and 428K spend (in CNY) during 8 days in the year 2016. The first 7 days and the last day are set as the training data and test data, respectively.

#### 4.1.2 Compared methods.

We compare ekRLB and ssRLB with two state-of-the-art baselines. **Lin** is a linear bidding strategy with bid price  $a_{Lin} = b_0 \theta(\mathbf{x})$ , where parameter  $b_0$  can be tuned on training data [24]. **RLB** is a model-based reinforcement learning bidding strategy [3], which achieves fine-grained pCTR-price mapping, and adaptively adjusts bid prices for different states. These two baselines only make use of  $r_{mean}(\mathbf{x})$ . Besides, we also consider two variants of our proposed reinforcement learning bidding framework, including constant risk tendency (**CRTRLB**) and constant uncertainty (**CURLB**). CRTRLB sets risk tendency to a constant for any resource states, i.e.,  $\beta(t, b) \equiv \beta_0$ , which is equivalent to the previous work [37]. CURLB assigns a constant risk for all ad impressions, i.e.,  $r_{std}(\mathbf{x}) \equiv r_0$ . These two variants can be used to validate the superiority of dynamic risk tendency and the necessity for risk estimation.

#### 4.1.3 Evaluation sketch.

Given a predefined budget and episode length for each bidding campaign, we evaluate different bidding strategies in terms of the total click number. Evaluation sketch including experiment flow, budget constraints, and more details as follows.

- *Experiment flow.* The bidding data is a list of ad impression records, each of which is accompanied by the bid request feature vector, the market price, and the user response (click) label. Based on Bayesian logistic regression, we obtain the estimation mean and risk of ad displaying the value for each impression record. Following the model-based reinforcement learning bidding in Section 3.5, we divide training and test datasets into episodes, each of which contains  $T$  impression records and budget  $B$ . All bidding strategies are evaluated in the episode level. Episode length  $T$  influences the computation complexities of bidding price and budget richness, which are required over all remaining impression numbers  $t$ . We set the size of episode length  $T = 1000$ . In the second-price auction platform, a bidding agent wins an ad impression if its price is higher than the market price, and only needs to pay the cost of the market price. Then the bidding

agent receives the reward of the click label and applies it to update the bidding strategy.

- *Budget constraints.* We allocate the budget as follows:  $B = CPM_{train} \times 10^{-3} \times T \times c_0$ , where  $CPM_{train}$  and  $c_0$  are the cost per mille impressions in the training dataset and budget coefficient, respectively. We compare the models using the budget coefficient set  $\{1/32, 1/16, 1/8, 1/4, 1/2\}$ .
- *More details.* For ekRLB, we tune the risk tendency hyperparameters  $\alpha$  and  $\hat{U}$  in Eq. (7) on the training dataset to optimize the number of total clicks. For ssRLB, the mapping function  $MLP(t, b; \mathbf{W}_{mlp})$  is realized by a four-layer MLP with 64 hidden units. Weight  $\mathbf{W}_{mlp}$  is trained with the Adam optimizer to minimize the mean square loss function at the learning rate of  $1 \times 10^{-3}$ . The experience buffer size and batch size are  $1 \times 10^5$  and 32, respectively. We update the buffer every 5 training episode.

## 4.2 Comparison Results

We present the total click number of different methods on the 9 campaigns of iPinYou and YOYI with budget coefficient  $c_0 = 1/2$  in Table 2 to answer question Q1. On iPinYou, ekRLB obtains the best performance on most campaigns, achieving the largest average total click number of 244.2. On YOYI, both ekRLB and ssRLB outperform RLB, and ssRLB achieves the largest click number of 914. The two variants CRTRLB and CURLB obtain fewer clicks on both datasets, which validates the benefits of considering both uncertainty and risk tendency in Eq. (2). Note that ssRLB obtains fewer clicks than ekRLB on iPinYou but more on YOYI probably because the training of self-supervised reinforcement learning is less stable. Nevertheless, we deem that ssRLB is still a valuable alternative since it does not require manual tuning of  $\alpha$  and  $\hat{U}$  over a large volume of historical data. We leave the improvement on risk tendency learning for future work.

