Reserve Price Optimization in First-Price Auctions via Multi-Task Learning

Achir Kalra
Forbes Media LLC.
New York, USA
akalra@forbestech.com

Chong Wang
Amazon Ads
New York, USA
munichong@gmail.com

Cristian Borcea
New Jersey Institute of Technology
New Jersey, USA
borcea@njit.edu

Yi Chen
New Jersey Institute of Technology
New Jersey, USA
yi.chen@njit.edu

***Abstract*—Online publishers typically sell ad impressions through auctions held in ad exchanges in real-time, i.e., real-time bidding (RTB). A publisher will accept the winning bid if it is higher than a given reserve price for an ad impression. Setting an appropriate reserve price for an ad impression is critical for publishers' revenue generation, but also challenging. While this problem has been studied for second-price auctions, it lacks studies for first-price auctions, the de facto industry standard since 2019. This paper proposes a machine learning model that determines the optimal reserve prices for individual ad impressions in real-time. It uses a multi-task learning framework to predict the lower bounds of the highest bids with a coverage probability, using only the data available to publishers. The experiments using data from a large international publisher show that the proposed model outperforms the comparison systems on generating revenue.**

***Index Terms*—computational advertising, neural networks, prediction interval estimation, survival analysis, proportional hazards model**

I. INTRODUCTION

Online display advertising is the most important revenue stream of most websites, which provide free information and services in exchange for money received from display ads. In display advertising, advertisers (e.g., Volkswagen) pay publishers (e.g., Forbes) for showing banners, videos, or text on their webpages. One display of an ad in a page view is called an *ad impression*.

One of the main ad selling methods is real-time bidding (RTB). An impression triggered in real-time is sent to ad exchanges with a reserve price provided by the publisher. The reserve price is the minimum price that the publisher would be willing to accept for this impression. The advertisers bid on the impression in an auction. The winning advertiser whose bid is higher than the reserve price is allowed to show the ad on the publisher's webpage.

Second-price auctions and first-price auctions are the two main auction types used in online display advertising. To improve bidding transparency [1], since 2019, most of the display advertising market has switched from second-price auctions to first-price auctions. Unlike second-price auctions in which winners are charged the price of the second highest bid, in first-price auctions, the winning advertisers pay the prices that they bid, as long as they outbid the reserve price. If all bids are lower than the reserve price (i.e., underbid), the impression is not sold.

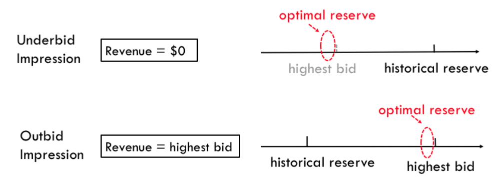

Fig. 1: Optimal Reserve Prices for Two Cases

Reserve prices directly impact revenue. Publishers' ad revenue in first-price auctions can be simply defined as $\sum_{i=0}^{N} v_i$, where N is the number of impressions and v_i is the revenue of the ith impression. Impression revenue calculation can be divided into two cases, as shown in Figure 1. Reserve prices can play an important role in both cases:

- Underbid Impressions: Underbid impressions occur when the reserve price is higher than the highest bid, resulting in zero revenue. To maximize the revenue, the publisher should lower the reserve price to match or be slightly below the highest bid. By optimizing reserve prices, publishers can obtain revenue from these impressions which are supposed to be underbid. However, the problem is difficult because the highest bids of historical underbid impressions are unknown in the transaction data at the publisher side.
- Outbid Impressions: Outbid impressions occur when the reserve price is equal or less than the highest bid, resulting in revenue equal to the highest bid. However, if the reserve price is set too low, advertisers may not be incentivized to bid higher. In such cases, advertisers who bid only slightly higher than the reserve price might win, resulting in sub-optimal impression revenue for publishers. Our analysis of preliminary data indicates a strong correlation between advertisers' bids and reserve prices. When a reserve price is outbid, there is a high likelihood that the final revenue will be only slightly higher than the reserve price. Therefore, setting an appropriate reserve price can optimize the publisher's revenue even in outbid impressions. By setting reserve prices slightly below the highest bids, advertisers can be stimulated to bid higher in the long run, within a matter of hours up to weeks. [2]

There are studies on reserve price prediction for second-price auctions done at publishers' side [3], [4], [5], [6],

[7], [8], [9], [10], [11]. None of them are suitable for first-price auctions due to different data censorship. In first price auctions, second prices are not visible and thus the relationship among the highest price, the second price, and the reserve price can no longer be utilized (Section II). The literature lacks studies on reserve price prediction for individual impressions in first-price auctions.

Predicting the optimal reserve price in first-price auctions is very challenging. First, directly predicting the highest bids is risky. In first-price auctions, the highest bids of outbid impressions are available (i.e., uncensored) to the publishers. Therefore, publishers could build a machine learning model to predict the highest bid of an impression using their historical outbid transaction data, and set the reserve price right below the predicted highest bid in order to be outbid and push the future bids higher. However, due to uncertainty in the ad market, data noise, and model mis-specification, the predicted highest bids may not be correct. This strategy is risky and often results in failing to sell the impressions. Second, in addition to outbid impressions, there are historic underbid impressions. Since their highest bids are censored (i.e., unavailable to publishers), they cannot be directly used to train a model for the highest bid prediction. Nevertheless, they carry valuable information regarding advertisers' bidding behavior. How to combine both uncensored and censored information for reserve price prediction is an open problem. Third, publishers have very limited information about users and advertisers. They do not have access to personally identifiable user information or behavior outside their own website. In contrast, advertisers and Demand-side Platforms (i.e., systems that help advertisers to buy impressions in real-time) can access much richer information about users. Their real-time bidding algorithms are black-boxes to publishers. Thus, it is not feasible to reverse-engineer these algorithms at the publisher's side for bidding price prediction.

To overcome the first challenge, instead of point estimation, we propose to use interval estimation, i.e., $[b_L, +\infty]$. That is, instead of predicting the highest bid b, we propose to predict the lower bound of the highest bid b_L, with a pre-specified confidence level, $(1-\alpha)\%$. Publishers can set the confidence level according to their revenue goal: decreasing the confidence level leads to a more aggressive strategy (higher risk), while increasing the confidence level leads to a more conservative strategy (less risk). For instance, suppose a publisher sets the risk level α to be 20%. That is, it can tolerate at most 20% likelihood that the reserve price is underbid and the ad is not sold. Given the confidence level, $(1-\alpha)\%$, i.e., 80%, suppose the model predicts $\hat{b}_L$=0.1. In this case, a publisher can set the reserve price to be \$0.1, expecting at least 80% likelihood of being outbid. To build such a model, we adapt the QD loss function [12] for computing prediction intervals, fed by outbid impressions.

To address the second challenge, we propose a multi-task learning framework, where the main task is to predict the highest bid lower bound b_L, and the auxiliary task is to predict the failure rate h of a given reserve price (i.e., the probability of being underbid). Unlike the main task that uses historical outbid impressions only, the auxiliary task leverages both historical outbid and underbid impressions, and boosts the main task's performance via shared learning parameters.

