

Simultaneous Optimization of Bid Shading and Internal Auction for Demand-Side Platforms

Yadong Xu¹, Bonan Ni¹, Weiran Shen², Xun Wang¹,
Zichen Wang³, Yinsong Xue³, Pingzhong Tang¹

¹Institute for Interdisciplinary Information Sciences, Tsinghua University

²Gaoling School of Artificial Intelligence, Renmin University of China

³ByteDance

xuyd17@mails.tsinghua.edu.cn, bricksern@gmail.com, shenweiran@ruc.edu.cn, wxhelloworld@outlook.com,
wangzichen.118ada@bytedance.com, xueyinsong@bytedance.com, kenshinping@gmail.com

Abstract

Online advertising has been one of the most important sources for industry’s growth, where the demand-side platforms (DSP) play an important role via bidding to the ad exchanges on behalf of their advertiser clients. Since more and more ad exchanges have shifted from second to first price auctions, it is challenging for DSPs to adjust bidding strategy in the volatile environment. Recent studies on bid shading in first-price auctions may have limited performance due to relatively strong hypotheses about winning probability distribution. Moreover, these studies do not consider the incentive of advertiser clients, which can be crucial for a reliable advertising platform. In this work, we consider both the optimization of bid shading technique and the design of internal auction which is ex-post incentive compatible (IC) for the management of a DSP. Firstly, we prove that the joint design of bid shading and ex-post IC auction can be reduced to choosing one monotone bid function for each advertiser without loss of optimality. Then we propose a parameterized neural network to implement the monotone bid functions. With well-designed surrogate loss, the objective can be optimized in an end-to-end manner. Finally, our experimental results demonstrate the effectiveness and superiority of our algorithm.

Introduction

Over the past two decades, online advertising has been proven to be one of the most effective methods of monetization for internet companies. Typically, for each of these companies, there is a main auction mechanism that is in charge of allocating relevant ads to target users and collecting payments from the winning advertisers. In practice, different mechanisms such as VCG (Vickrey–Clarke–Groves) (Vickrey 1961; Clarke 1971; Groves et al. 1973) and GSP (generalized second-price) (Edelman, Ostrovsky, and Schwarz 2007) and their variations are widely adopted and advertisers exhibit different bidding behaviors in different mechanisms.

For ad platforms such as Google or Tiktok, there are billions of ad auctions, each executed within a few milliseconds, conducted daily. With such stringent latency requirements, ad exchanges have emerged as a popular form of online advertising marketplaces in real time. Specifically, an

ad exchange runs auctions to sell user impressions from diverse publishers to advertisers. As can be imagined, it is challenging to bid well for each individual advertiser in such complex ad exchanges. As a result, the DSPs (demand-side platforms) are widely adopted to help advertiser clients target their desired users and optimize the bid price for corresponding valuable impressions. Besides, the DSPs allow advertisers to set high-level goals (e.g., effective cost per action (eCPA)), and then bid on behalf of the managed advertisers in the ad exchange. In this paper, we consider the problem of optimizing trade-offs between surplus and profit for a target DSP, while meeting the goals for every advertiser client.

In the ad exchange literature, second-price auction has once been the most widely used online auction format, due to its advantage in dominant strategy incentive compatibility. That is, an advertiser will bid exactly her valuation for the slot regardless of the bidding strategies of others. In recent years, however, a large number of ad exchanges switch their auctions to first-price auction (Google 2021) for a variety of reasons, including (i) in first price auction the winner’s payment is not affected by the uncertainty about the opponents’ bids, which not only guarantees the implementation’s credibility (Akbarpour and Li 2018), but also possibly simplifies the design of certain bidding strategy, especially when the advertiser has a budget or ROI (return over investment) constraint, and (ii) an advertiser’s exact valuation for the slot may be confidential, which makes the truth-telling equilibrium of second-price auction less attractive (Niu et al. 2017).