In the left part of Figure 3, we further show the performance improvements of RLB and ekRLB over Lin on iPinYou under the entire budget coefficient set. Two notable findings can be observed. First, the two reinforcement learning based strategies outperform Lin consistently over all budget settings, which validates the benefits of modeling the bidding process as an MDP. Given a specific budget constraint, all bid requests are inherently correlated instead of independent in Lin since previous bid prices determine the remaining budget. Second, ekRLB gets more clicks than the traditional RLB, especially in the case with larger budgets. Such observation demonstrates the advantage of explicitly modeling prediction uncertainty within a risk-aware reinforcement learning framework.

## 4.3 Ablation Study

To answer Q2, we study the individual contributions of prediction uncertainty and risk tendency by comparing with constant uncertainty and risk tendency. Regarding ekRLB as the benchmark method, the right part of Figure 3 shows the performance degradation of the two variants CRTRLB and CURLB. Along with the results in Table 2, we can find that (i) both CRTRLB (with constant risk tendency) and CURLB (with constant uncertainty) perform worse than ekRLB, which proves that it is critical to consider both prediction uncertainty and risk tendency to achieve an optimal bid optimization framework. (ii) CURLB achieves bet-




Figure 3: Left: The click number improvements of RLB and ekRLB over Lin in all budget coefficients on iPinYou. Right: The click number diminution of CRTRLB and CURLB over ekRLB on iPinYou.

ter performance than CRTRLB, suggesting that modeling risk tendency is even more important than prediction uncertainty. This is because risk tendency determines how to use (e.g., add or subtract) prediction uncertainty.

## 4.4 Hyperparameter Study

To answer Q3, we give a comprehensive hyperparameter study to investigate how hyperparameters slope  $\alpha$ , constant uncertainty  $r_0$ , and constant risk tendency  $\beta_0$  affect the performance of the methods ekRLB, CURLB, and CRTRLB, respectively. Specifically, we report the experimental results on campaign 1458 of the iPinYou dataset with budget coefficient  $c_0 = 1/2$  in Table 3. The performance metrics include click number and budget consumption ratio, where the budget consumption ratio is defined as the cumulative cost over the overall budget. The detailed experimental results and analysis are summarized as follows.

*Slope  $\alpha$  for ekRLB.* Hyperparameter  $\alpha$  controls the slope of risk tendency, and then influences the bid price of ekRLB. Note that, in the case  $\alpha = 0$ , we have 0 risk tendency based on Eq. (6) of the main text, and ekRLB degenerates to RLB. We can clearly observe that the total click number reaches the highest value at a medium slope scale of  $\alpha = 0.1$ . This result implies that the total click number is not sensitive w.r.t  $\alpha$  and that ekRLB is robust on the hyperparameter  $\alpha$ .

*Constant uncertainty  $r_0$  for CURLB.* The method CURLB achieves relatively comparable performance with ekRLB even though a constant uncertainty is applied. We further study the influence of constant uncertainty  $r_0$  by varying it in range  $[0, 1.8 \times \bar{r}_{std}]$ . Here we set constant uncertainty  $r_0$  as a constant coefficient multiplying with the average uncertainty  $\bar{r}_{std}$  in the training dataset. For the case  $r_0 = 0$ , we have a prediction uncertainty of 0 for all ad impressions, when CURLB degenerates to RLB. We observe that the hyperparameter  $r_0 = 0.2 \times \bar{r}_{std}$  achieves the best performance in terms of both click numbers since higher/lower risks lead to ad impressions with overly large/small prices that decrease the total click number.

*Constant risk tendency  $\beta_0$  for CRTRLB.* Compared with ekRLB, it can be observed that a random selection of constant risk tendency  $\beta_0$  in CRTRLB greatly damages model performance. We further study its influences by considering  $\beta_0$  within range  $[-0.5, 0]$ . We only use the negative risk tendencies since a positive one tends to overestimate an ad impression and usually results in worse performance. For the case  $\beta_0 = 0$ , we remove risk tendency and degenerate CRTRLB to RLB. We observe that  $\beta_0 = 0.0$  achieves the most clicks and that when  $\beta_0$  equals  $-0.4$  to  $-0.5$ , CRTRLB extremely underestimates the ad impression value and bids

Table 3: Hyperparameter study on slope  $\alpha$  (ekRLB), uncertainty  $r_0$  (CURLB) and risk tendency  $\beta_0$  (CRTRLB).