To tackle the third challenge, we use a deep neural network (DNN) to capture the complex joint effect between the features. The DNN models the features for users, pages, and ad placements using embeddings. These latent vectors learn latent features from massive historical transaction data at the publisher side.

To summarize, this paper has five main contributions. First, this is the first work that optimizes the reserve price (also known in the industry as the publisher bid) of each individual ad impression in real-time in first-price auctions. Second, we propose to use interval estimation, instead of point estimation, to quantify the uncertainty in the highest bid prediction. This allows publishers to adjust the risk level based on their business strategies. This is also the first attempt to adapt the QD method [12] for ad bidding interval estimation. Third, we propose a multi-task learning framework, based on a DNN, that predicts the highest bid lower bounds using the QD loss function and the reserve price failure rates, such that both historic outbid and underbid impressions can be utilized in model learning. The proposed method does not rely on any distribution assumptions on data. Fourth, we not only learn the bid distribution information, but also study how to set the reserve price with different risk levels to optimize the expected revenue. Finally, we trained and evaluated the proposed leaning model on real-life first-price auction impression transaction data from a large publisher, i.e., Forbes Media, which contain tens of millions of transactions. The experimental results demonstrate that the proposed model outperforms the comparison systems. It achieves *higher expected revenue* by predicting high values of reserve prices with sufficient coverage to cover advertisers' highest bids. Using the real-life RTB setting and data available at the publisher's side, the proposed method is practical and can be deployed on any online publisher platform for revenue benefits.

II. Related Work

Reserve price optimization has been studied in the past few years due to its importance to online advertising. As far as we know, Feng et al. [2] is the only study on reserve price prediction for first-price auctions. It proposes a gradient-based algorithm to adaptively optimize reserve prices based on estimates of bidders' responsiveness to publishers' announced reserve prices in experimental shocks in a synthetic dataset. There are three major differences with our study. First, [2] predicts the traditional reserve price (i.e., *public* reserve prices visible to advertisers). It does not work for *secret* reserve prices invisible to advertisers, but our method does. Secret reserve prices, also known as *publisher bids*, are a new type of reserve price for first-price auctions, which are now widely adopted by publishers. According to the econometrics literature [13], secret reserve prices are perceived favorably by sellers who would like to post a high reserve price, without discouraging participation. Our proposed model is evaluated using a dataset with secret reserve prices. Nevertheless, it can also be used to set public reserve prices if publishers prefer so and the ad exchange allows it. Publishers may periodically retrain the models to dynamically capture market changes. Second, Feng et al. [2] set a single reserve price for all impressions, whereas our model predicts a reserve price for each *individual* impression, which is the typical setting in publisher bids. Third, [2] is evaluated on a synthetic dataset, while our model is evaluated on a real-life dataset.

Other existing studies consider second-price auctions, which had been the dominant auction type until 2019. None of these techniques are suitable for first-price auctions due to different data censorship. For bidding price prediction at publishers' side, many studies [14], [5], [15], [7], [8], [9], [10], [11] assume that publishers know the highest and the second highest bids of historical impressions, with evaluation on synthetic data or ad exchange data. The assumption, and thus the solutions, are applicable only for publishers who own an ad exchange, e.g., Google and Yahoo!. Most publishers, e.g., Forbes, New York Times, etc. do not observe any bids for underbid impressions in either first- or second-price auctions. Considering data censorship at the publishers' side, Alcobendas et al. [6] propose a game-theoretic-based model to optimize reserve prices with uncensored data and left-censored data. The model considers auction details (e.g., the number of bids higher than the reserve price), typically available only for publishers that own an ad exchange. Chahuara et al. [4] uses the Aalen's additive models to estimate the first and the second bids' distributions to maximize a revenue function for second-price auctions for publishers. However, unlike in second-price auctions, the publishers do not know the second prices in first-price auctions. Kalra et al. [3] uses a parametric survival model and header bidding to optimize reserve prices with left- and right-censored data. The model relies on the assumption that the probability of no advertiser bidding higher than the reserve price follows a known distribution.

Other works on ad bidding prices study the problem from the advertisers' side, estimating the winning prices [16], [17], [18] or bid landscape to help advertisers win RTB auctions [19], [20], [21], [22]. Unlike publishers, advertisers have much more detailed data about users. They also have different data censorship. For example, advertisers always know their own bids, in addition to publishers' reserve prices. In contrast, for underbid impressions, publishers only know the reserve prices but not the advertisers' bids. Some studies [16], [17] predict winning bids and do not provide probabilistic predictions or distributional information. Another study [18] predicts distributional information considering advertisers' campaign budgets and paces of bidding over time. However, budgets and paces are not applicable at the publishers' side because they are unknown. Bid landscape prediction forecasts the distribution of advertisers' bids of an impression. For example, [19], [20] forecast the bid landscape based on statistical counts derived from segmented samples per advertisers' campaigns or by targeting impression attributes. These methods rely on a vast number of attributes about impressions (e.g., user age and gender), which are not available to publishers. The works in [21], [22] discretize bid prices into hundreds of bins and make predictions using neural networks for every bin. While effective, it is computationally expensive and thus unfeasible in real-time. In summary, the solutions for advertisers are not applicable to publishers due to the lack of data.

Another family of related research is auction theories [23], [24], which typically provide solutions under ideal settings. However, such studies cannot be applied directly in the ad industry settings and are not evaluated with real-life datasets. For example, [23] focuses on general pricing auctions and requires information about the number of bidders and their identities across auctions. In real-life, publishers do not know how many and which advertisers join the auctions. Publishers also cannot track advertisers' bidding behavior over time. The work in [24] shows that shading strategies lead to large increase of revenue for the advertisers in second-price auctions. It requires information that is not available at the publisher's side, such as the number of advertisers.

Different from all the related work, our paper propose machine learning techniques to predict the bidding prices for individual ad impressions for publishers in first-price auctions, which became the de facto industry standard since 2019. Furthermore, our work uses only data available to publishers and works well with real-life data. Thus, our solution is practical and can be used by most online publishers to increase ad revenue. In addition, our approach is general and does not rely on any distribution assumption.

III. Proposed Method

As discussed in Section I, directly predicting the exact highest bids is risky. Instead of the traditional point estimation, we propose to predict the lower bound of the bidding prices with a given confidence level, that is, the highest bids will be higher than the lower bound with the confidence level. The problem is defined as follows.

Definition 1. *[Main Task: Highest Bid Lower Bound Prediction] Given an ad impression and a risk level $\alpha\%$, predict the highest bid lower bound (i.e., the recommended reserve price) b_L such that the predicted reserve price will be outbid with at least a probability of $(1-\alpha)\%$.*

Since the highest bid lower bound prediction requires true highest bids in a model that infers the lower bounds, only historical outbid impressions are useful for this task because their highest bids are known to the publisher (i.e., uncensored). Historical underbid impressions cannot be used for this task because their highest bids are left-censored, and only the reserve prices are known. However, historical underbid impressions also carry valuable information on advertiser bidding patterns. Ignoring underbid data may lose significant information. To leverage underbid impressions, we propose the second task, reserve price failure rate prediction.