However, the shift from second-price to first-price auction makes it difficult for a DSP to adapt to the volatile auction environment, as the DSP needs to estimate the opponents’ bid prices to optimize its bidding strategy. The well-known *bid shading* technique, which has been widely used in US treasury auctions (Hortaçsu, Kastl, and Zhang 2018) and FCC spectrum auctions (Chakravorti et al. 1995), is then applied in first-price auction to adjust the bid price. The challenge of bid shading lies in the trade-off between payment and winning probability. Concretely, a lower bid is more likely to lose but charged less when obtaining the impression, while a higher bid is more likely to win but charged more when winning the impression. Most studies on bid shading (Pan et al. 2020; Zhou et al. 2021) rely heavily on the exact estimation of winning probability, but Zhang

et al. (2021) provides evidence on the inaccurate estimation of these parametric distributions, showing necessity of further research on bid shading.

Different from ad exchanges, a DSP manages the purchases of impressions on behalf of its advertiser clients. Therefore, when using an auction to allocate impressions and extract payments, a DSP needs to convince each advertiser that all its purchases are made optimally, making it crucial to run an incentive compatible (IC) auction so that the bids are reliable and predictable.

In this work, we consider both the optimization of bid shading technique and the design of ex-post IC auction for a target DSP participating in open (non-censored) first-price auctions, which are adopted by many ad exchanges nowadays (Gligorijevic et al. 2020). In open first-price auctions, feedback denoting the minimum winning price (i.e., the highest bid of the opponents) is sent to all participants no matter whether they win or not. Our contributions can be summarized as follows.

- We first formalize the DSP’s design problem into the joint optimization of (i) a bid shading strategy for the first-price auction and (ii) an internal auction that is ex-post IC and used once a slot is won from the first-price auction. Furthermore, we derive a reduction from joint optimization of bid shading and ex-post IC auction to learning of monotone bid functions without loss of optimality.
- We propose a partially monotone and invertible parameterized neural network to implement the monotone bid functions. With well-designed surrogate loss function, we then optimize our objective in an end-to-end manner.
- Based on both offline data and online A/B test, extensive experiments are conducted to demonstrate the effectiveness and superiority of our algorithm.

Additional Related Work

One closely related topic is automated mechanism design, where learning-based techniques are adopted for design of desired mechanism. Conitzer and Sandholm (2002) proposes the automated mechanism design approach, where the mechanism is automatically created for specific preference instance at hand. Since then, there is a growing body of research on learning-theoretic characterizations in auction settings, such as sample complexity of revenue maximization (Morgenstern and Roughgarden 2016; Cole and Roughgarden 2014; Balcan, Sandholm, and Vitercik 2018). Besides, there is also a stream of research that uses parametric function approximators to represent mechanisms and then design learning approaches for the optimization problem. Firstly, some works develop learning methods for optimization of specific parameters in auctions, such as reserve price (Mohri and Medina 2014) and boosted value (Golrezaei et al. 2021). And reinforcement learning is adopted for dynamic settings (Shen et al. 2020; Tang 2017). Secondly, a plethora of works, which belong to *differentiable economics*, aim to design specific network architecture for the more general case. For instance, some works, e.g., RegretNet (Dütting et al. 2019), EquivariantNet (Rahme et al. 2021) and CITransNet (Duan et al. 2022), relax the incentive

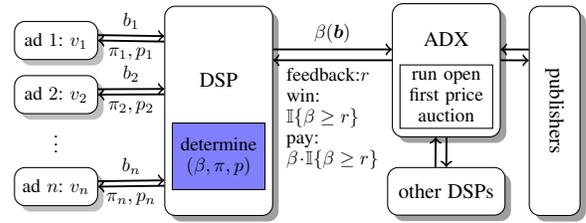


Figure 1: Illustration of the procedure. Our work focuses on the design of (β, π, p) in the blue rectangle.

constraint and design revenue-maximizing mechanisms via deep learning. In contrast, MenuNet (Shen, Tang, and Zuo 2019) and MyersonNet (Duetting et al. 2021) also leverage deep learning but design mechanisms which are exactly IC and revenue-maximizing. Among the frameworks of differentiable economics, our work is most related to the work by Liu et al. (2021), which extends MyersonNet via integrating context features and uses partially monotone networks to design IC mechanisms for multi-objective optimization in e-commerce advertising.