Hyperparameters		#click	consumption ratio
$\alpha$	0.0	1928	98.70%
	0.001	1925	98.71%
	0.01	1927	98.86%
	0.1	<b>1930</b>	99.54%
	0.2	1923	99.00%
	0.3	1923	99.11%
	0.4	1924	99.20%
	0.5	1928	99.29%
$r_0$	0	1928	99.83%
	$0.2 * \bar{r}_{std}$	<b>1932</b>	99.88%
	$0.4 * \bar{r}_{std}$	1922	99.89%
	$0.6 * \bar{r}_{std}$	1919	99.90%
	$0.8 * \bar{r}_{std}$	1910	99.90%
	$1.0 * \bar{r}_{std}$	1898	99.91%
	$1.2 * \bar{r}_{std}$	1896	99.91%
	$1.4 * \bar{r}_{std}$	1884	99.91%
	$1.6 * \bar{r}_{std}$	1878	99.91%
$1.8 * \bar{r}_{std}$	1876	99.91%	
$\beta_0$	0.0	<b>1928</b>	98.69%
	-0.001	1925	98.68%
	-0.01	1922	98.55%
	-0.1	1857	96.50%
	-0.2	1629	90.00%
	-0.3	845	32.77%
	-0.5	159	5.25%
	-0.5	0	0.06%

with a small price, leading to losing almost all ad impressions and low budget consumption ratio.

#### 4.5 Visualization of Risk Tendency

Risk tendency reflects the risk preference of a rational bidder in given states. We designed two instantaneous to learn risk tendency, one based on expert knowledge and the other based on self-supervised learning. We visualize these two methods' risk tendencies in Figure 4. It can be observed that both approaches have similar trends in mapping states to risk tendency, which implies that our self-supervised learning algorithm aligns well with the expert knowledge. Specifically, risk tendencies are negative for those resource-limited states at the upper left corner, where a bidder has a small budget for a large number of remaining auctions. In such a state, the bidder prefers to bid with conservative prices. On the other hand, risk tendencies change to positive for those resource-rich states at the bottom right corner because the bidder has a sufficient budget for remaining auctions to support more aggressive bid prices.

### 5. RELATED WORK

**User response prediction.** User response prediction can be modeled as a probability estimation task, e.g., click-through rate (CTR) [10], conversion rate (CVR) [21]. Click-through rate (CTR) prediction is widely studied to estimate the value of an ad impression. It plays a key role in display advertising [18; 21; 30; 10; 14]. The traditional solutions use linear models [18], gradient boosting decision trees (GBDT) [29] or factorization machines [23] to predict CTR. Deep learning based methods [36; 27; 5; 12; 25; 13; 33; 2;




Figure 4: Left: Risk tendency obtained by expert knowledge. Right: Risk tendency learned from self supervision. X-axis and Y-axis represent the remaining budget  $b$  and remaining auctions  $t$ , respectively.

41; 6; 38] are also used for CTR prediction tasks. Many of them follow the Embedding&MLP paradigm. All these methods aim to only predict the CTR without any information guiding the confidence level. [37] is the only study that simultaneously captures the estimation and risk of pCTR prediction based on Bayesian logistic regression.

**Risk Management Techniques.** The goal of risk management in finance is to guarantee that the risk does not hurt the business profit [22]. The risk and expected return of each asset can be quantified by the mean return of the asset and corresponding covariance matrix [20; 16]. The mean-variance analysis is adopted to achieve the trade-off between the risk and the expected return [28]. [39] measures the campaign-level risk and returns in a special case of arbitrage. As for impression-level risk management, [37] proposes a risk-aware bidding strategy based on the value at risk (VaR) with campaign-level profit gain.

**Bidding strategy.** Truthful bidding is the most fundamental strategy that has been proven to be optimal in second-price auctions with unlimited budgets [17]. In practice, every ad campaign has a budget constraint, where linear [24] and model-based/model-free reinforcement learning-based strategies [3; 31] can be adopted to determine bid prices. Other advanced methods, such as Recurrent Neural Networks [11] optimization [15], are developed for bidding with budget constraints. However, these methods inevitably assume the user response prediction is perfect and ignore the inherent estimation uncertainty. [37] is most relevant to our work, which proposes a risk management algorithm based on value at risk, but ignores intrinsic interactions between the market environment and the state of a bidder. In our work, we propose a risk-aware reinforcement learning based bidding strategy that explicitly considers such interactions.