Definition 2. *[Auxiliary Task: Reserve Price Failure Rate Prediction] Given an ad impression and a reserve price,* *predict the failure rate h: how likely it is the given reserve price will be underbid.*

Predicting the highest bid lower bound b_L and predicting the failure rate h of a reserve price are closely related: given the information of an ad impression, the former is to predict the reserve price which has a likelihood of $(1-\alpha)\%$ to be outbid, while the latter is to predict the likelihood that a given reserve price will be underbid. The output of the highest bid lower bound prediction is what the publisher eventually wants. Adding failure rate prediction with shared learning parameters of the main task can boost performance of the main task by utilizing not only outbid impressions but also underbid impressions. Given these two related tasks, we propose to use a multi-task learning framework.

The loss functions of each task and the combined loss are introduced in Sections III-A, III-B, and III-C, respectively. The highest bid lower bound and the failure rate are estimated using a neural network learned from historical impressions, as presented in Section III-D. The loss of the lower bound estimation is computed using outbid impressions, while that of the failure rate estimation is computed using both outbid and underbid impressions. We discuss the selection of risk tolerance level in Section III-E.

A. Loss of Highest Bid Lower Bound Prediction

This section introduces the loss function of the main task: predicting lower bounds of highest bids with a given confidence level. We adapt the QD loss function proposed in [12] for Quality-Driven prediction interval estimation. QD has been applied in existing studies, such as [25], [26], [27] for wind power interval prediction. QD can generate high-quality prediction intervals which are as narrow as possible and meanwhile capture some specified proportion of data points, which is called the High-Quality (HQ) principle [28]. Compared with traditional prediction interval (PI) construction methods, which minimize prediction errors, QD directly improves PI quality. The constructed PIs are guaranteed to be optimal in terms of their key characteristics: width and coverage probability.

The QD loss function has two components: Mean Prediction Interval Width (MPIW) and Prediction Interval Coverage Probability (PICP). The overall loss is the sum of MPIW and PICP. MPIW measures the width of the average prediction intervals. The assumption is that a wide prediction interval (e.g., $[0,+\infty]$) is not at all informative and useful. Therefore, a good prediction interval should be as narrow as possible. QD defines a captured MPIW, denoted as $MPIW_{capt.}$, which represents the average width of prediction intervals that correctly includes the ground truth labels:

$$MPIW_{capt.} = \frac{1}{\sum_{i=1}^{n} k_i}\left(\hat{b}_{U_i} - \hat{b}_{L_i}\right)\cdot k_i \tag{1}$$

where k_i is a Boolean indicating if the ground truth label of the ith sample out of n samples is correctly captured in the estimated PI. $\hat{b}_{U_i}$ and $\hat{b}_{L_i}$ are the upper and lower bounds of the estimated PI, respectively. The higher the $MPIW_{capt.}$, the better the PI quality.

We modify the original QD loss function to use it for reserve price optimization: the upper bound included in the original loss function is canceled because the ranges of the highest bids are one-sided, i.e., the upper bound is $+\infty$. In particular, since we only care about the the lower bound, b_{U_i} is $+\infty$ (i.e., the interval is $[b_{L_i},+\infty]$). b_{U_i} is removed from Equation 1. Note that $k_i=1$ if $b_{L_i}\leqslant b_i$, where b_i is the actual highest bid:

$$MPIW_{capt.} = -\frac{\hat{b}_{L_i}k_i}{\sum_{i=1}^{n} k_i} \tag{2}$$

PICP measures the coverage probability of the estimated PIs, i.e., how many ground truth labels are correctly captured:

$$PICP = \frac{1}{n}\sum_{n}^{i=1} k_i \tag{3}$$

PICP (Equation. 3) is the most important indicator of the quality of PIs. We learn the parameters θ that can minimize $L_\theta = L(\theta|\mathbf{k},\alpha)$, where α is the risk level specified by publishers (Definition 1). This can be further represented by a binomial distribution: $L_\theta = \binom{n}{c}(1-\alpha)^c\alpha^{n-c}$, where $c=\sum_{i=1}^{n}k_i$. Using the de Moivre-Laplace theorem, it can further be approximated by a normal distribution. Therefore, the negative log likelihood is updated to:

$$-logL_\theta \propto \frac{n}{\alpha(1-\alpha)}((1-\alpha)-PICP)^2 \tag{4}$$

Putting the MPIW and PICP terms together, the loss of highest bid lower bound prediction considers both width and coverage (Equation (5)). λ is a parameter controlling the importance of PICP.

$$\begin{aligned} Loss_{qd} &= MPIW_{capt.} + \lambda\cdot PICP \\ &= -\frac{\hat{b}_{L_i}k_i}{\sum_{i=1}^{n}k_i} + \lambda\frac{n}{\alpha(1-\alpha)}max(0,(1-\alpha)-\frac{1}{n}\sum_{n}^{i=1}k_i)^2 \end{aligned} \tag{5}$$

B. Loss of Failure Rate Prediction

Section III-A used only historical outbid impressions. However, historical underbid impressions also carry important information on advertisers bidding patterns. Since their highest bids are censored, the QD estimation is not applicable. To leverage both outbid impressions and underbid impressions, we use survival analysis models to predict the failure rate of a reserve price. Survival analysis is widely used in many areas such as public health, e-commerce, credit risk, and so on for applications where the time to the event is of interest. In conjunction with the core objective, failure rate prediction can act as a auxiliary task and enhance the model's comprehension of advertisers' bidding tendencies in not only outbid impressions, but also underbid impressions.

To apply survival analysis to predict the failure rate of a given reserve price, we make the following analogy: one impression is an instance, which has a set of features. The event of interest is that all advertisers bid lower than the given reserve price r (i.e., underbid). The time to event is the reserve price. Left-censored instances are underbid impressions (only the reserve price r is known), while uncensored instances are outbid impressions (the highest bid b is known). With the increase in the reserve price (i.e., time to event), the event of interest also increases. When the reserve price is \$0, it is most likely that the event of reserve price failure does not occur. If the reserve price is \$100, it hardly receives a higher bid. Let us note that hazard rate in survival analysis is referred as failure rate in our application.

We propose a loss function for failure rate prediction incorporating the loss function of a Cox's proportional hazards model (the Cox PH model) [29] in Equation 6. It gives an expression for the hazard at time t for an individual with a given specification of a set of explanatory variables.

$$h(t, X_i) = h_0(t) e^{\hat{y}_i} \tag{6}$$

The Cox PH model consists of two parts. One is the underlying baseline hazard function, $h_0(t)$. This is the cumulative hazard rate, i.e., the percentage of the training instances whose events have already occurred at t. $h_0(t)$ describes how the risk of an event per time unit changes over t at baseline levels of explanatory variable. The second part is the exponential expression $e^{\hat{y}_i}$, which is computed from trainable parameters θ and explanatory variables X_i. $\hat{y}_i$ describes how the hazard varies in response to explanatory variables. The output $h(t, X_i)$ is the hazard rate of X_i at time t.