Another closely related line of works includes optimization of bid shading technique. In first-price auctions, the general approaches for bid shading can be divided into three categories. Firstly, point estimation is used to directly predict the minimum winning price or its derivatives. For example, Gligorijevic et al. (2020) uses models based on Factorization Machines to estimate the bid shading factor (i.e., the ratio of the minimum winning price to the original bid price). Secondly, distribution estimators are built using bid landscape forecasting to predict the winning probability distribution (Cui et al. 2011; Ren et al. 2019; Wu, Yeh, and Chen 2018; Ghosh et al. 2019), and then the optimal bid is determined by searching for maximum expected surplus (Pan et al. 2020; Zhou et al. 2021). These estimators all admit an efficient search for optimal bid but differ in the assumption of the distribution families, such as gamma and log-normal distribution. Thirdly, instead of point and distribution estimations, Zhang et al. (2021) uses online learning to design nonparametric bid shading algorithms. Outside of bid shading in first-price auction, our work is also related to the work by Nedelec, El Karoui, and Perchet (2019), which designs learning-based shading strategies to do manipulation in incentive compatible though prior-dependent revenue-maximizing auctions.

Preliminaries

In this section, we first describe the basic flow of events between the target DSP and the ad exchanges that adopt open first-price auctions. Then we characterize the ex-post IC and IR auctions for the target DSP under this scenario. Finally, we formulate the joint optimization problem of bid shading and auction for trade-offs between surplus and profit.

Procedure Description

We consider that a DSP interacts with ad exchanges in open first-price auctions. When an impression is submitted to the

ad exchange, the DSP competes for this impression on behalf of n advertiser clients in a real-time first-price auction. If the DSP wins the impression, it will allocate the impression to the clients and extract payment. We use v_i to denote the value of impression for advertiser i and use $\mathbf{v} = (v_1, \dots, v_n)$ to represent the value profile. Besides, we assume that v_i is drawn from a publicly known distribution F_i , with the corresponding density function f_i . The flow of events in our model is divided into three phases and is illustrated in Fig. 1.

Phase 1 Each advertiser client i submits b_i to the DSP. Then the DSP uses some bidding strategy to calculate $\beta(\mathbf{b})$ based on the bid profile $\mathbf{b} = (b_1, \dots, b_n)$ and submits $\beta(\mathbf{b})$ to the ad exchange.

Phase 2 The ad exchange receives bids from different DSPs and runs open first-price auction to determine who gets to display an advertisement. Specifically, let r denote the maximum bid of other DSPs, then the target DSP wins only if $\beta(\mathbf{b}) \geq r$ and will be charged $\beta(\mathbf{b})$. Otherwise, the target DSP loses and pays nothing. Moreover, unlike canonical first-price auction, the target DSP will always obtain r as feedback regardless of the auction outcome, which means the minimum winning price.

Phase 3 The target DSP determines how to allocate the impression and how much each advertiser has to pay according to some allocation rule $\pi(\mathbf{b}, r; \beta)$ and payment rule $p(\mathbf{b}, r; \beta)$. To be specific, the allocation rule π outputs an n -dimensional vector indicating the quantity of impressions allocated to each advertiser while the payment rule p outputs an n -dimensional non-negative vector specifying the payment for each advertiser.

It is worth noting that the auction (π, p) depends not only on the bid profile \mathbf{b} , but also on the bidding strategy β and feedback r , which is crucial for our design of optimal incentive compatible auction. Furthermore, we assume that r follows some known cumulative distribution $W(\cdot|\mathbf{x})$, which is determined conditionally on \mathbf{x} . Here \mathbf{x} denotes the aggregated context information, including both the features of impression and the features of all advertisers. Then the expectation of $\mathbb{I}\{\beta(\mathbf{b}) \geq r\}$ can be written as $W(\beta(\mathbf{b})|\mathbf{x})$. For simplicity, we use $W(\cdot)$ and $W(\cdot|\mathbf{x})$ interchangeably. We emphasize that \mathbf{x} does not include the bids information, thus we have a natural assumption that v_i and r are *conditionally independent* given \mathbf{x} .