### 6. CONCLUSION

In this paper, we investigated the bid optimization problem in RTB with the benefits of risk information. For the bid price determination with budget constraint, we, to the best of our knowledge, firstly consider both estimation uncertainty and the dynamic risk tendency. Specifically, we first theoretically analyze the relation between prediction uncertainty and the risk tendency of a bidder, and then proposed an adaptive risk-aware bidding algorithm with budget constraint. Subsequently, we developed two instantiations to determine risk tendency based on expert knowledge or self-supervised learning. Experimental results on real datasets demonstrate that RRLB making use of both prediction uncertainty and risk tendency achieves better cumulative performance than representative competitors.



## 7. REFERENCES

- [1] S. A. Armstrong. *A meta-analysis of randomness in human behavioral research*. Louisiana State University and Agricultural & Mechanical College, 2004.
- [2] N. Bhamidipati, R. Kant, and S. Mishra. A large scale prediction engine for app install clicks and conversions. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 167–175, 2017.
- [3] H. Cai, K. Ren, W. Zhang, K. Malialis, J. Wang, Y. Yu, and D. Guo. Real-time bidding by reinforcement learning in display advertising. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 661–670, 2017.
- [4] O. Chapelle. Modeling delayed feedback in display advertising. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1097–1105, 2014.
- [5] H.-T. Cheng, L. Koc, J. Harmsen, T. Shaked, T. Chandra, H. Aradhye, G. Anderson, G. Corrado, W. Chai, M. Ispir, et al. Wide & deep learning for recommender systems. In *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems*, pages 7–10, 2016.
- [6] Y. Feng, F. Lv, W. Shen, M. Wang, F. Sun, Y. Zhu, and K. Yang. Deep session interest network for click-through rate prediction. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 2301–2307, 2019.
- [7] A. Ghosh, S. Mitra, S. Sarkhel, and V. Swaminathan. Optimal bidding strategy without exploration in real-time bidding. In *Proceedings of the 2020 SIAM International Conference on Data Mining*, pages 298–306. SIAM, 2020.
- [8] A. Ghosh, B. I. Rubinstein, S. Vassilvitskii, and M. Zinkevich. Adaptive bidding for display advertising. In *Proceedings of the 18th international conference on World wide web*, pages 251–260, 2009.
- [9] Google. The arrival of real-time bidding and what it means for media buyers. In *Google White Paper*, 2012.
- [10] T. Graepel, J. Q. Candela, T. Borchert, and R. Herbrich. Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft’s bing search engine. Omnipress, 2010.
- [11] N. Grislain, N. Perrin, and A. Thabault. Recurrent neural networks for stochastic control in real-time bidding. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2801–2809, 2019.
- [12] H. Guo, R. Tang, Y. Ye, Z. Li, and X. He. Deepfm: a factorization-machine based neural network for ctr prediction. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 1725–1731, 2017.
- [13] X. He and T.-S. Chua. Neural factorization machines for sparse predictive analytics. In *Proceedings of the 40th International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 355–364, 2017.
- [14] X. He, J. Pan, O. Jin, T. Xu, B. Liu, T. Xu, Y. Shi, A. Atallah, R. Herbrich, S. Bowers, et al. Practical lessons from predicting clicks on ads at facebook. In *Proceedings of the Eighth International Workshop on Data Mining for Online Advertising*, pages 1–9, 2014.
- [15] Y. He, X. Chen, D. Wu, J. Pan, Q. Tan, C. Yu, J. Xu, and X. Zhu. A unified solution to constrained bidding in online display advertising. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 2993–3001, 2021.
- [16] J. Hull. *Risk management and financial institutions*, volume 733. John Wiley & Sons, 2012.
- [17] V. Krishna. *Auction theory*. Academic press, 2009.
- [18] K.-c. Lee, B. Orten, A. Dasdan, and W. Li. Estimating conversion rate in display advertising from past performance data. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 768–776, 2012.
- [19] T. J. Linsmeier and N. D. Pearson. Value at risk. *Financial Analysts Journal*, 56(2):47–67, 2000.
- [20] H. Markowitz. Harry m. markowitz. *Portfolio selection, Journal of Finance*, 7(1):77–91, 1952.
- [21] H. B. McMahan, G. Holt, D. Sculley, M. Young, D. Ebner, J. Grady, L. Nie, T. Phillips, E. Davydov, D. Golovin, et al. Ad click prediction: a view from the trenches. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1222–1230, 2013.
- [22] J. Mun. *Modeling risk: Applying Monte Carlo simulation, real options analysis, forecasting, and optimization techniques*, volume 347. John Wiley & Sons, 2006.
- [23] R. J. Oentaryo, E.-P. Lim, J.-W. Low, D. Lo, and M. Finegold. Predicting response in mobile advertising with hierarchical importance-aware factorization machine. In *Proceedings of the 7th ACM International Conference on Web Search and Data Mining*, pages 123–132, 2014.
- [24] C. Perlich, B. Dalessandro, R. Hook, O. Stitelman, T. Raeder, and F. Provost. Bid optimizing and inventory scoring in targeted online advertising. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 804–812, 2012.
- [25] S. Rendle. Factorization machines with libfm. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 3(3):1–22, 2012.
- [26] R. T. Rockafellar, S. Uryasev, et al. Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42, 2000.