$h_0(t)$ is an unspecified function, which can be computed from existing observations, without any assumption of the distribution of the baseline hazard. This property makes the Cox PH model a semi-parametric model, a key reason that the Cox PH model is widely used. The fact that the model does not make any assumption about the distribution of the baseline hazard is important in online display advertising because the ad market is very complex and dynamic. Thus, the highest bids may not always be drawn from a specific distribution.

To handle both uncensored and censored data, θ are estimated using the Cox PH partial likelihood function. The partial likelihood function considers probabilities only for outbid impressions and does not explicitly consider those for underbid ones. In particular, for an underbid impression A_i whose reserve price is r_i, we find all outbid impressions A_j whose highest bid is b_j, where $b_j \geqslant r_i$. It is known that A_i was underbid before r_i and A_j was not underbid at r_i. This is the only case in our data, in which the relationship of two censored impressions is known, i.e., A_j is more valuable than A_i. The goal is to find θ that can maximize $h(r_i, X_i) - h(r_i, X_j)$. Therefore, the partial likelihood of the event at a price r_i is:

$$L_i = \frac{h(r_i, X_i)}{\sum_{j: b_j \geqslant r_i} h(r_i, X_j)} = \frac{h_0(r_i) e^{\hat{y}_i}}{\sum_{j: b_j \geqslant r_i} h_0(r_i) e^{\hat{y}_i}} = \frac{e^{\hat{y}_i}}{\sum_{j: b_j \geqslant r_i} e^{\hat{y}_i}} \tag{7}$$

Treating the impressions as if they are statistically independent, the joint probability of all uncensored cases is $L_\theta = \prod_{A_i \in U} L_i$, where U is the set of underbid impressions. The loss of the failure rate prediction (negative log partial likelihood) is in Equation 8:

$$Loss_{cox} = \sum_{A_i \in U} (log \sum_{j: b_j \geqslant r_i} e^{\hat{y}_i} - \hat{y}_i) \tag{8}$$

C. Loss of Multi-task Learning

Incorporating Equations (5) and (8) together, the loss function for the multi-task learning is in Equation 9.

$$\begin{aligned} Loss &= Loss_{qd} + \mu Loss_{cox} \\ &= -\frac{\hat{b}_{L_i} k_i}{\sum_{i=1}^{n} k_i} + \lambda \frac{n}{\alpha(1-\alpha)} max(0, (1-\alpha) - \frac{1}{n} \sum_{n}^{i=1} k_i)^2 \\ &+ \mu \sum_{A_i \in U} (log \sum_{j: b_j \geqslant r_i} e^{\hat{y}_i} - \hat{y}_i) \end{aligned} \tag{9}$$

where μ is a parameter controlling the importance of the failure rate prediction, and $\hat{b}_{L_i}$ and $\hat{y}_i$ are computed from the mapping of θ and features X.

D. Predicting Highest Bid Lower Bounds and Failure Rates

The next task is to predict $\hat{b}_{L_i}$ and $\hat{y}_i$ that minimize the loss function defined in Equation 9. As discussed in Section I, it is very challenging to predict advertisers' bidding behavior based on data available at the publisher's side. Most advertisers either collect data about the users or buy such data from third-party companies. They use massive user datasets and build algorithms to determine their bids in real-time. However, these algorithms are black-boxes to publishers, who do not have detailed user data nor access to advertisers' bidding algorithms. Our aim is to provide a practical solution that works for publishers, using their limited available information.

The highest bid reflects the value of an ad impression. Advertisers will bid high if they believe the ad impressions will benefit their advertising campaigns. The value of an impression is determined by four factors. 1) User interest. If advertisers think the user has interests in their products, they are willing to pay more for the impression. Our model uses the following user features available to publishers from the user cookies: user IDs, state-level locations, operating systems, Internet browser types, network bandwidths, and devices. Unlike advertisers, publishers typically do not have access to user personal data. 2) Ad placement. Studies show that an ad at the top of a page is typically much more viewable than one at the bottom [30] and attracts higher bids. We consider two ad placement features: ad unit size (e.g., 123x324 in pixels) and ad position. 3) Page information. An ad opportunity in an article about electronic products may be more attractive than one in a political article because the user who reads the former is more likely to have a shopping intent. We consider the following page features: page URLs, channels (e.g., business, lifestyle), sub-channels/sections, and the trending status of the page labeled by the publishers' editors. 4) Context. Context features include hour of the day and referrer URLs, i.e., in which page the request for the current page originated.

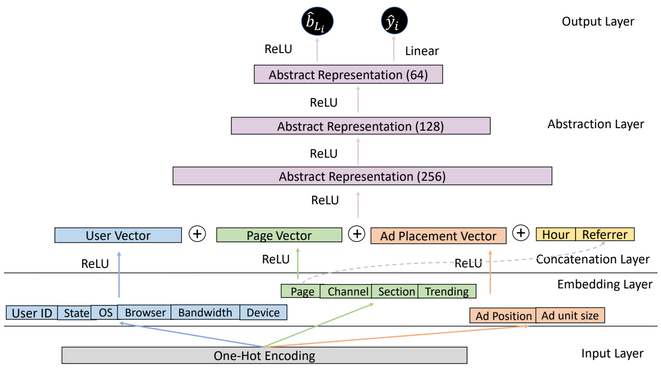

Fig. 2: Architecture of the Multi-task Learning Network

Due to the complexity of advertisers' real-time bidding algorithms and the uncertainty in the ad market, we propose to learn latent features, instead of using explicit features, from historical data. It is also important to capture interactions among latent features. Thus, simple linear models are not sufficient for this task. We propose to use deep neural networks, which have been proven successful for such situations.

To reduce model complexity and improve prediction performance, we adopt a multi-task learning framework, which predicts $\hat{b}_{L_i}$ and $\hat{y}_i$ using a set of shared parameters. Figure 2 presents the architecture of the proposed learning framework. In the following, we discuss each layer in the architecture.

Input Layer: All input features are either categorical or can easily be converted to categorical variables. For instance, the ad unit size is converted to a string feature, such as "747x413". Since a publisher usually has a fixed number of standard ad unit sizes on page templates, using one-hot encoded string features captures the interaction between width and height.

Embedding Layer: User, page, and ad placement features are represented by latent embedding vectors, which can significantly enhance model expressibility and the ability of capturing sophisticated ad transaction data. The embedding layer retrieves feature embeddings based on the one-hot input. The lengths of the embeddings are pre-specified parameters that can be tuned empirically. For each factor (i.e., user, page, and ad placement), the feature embeddings are concatenated and then mapped by the commonly-used ReLU function into a unified vector (i.e., a user vector, a page vector, and an ad placement vector). This is written as $f(x) = max(0, x)$, returning 0 for a negative input, and returning the actual value for a positive input. It allows the model to account for non-linearities and interactions. In theory, ReLU is able to approximate any function. Our framework uses it to mimic advertisers' black-box bidding algorithms.

Concatenation Layer: The user, page, ad placement, and context vectors are concatenated together. To reduce overfitting and decrease training cost, Page and referrer share the same embedding matrix.

Abstraction Layer: The concatenated vector is mapped to denser representations by fully connected layers with ReLU activation.

Output Layer: Two values are output: 1) the lower bound of the highest bid $\hat{b}_{L_i}$, whose range is $[0, +\infty]$. Since the range of ReLU is $[0, +\infty]$, a ReLU function is used as the activation function. 2) the exponential of the Cox PH model $\hat{y}_{U_i}$, whose range is $[-\infty, +\infty]$. Hence, a linear function is used as the activation function. $\hat{b}_{L_i}$ and $\hat{y}_{U_i}$ are then fed into the final loss function (Equation 9).

E. Risk Level Selection

The proposed model allows publishers to specify a risk level $\alpha \in (0, 1)$, which is the percentage of impressions that are tolerated to be underbid, and thus unsold. α is used in the loss function of the highest bid lower bound prediction (Equations 4, 5, and then 9). A high risk level will cause many underbid impressions. On the other hand, a low risk level may result in low prices of the highest bids in the long run, since advertisers will adjust their biding prices according to the reserve price. The optimal risk level can be determined based on the publisher's strategy (i.e., aggressive or conservative). It can also be set empirically by A/B testing (i.e., select the one that maximizes the total revenue over a period of time).

IV. Evaluation

A. Data and System Implementation

The data used for evaluation are first-price auctions collected on the Forbes Media's website in early 2021. We used a sample of one-month data for preliminary analysis, including feature selection. For sparse categorical features we only keep the highly frequent values. The infrequent ones are all assigned to a special value, i.e., ⟨feature_name⟩_others. To avoid data leakage and feature churn in sparse categorical features, we keep the preliminary analysis dataset and the evaluation dataset separated. After the month for the preliminary dataset collection, we collect the evaluation dataset in the following four days. The evaluation dataset contains nearly 60 million impressions on average per day, in which the ratio of outbid impressions and underbid impressions is about 3:2.

Both outbid and underbid impressions are available for model training, but only outbid impressions are used for evaluation due to the availability of the highest bids as ground truth. A model is trained using the impressions in one day and tested on the outbid impressions of the next day. As we have four-day data, each model is tested on the second, third, and fourth day, respectively.

Our model, denoted as QD+Cox, is implemented using Tensorflow. The experiments are run on a desktop with an i7 3.60Hz CPU, 32GB RAM, and an NVIDIA GeForce GTX 1060 6G GPU. The training goal is to minimize the total loss in Equation 9. Since the training dataset fits the memory, we adopt the Stochastic Gradient Descent (SGD) optimizer with a learning rate of 10^{-3}. The training batch size is set to 256. To avoid overfitting, across all 10 epochs, the best model on the validation data is applied to the test data. Unless otherwise specified, the following parameter values are used in the experiments. The risk level α is set to 30%. Thus, the prediction intervals, i.e., $[\hat{b}_L, +\infty]$, are expected to cover the highest bids of at least 70% impressions. The parameter λ is set to 10, and μ is set to 0.1. The widths of the embeddings are empirically set to 128. The abstraction layers are empirically set to 256, 128, and 64. We adapt three abstraction layers.

B. Evaluation Metrics

Predicting the highest bid lower bounds b_L with a risk level α is a special case of prediction interval estimation, where the upper bound is $+\infty$. Thus, we adapt an evaluation metric in the HQ principle [28], [12]: **Prediction Interval Coverage Probability (PICP)**. PICP, defined in Equation 3, measures the coverage of the estimated PIs. Its values are expected to be more than $(1-\alpha)\%$.

We also use the **Median Outbid Reserve Price (MORP)** metric, which measures the median of the predicted prices of all testing outbid impressions. Publishers require a tight lower bound of the highest bid in order to set the reserve price as high as possible for revenue optimization. MORP reflects how high the predicted reserve prices are. We use MORP instead of MPIW (defined in Equation 2) because MPIW is less intuitive. This is because MPIW is a negative value in our application, since the upper bound is $+\infty$. In addition, as the distribution of the reserve prices and the highest bids are highly left skewed, the median is a better measure than the mean in our application. Formally, MORP is the median of all $\hat{b}_L$ which $\hat{b}_{L_i} \leqslant b_i$ (i.e., $median(\left\{\hat{b}_{L_i} | \hat{b}_{L_i} \leqslant b_i\right\})$). MORP has the same spirit as MPIW, but uses median of the predicted prices instead of mean, and then takes the opposite value (i.e., the additive inverse). Thus, MORP is a positive value. Higher MORP means higher optimal reserve prices and higher potential revenue. In our application, a higher MORP indicates that higher reserve prices, if outbid, achieve higher revenue. On the other hand, a higher MORP also means higher a lower bound, thus the reserve price is more likely to fail to be outbid. A higher PICP means a lower lower bound of the prediction intervals, thus the reserve price is more likely to be outbid. Increasing PICP often results in decreasing MORP, and vice versa.

To address this issue, we use **Covered Outbid Reserve Price (CORP)**, as a unified evaluation metric, defined based on both PICP and MORP. A good model should achieve high CORP by balancing PICP and MORP. PICP signals the percentage of the outbid impressions among all impressions, and MORP reflects the median reserve price of the outbid impressions. CORP is defined as $CORP = (PICP \cdot N) \cdot MORP \sim PICP \cdot MORP$, where N is the total number of testing impressions. We omit N in the calculation of the metric because it is reliant on the size of the testing data and does not reflect the model performance. Intuitively, CORP is MORP weighted by PICP, i.e., the average predicted reserve prices of the testing outbid impressions whose predicted reserve prices are below the observed highest bids. CORP can also serves as a lower bound of publishers' revenue due to the following reasons: 1) the reserve prices calculated in CORP are generally less than the highest bids (i.e., impression revenue); 2) the ad revenue of the underbid impressions cannot be included in offline evaluation because their highest bids are unknown. The higher the CORP, the better the model.

Since the highest bids are available only for the outbid impressions on the publisher's side, we use the outbid impressions to perform the evaluation. During the offline evaluation phase, it is not possible to accurately calculate the actual revenue increase before conducting an online A/B test. This is because we can only calculate the revenue of previously outbid impressions, and not the effect of the underbid impressions due to the absence of highest bids. However, we consider any transition from underbid to outbid as a win since the revenue of historical underbid impressions is zero, indicating already the worst case.

C. Comparison Systems

As discussed in Section II, this is the first study that predicts individual impression's reserve price in first-price auctions. The prediction uses only the data available to publishers. Since there is no existing system addressing this problem to compare with, we use three widely-used methods of prediction interval (PI) estimation, namely MVE, Bootstrap and LUBE as comparison systems. Since our system (called *QD+Cox* in experiments) uses QD [12], we also compare against this method. All systems use different loss functions with the same set of features, which are constructed using the deep neural network proposed in Section III-D.

MVE: This method of constructing PIs has been used in several applications, such as wind power forecasts [31]. It assumes that errors are normally distributed around the true mean of targets, $y(x)$. It also assumes the dependence of the target variance on the set of inputs. It estimates the target variance using a dedicated neural network, whose outputs are the predicted mean $\hat{\mu}$ and the predicted variance $\hat{\sigma}^2$ of the normal distribution. The final reserve price is $\hat{b}_{L_i}$, which makes $\Phi(\hat{b}_{L_i}) = \alpha$; Φ is the cumulative distribution function (CDF) of the standard normal distribution.