- [27] Y. Shan, T. R. Hoens, J. Jiao, H. Wang, D. Yu, and J. Mao. Deep crossing: Web-scale modeling without manually crafted combinatorial features. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 255–262, 2016.
- [28] W. F. Sharpe. The sharpe ratio. *Journal of Portfolio Management*, 21(1):49–58, 1994.
- [29] I. Trofimov, A. Kornetova, and V. Topinskiy. Using boosted trees for click-through rate prediction for sponsored search. In *Proceedings of the Sixth International Workshop on Data Mining for Online Advertising and Internet Economy*, pages 1–6, 2012.
- [30] J. Wang and S. Yuan. Real-time bidding: A new frontier of computational advertising research. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*, pages 415–416, 2015.
- [31] D. Wu, X. Chen, X. Yang, H. Wang, Q. Tan, X. Zhang, J. Xu, and K. Gai. Budget constrained bidding by model-free reinforcement learning in display advertising. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 1443–1451, 2018.
- [32] W. C.-H. Wu, M.-Y. Yeh, and M.-S. Chen. Predicting winning price in real time bidding with censored data. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1305–1314, 2015.
- [33] J. Xiao, H. Ye, X. He, H. Zhang, F. Wu, and T.-S. Chua. Attentional factorization machines: learning the weight of feature interactions via attention networks. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 3119–3125, 2017.
- [34] S. Yuan, J. Wang, and X. Zhao. Real-time bidding for online advertising: measurement and analysis. In *Proceedings of the Seventh International Workshop on Data Mining for Online Advertising*, pages 1–8, 2013.
- [35] D. Zha, K.-H. Lai, K. Zhou, and X. Hu. Experience replay optimization. In *In Proceedings of the 2019 International Joint Conference on Artificial Intelligence*, page 4243–4249, 2019.
- [36] S. Zhai, K.-h. Chang, R. Zhang, and Z. M. Zhang. Deepintent: Learning attentions for online advertising with recurrent neural networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1295–1304, 2016.
- [37] H. Zhang, W. Zhang, Y. Rong, K. Ren, W. Li, and J. Wang. Managing risk of bidding in display advertising. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 581–590, 2017.
- [38] W. Zhang, T. Du, and J. Wang. Deep learning over multi-field categorical data. In *European Conference on Information Retrieval*, pages 45–57, 2016.
- [39] W. Zhang and J. Wang. Statistical arbitrage mining for display advertising. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1465–1474, 2015.
- [40] W. Zhang, S. Yuan, and J. Wang. Optimal real-time bidding for display advertising. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1077–1086, 2014.
- [41] G. Zhou, X. Zhu, C. Song, Y. Fan, H. Zhu, X. Ma, Y. Yan, J. Jin, H. Li, and K. Gai. Deep interest network for click-through rate prediction. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1059–1068, 2018.