Bootstrap: This method builds B neural network models using different subsets of the parameter space, and then makes collective decisions by the ensemble of neural networks [32]. The predicted mean is $\hat{y} = \sum_{B}^{h=1} \hat{y}_h$. The predicted variance $\hat{\sigma}^2_{\hat{y}} = \frac{1}{B-1} \sum_{B}^{h=1} (\hat{y}_h - \hat{y})$. One separate neural network is built to estimate the variance of errors $\hat{\sigma}^2_{\epsilon}$. Once both $\hat{\sigma}^2_{\hat{y}}$ and $\hat{\sigma}^2_{\epsilon}$ are known, the ith PI with a confidence level of $(1-\alpha)\%$ can be constructed: $\hat{y} \pm t_{1-\frac{\alpha}{2},df} \sqrt{\hat{\sigma}^2_{\hat{y}} + \hat{\sigma}^2_{\epsilon}}$, where $t_{1-\frac{\alpha}{2},df}$ is the $1-\frac{\alpha}{2}$ quantile of a cumulative t-distribution function with df degrees of freedom. df is defined as the difference between the number of training samples and the number of parameters of neural networks. The final reserve price is the lower bound: $\hat{y} - t_{1-\frac{\alpha}{2},df} \sqrt{\hat{\sigma}^2_{\hat{y}} + \hat{\sigma}^2_{\epsilon}}$. This method was used in applications such as healthcare [33] and solar energy [34].

LUBE: This method [28] was developed based on the HQ principle (described in Section III-A), which is used in applications such as wind energy prediction [35] and sediment load estimation [36]. It considers PICP and normalized MPIW (NMPIW). NMPIW is equal to MPIW divided by

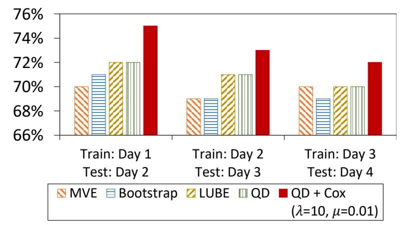

(a) PICP

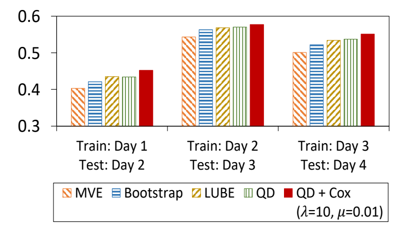

(b) MORP

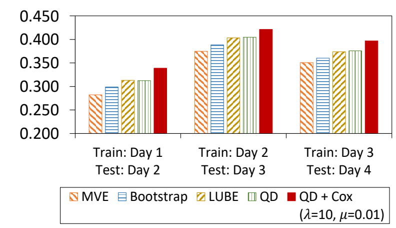

(c) CORP

Fig. 3: Performance Comparison for the Highest Bid Lower Bound Prediction

the range of the underlying target. Since the target range in our application is infinite, in the experiments, we use MPIW instead. LUBE attempts to minimize the coverage width-based criterion (CWC) for evaluation of PIs: $CWC = MPIW(1 + \gamma(PICP)e^{-\eta(PICP-\mu)})$, where $MPIW = \frac{1}{n}\sum_{i=1}^{n}\hat{b}_{U_i} - \hat{b}_{L_i}$, PICP is the same as the one in the proposed method (Equation 3). The constants η and μ are two hyperparameters determining how much penalty is assigned to PIs with a lower coverage probability. $\gamma(PICP)$ is a step function that is 1 if $PICP \geq \mu$; otherwise 0. Since LUBE is not differentiable everywhere, it uses Simulated Annealing (SA) as the training method.

QD: QD is a quality-driven distribution-free loss function proposed in [12] and based on LUBE. It was discussed in Section III-A, with definition in Equation (5).

QD+Cox: This is our proposed model that adapts QD and uses a multi-task learning framework, discussed in Section III.

D. Performance of Highest Bid Lower Bound Prediction

Figure 3 shows the comparison of all systems on PICP, MORP and CORP of the outbid impressions. Although the performance metrics are varying on different days due to daily data variance, our method QD+Cox performs best across all metrics. We observe that it significantly outperforms QD in PICP, achieves much higher CORP, and has a slightly better MORP. This demonstrates the significant benefit of multi-task learning with reserve price failure rate prediction. Adding failure rate prediction enables the model to leverage not only outbid impressions but also underbid impressions. Intuitively, outbid impressions provide information about the highest bids in the past, while underbid impressions show where the model previously underestimated the reserve price. Learning from both correct and incorrect predictions is important for improving the model's performance. From a technical perspective, QD is capable of accurately predicting whether a PI, with a reserve price as its lower-bound, will include the highest observed bid of past outbid impressions. In contrast, the failure rate prediction (Cox) seeks to understand the likelihood of a reserve price exceeding the highest bid. These two tasks are interconnected, and Cox can serve as a beneficial addition to the primary objective. This is the intrinsic reason that allow our method to outperform the model with QD alone.

For further analysis, Figure 3a shows the PICPs of most systems are higher than 70% when $\alpha = 30\%$, which satisfies the minimum coverage requirement. In other words, the highest bids of 70% impressions are higher than or equal to the corresponding $\hat{b}_L$. The exceptions are the PICPs of MVE on test day 3 and Bootstrap on test days 3 and 4, which are 69%. MVE and Bootstrap are inferior to the other models on MORP, as well (Figure 3b). Surprisingly, they do not get high MORP in return for their low PICPs. This indicates that their predicted $\hat{b}_L$ are mostly low, thus resulting in low MORP. However, for certain impressions, the predicted $\hat{b}_L$ values are too high to be outbid, thus resulting in low PICPs. In theory, MVE has a strong assumption that the variance of the highest bids follows Gaussian distribution. This is an improper assumption in our application because bids are skewed to the lower left tail. On the other hand, Bootstrap does not have any assumption. However, the limited number of bootstrap neural networks and the potential bias lead to an inaccurate estimation of the model mis-specification variance [37]. This may lead to PIs being either too wide or too narrow. Finally, As LUBE and QD have similar objectives [12], they have similar performance on our data, with QD being slightly better on average.

E. Performance with Varying λ and μ

Figure 4 shows the PICP, MORP, and CORP of our method for different combinations of λ and μ. The PICPs of all combinations are higher than 70%, which satisfies the minimum coverage requirement. High λ awards the model to focus more on the PICP part. In particular, high λ makes $Loss_{qd}$ more sensitive to the PICP part in Equation 5. When the coverage $\frac{1}{n}\sum_{n}^{i=1} k_i$ is low, the minimum coverage $max(0, (1-\alpha) - \frac{1}{n}\sum_{n}^{i=1} k_i)^2$ becomes positive. With a higher λ, small increases of its values can increase $Loss_{qd}$ more and thus penalize the training process more. Likewise, high μ stimulates the model to predict failure rate more accurately. Accurate failure rate prediction makes the model classify better if a reserve price will be outbid, which is related to what PICP measures. Thus, increasing λ and μ promotes coverage, thereby enhancing PICP.

Our method receives the best MORP when $\lambda = 10$ and $\mu = 0.01$ across all three days. In terms of CORP, QD+Cox(10, 0.01) performs best in the first two days, while QD+Cox(20, 0.001) in the last one, as its PICP in the last day is much higher. On average, QD+Cox(10, 0.01) has the

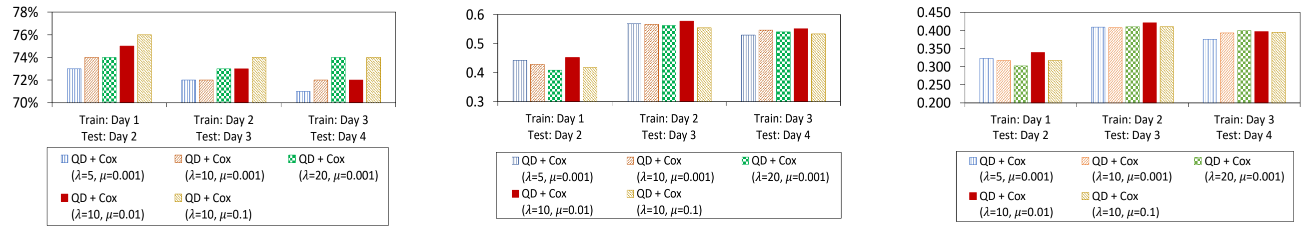

(a) PICP (b) MORP (c) CORP

Fig. 4: Performance of our Method for Different Values of λ and μ

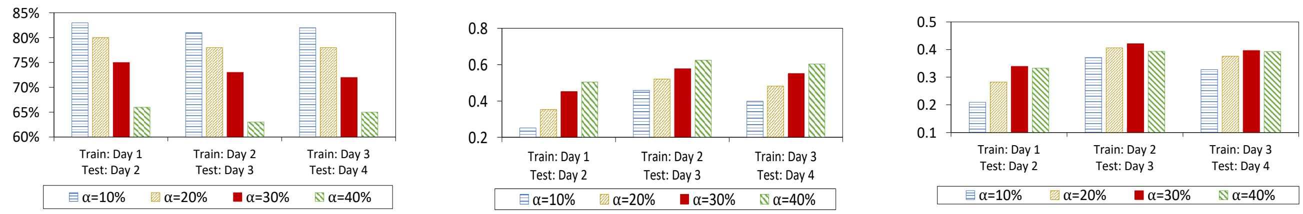

(a) PICP (b) MORP (c) CORP

Fig. 5: Performance of our Method for Different Values of Risk Level α

highest mean CORP over three test days. Based on all these results, $\lambda = 10$ and $\mu = 0.01$ are the recommended parameters in practice for our dataset.

F. Performance with Different Risk Levels

In this experiment, we evaluate the model performance of QD+Cox(10, 0.01) with different risk level α, which determines the required model coverage. In theory, there is often an inverse relationship between a model's coverage (i.e., PICP) and its lower bounds (i.e., MORP, final reserve prices). Reducing the lower bounds (i.e., the reserve price) increases the chance of covering the actual highest bids, and thus result in more outbid impressions. Conversely, a higher reserve price likely results in more underbid impressions. Thus, for publishers, setting α is a trade-off between selling more impressions and motivating advertisers to bid higher in the long-term. The results of this experiment are shown in Figure 5.

Figure 5a shows that, when α=30% and 40%, the PICPs of QD+Cox(10, 0.01) can satisfy the required minimum coverage (i.e., 70% and 60%, respectively). However, when α=10% and 0%, the PICPs on the test data (except for α=20% on Day 2) are less than the required minimum coverage (i.e., 90% and 80%, respectively). The reason is that covering more than 80% on test data is very difficult due to the high variance and uncertainty in the data. Our objective function (Equation 9) includes PICP. Thus, the model training is expected to maximize the coverage. However, the coverage may be lower than the minimum required, if the data is noisy and/or α is set too low). Thus, the publishers need to be aware of this issue and set α accordingly.

Figures 5b and 5c present the MORPs and CORPs for different α values. We observe that α=40% has the highest MORPs. This is because the associated required minimum coverage is only 60%, and thus the model has large room to push the predicted lower bounds higher. The model can give up more impressions in order to pursue higher lower bounds (i.e., predicted reserve prices). In contrast, when α=10%, to reach the minimum coverage 90%, the model has a low MORP. This low value is due to prioritizing coverage instead of the reserve price, which forces the model to shrink the predicted lower bounds significantly. CORP is highest when α=30%. In this case, although the individual reserve prices are lower (i.e., MORPs are lower), a lot more impressions are outbid (i.e., PICPs are higher), compared with the situation When α=40%, thus achieving a higher CORP.

As we can see, setting α is a trade-off between harvesting many outbid impressions and increasing the unit price of impressions. A high α causes many underbid impressions. However, once an impression gets outbid by advertisers, the publisher can earn more revenue. On the other hand, a low α leads to many sold impressions, but lower reserve prices cannot motivate advertisers to make their bids higher, as advertisers' RTB algorithms may quickly learn that high bids are unnecessary and will try to bid lower the next time.

We also studied the impact of the training sizes on the model performance, which for space reasons is not included in the paper. Experiments show that training on one, two, or three days generates similar results. This indicates that publishers, especially those with limited computational resources, may only need to train on the previous day to get good performance.

V. Conclusions

Since the display advertising industry switched to first-price auctions, accurate estimations of the highest bids of future impressions can be used by publishers to set the reserve prices to optimize the outbid rate and to motivate advertisers to bid higher in the future. This study proposes a model to predict the highest bid lower bound, which can be used as the recommended reserve price for publishers, given a risk tolerance level α. The actual highest bids have a likelihood of $(1-\alpha)\%$ to be higher than the predicted lower bounds. Our multi-task learning model predicts the failure rate of a reserve price, using both historic outbid ad impressions and underbid ad impressions. Learning these closely related tasks with shared model parameters enables better data utilization and boosts model performance. Furthermore, the model uses deep neural networks to capture user, page, ad placement, and context features, as well as the interactions among these features. The experiments on data from a large online publisher show that the proposed method significantly outperforms the comparison systems.

ACKNOWLEDGEMENT

This work was partially supported by the National Science Foundation under Grants No. CNS 2237328 and DGE 2043104, as well as by the Martin Tuchman '62 Chair Endowment and the Leir Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

REFERENCES

[1] S. Despotakis, R. Ravi, and A. Sayedi, "First-price auctions in online display advertising," *Available at SSRN 3485410*, 2019.
[2] Z. Feng, S. Lahaie, J. Schneider, and J. Ye, "Reserve price optimization for first price auctions in display advertising," in *ICML'21*. PMLR, 2021, pp. 3230–3239.
[3] A. Kalra, C. Wang, C. Borcea, and Y. Chen, "Reserve price failure rate prediction with header bidding in display advertising," in *KDD'19*, 2019.
[4] P. Chahuara, N. Grislain, G. Jauvion, and J.-M. Renders, "Real-time optimization of web publisher rtb revenues," in *KDD'17*, 2017, pp. 1743–1751.
[5] A. M. Medina and M. Mohri, "Learning theory and algorithms for revenue optimization in second price auctions with reserve," in *ICML'14*, 2014, pp. 262–270.
[6] M. A. Alcobendas Lisbona, S. Chammas, and K.-c. Lee, "Optimal reserve prices in upstream auctions: Empirical application on online video advertising," in *KDD'16*, 2016, pp. 1395–1404.
[7] Z. Xie, K.-C. Lee, and L. Wang, "Optimal reserve price for online ads trading based on inventory identification," in *The AdKDD Workshop in conjunction with KDD'17*, 2017, p. 6.
[8] N. Cesa-Bianchi, C. Gentile, and Y. Mansour, "Regret minimization for reserve prices in second-price auctions," *IEEE Transactions on Information Theory*, vol. 61, no. 1, pp. 549–564, 2015.
[9] D. Austin, S. Seljan, J. Monello, and S. Tzeng, "Reserve price optimization at scale," in *DSAA'16*, 2016, pp. 528–536.
[10] A. Munoz and S. Vassilvitskii, "Revenue optimization with approximate bid predictions," in *NIPS'17*, 2017, pp. 1858–1866.
[11] G. Jauvion, N. Grislain, P. Dkengne Sielenou, A. Garivier, and S. Gerchinovitz, "Optimization of a ssp's header bidding strategy using thompson sampling," in *KDD'18*, 2018, pp. 425–432.
[12] T. Pearce, A. Brintrup, M. Zaki, and A. Neely, "High-quality prediction intervals for deep learning: A distribution-free, ensembled approach," in *ICML'18*. PMLR, 2018, pp. 4075–4084.
[13] P. Jehiel, L. Lamy *et al.*, "Absolute auctions and secret reserve prices: Why are they used?" *Levine's working paper archive*, 2011.
[14] J. Li, X. Ni, Y. Yuan, R. Qin, X. Wang, and F.-Y. Wang, "The impact of reserve price on publisher revenue in real-time bidding advertising markets," in *IEEE SMC'17*, 2017, pp. 1256–1261.
[15] S. Yuan, J. Wang, B. Chen, P. Mason, and S. Seljan, "An empirical study of reserve price optimisation in real-time bidding," in *KDD'14*, 2014, pp. 1897–1906.
[16] W. Zhang, T. Zhou, J. Wang, and J. Xu, "Bid-aware gradient descent for unbiased learning with censored data in display advertising," in *KDD'16*, 2016, pp. 665–674.
[17] W. Wu, M.-Y. Yeh, and M.-S. Chen, "Deep censored learning of the winning price in the real time bidding," in *KDD'18*, 2018, pp. 2526–2535.
[18] W. C.-H. Wu, M.-Y. Yeh, and M.-S. Chen, "Predicting winning price in real time bidding with censored data," in *KDD'15*, 2015, pp. 1305–1314.
[19] Y. Cui, R. Zhang, W. Li, and J. Mao, "Bid landscape forecasting in online ad exchange marketplace," in *KDD'11*, 2011, pp. 265–273.
[20] Y. Wang, K. Ren, W. Zhang, J. Wang, and Y. Yu, "Functional bid landscape forecasting for display advertising," in *ECML PKDD'16*, 2016, pp. 115–131.
[21] A. Ghosh, S. Mitra, S. Sarkhel, J. Xie, G. Wu, and V. Swaminathan, "Scalable bid landscape forecasting in real-time bidding," in *ECML-KDD'19*. Springer, 2019, pp. 451–466.
[22] K. Ren, J. Qin, L. Zheng, Z. Yang, W. Zhang, and Y. Yu, "Deep landscape forecasting for real-time bidding advertising," in *KDD'19*, 2019, pp. 363–372.
[23] S. Chawla, J. D. Hartline, and R. Kleinberg, "Algorithmic pricing via virtual valuations," in *Proceedings of the 8th ACM Conference on Electronic Commerce*, 2007, pp. 243–251.
[24] T. Nedelec, M. Abeille, C. Calauzènes, N. El Karoui, B. Heymann, and V. Perchet, "Thresholding the virtual value: a simple method to increase welfare and lower reserve prices in online auction systems," *arXiv preprint arXiv:1808.06979*, 2018.
[25] T. S. Salem, H. Langseth, and H. Ramampiaro, "Prediction intervals: Split normal mixture from quality-driven deep ensembles," in *UAI'20*. PMLR, 2020, pp. 1179–1187.
[26] J. Hu, Y. Lin, J. Tang, and J. Zhao, "A new wind power interval prediction approach based on reservoir computing and a quality-driven loss function," *Applied Soft Computing*, vol. 92, p. 106327, 2020.
[27] F. Liu, C. Li, Y. Xu, G. Tang, and Y. Xie, "A new lower and upper bound estimation model using gradient descend training method for wind speed interval prediction," *Wind Energy*, 2020.
[28] A. Khosravi, S. Nahavandi, D. Creighton, and A. F. Atiya, "Lower upper bound estimation method for construction of neural network-based prediction intervals," *IEEE Transactions on Neural Networks*, vol. 22, no. 3, pp. 337–346, 2010.
[29] D. G. Kleinbaum and M. Klein, *Survival Analysis*. Springer, 2010.
[30] C. Wang, A. Kalra, L. Zhou, C. Borcea, and Y. Chen, "Probabilistic models for ad viewability prediction on the web," *TKDE*, vol. 29, no. 9, pp. 2012–2025, 2017.
[31] A. Khosravi and S. Nahavandi, "An optimized mean variance estimation method for uncertainty quantification of wind power forecasts," *International Journal of Electrical Power & Energy Systems*, vol. 61, pp. 446–454, 2014.
[32] A. Khosravi, S. Nahavandi, D. Srinivasan, and R. Khosravi, "Constructing optimal prediction intervals by using neural networks and bootstrap method," *IEEE Transactions on neural networks and learning systems*, vol. 26, no. 8, pp. 1810–1815, 2014.
[33] Y. Liao, L. Zhang, and C. Liu, "Uncertainty prediction of remaining useful life using long short-term memory network based on bootstrap method," in *ICPHM'18*. IEEE, 2018, pp. 1–8.
[34] K. Li, R. Wang, H. Lei, T. Zhang, Y. Liu, and X. Zheng, "Interval prediction of solar power using an improved bootstrap method," *Solar Energy*, vol. 159, pp. 97–112, 2018.
[35] P. Wen, S. Zhang, Y. Xing, L. Huo, and N. Bohlooli, "A novel method based on lower–upper bound approximation to predict the wind energy," *Journal of Cleaner Production*, vol. 259, p. 120458, 2020.
[36] X.-Y. Chen and K.-W. Chau, "Uncertainty analysis on hybrid double feedforward neural network model for sediment load estimation with lube method," *Water Resources Management*, vol. 33, no. 10, pp. 3563–3577, 2019.
[37] A. Khosravi, S. Nahavandi, D. Creighton, and A. F. Atiya, "Comprehensive review of neural network-based prediction intervals and new advances," *IEEE Transactions on Neural Networks*, vol. 22, no. 9, pp. 1341–1356, 2011.