Ex-post IC and IR Auction Design

Incentive compatibility (IC), which roughly means advertisers have no incentives to misreport their values, is of vital importance for the design of auction mechanism. For one thing, without loss of generality, the target DSP only needs to consider incentive compatible mechanisms according to the revelation principle (Myerson 1981). For another thing, since advertisers do not need to compute complex bidding strategy in IC mechanisms, the IC property ensures the stability of the advertising system in that the bids from advertiser clients are reliable and predicible. Therefore, in what follows we do not distinguish between advertisers' values \mathbf{v} and bids \mathbf{b} when there is no confusion. As discussed in the

work by Babaioff et al. (2020) and Ni and Tang (2022), it is without loss of generality to consider the classic *quasi-linear utility maximization model*, where each advertiser aims to maximize its quasi-linear utility $u_i = v_i \pi_i - p_i$ (More explanations are given in the full version). Then the ex-post incentive compatibility (IC) and individual rationality (IR) can be defined formally as follows.

Definition 1. Given β , an auction mechanism (π, p) is *ex-post IC* if for all i, v_i, b_i, b_{-i} and r , the following inequality holds.

$$\begin{aligned} & v_i \pi_i(v_i, b_{-i}, r; \beta) - p_i(v_i, b_{-i}, r; \beta) \\ & \geq v_i \pi_i(b_i, b_{-i}, r; \beta) - p_i(b_i, b_{-i}, r; \beta). \end{aligned} \quad (1)$$

Besides, (π, p) is *ex-post IR* if for all i, v_i, b_{-i} and r , the utility is non-negative, that is, $v_i \pi_i(v_i, b_{-i}, r; \beta) - p_i(v_i, b_{-i}, r; \beta) \geq 0$.

Note that the definition of ex-post IC and IR depends on the realization r of the distribution $W(\cdot)$, which is the feedback of open first-price auction. Following the work by Myerson (1981), we can fully characterize the set of ex-post IC and IR auctions as stated in Proposition 2 and 3, respectively. Moreover, Proposition 2 points out that with the help of payment identity (2), we can only concentrate on the design of allocation function.

Proposition 2. Given β , an auction mechanism (π, p) is *ex-post IC* if and only if for all i, v_{-i} and r , the following two conditions hold.

- (*Monotonicity*) The allocation $\pi_i(v_i, v_{-i}, r; \beta)$ for advertiser i is monotone increasing in v_i .
- (*Payment identity*) The payment for advertiser i satisfies

$$\begin{aligned} & p_i(v_i, v_{-i}, r; \beta) - p_i(0, v_{-i}, r; \beta) \\ & = v_i \cdot \pi_i(v_i, v_{-i}, r; \beta) - \int_0^{v_i} \pi_i(s, v_{-i}, r; \beta) ds. \end{aligned} \quad (2)$$

Proposition 3. Given β , an *ex-post IC* auction (π, p) is *ex-post IR* iff. for all i, v_{-i} and r , $p_i(0, v_{-i}, r; \beta) \leq 0$ holds.

Problem Formulation

In this work, we consider the joint optimization problem of bid shading strategy β and ex-post IC and IR auction design (π, p) from the perspective of the target DSP. Besides, the objective we optimize is the linear combination of surplus and profit. To be specific, the surplus of the DSP is given by the difference between the obtained value for the advertiser clients and the payment to ad exchanges, i.e., $S(\mathbf{v}, r; \pi, p, \beta) = \sum_{i=1}^n v_i \pi_i(\mathbf{v}, r; \beta) - \beta(\mathbf{v}) \cdot \mathbb{I}\{\beta(\mathbf{v}) \geq r\}$, while the profit is defined as the difference between the revenue from the managed advertisers and the payment to ad exchanges, i.e., $R(\mathbf{v}, r; \pi, p, \beta) = \sum_{i=1}^n p_i - \beta(\mathbf{v}) \cdot \mathbb{I}\{\beta(\mathbf{v}) \geq r\}$. Therefore, we can formulate the corresponding joint optimization problem.

$$\begin{aligned} & \max_{\beta, \pi, p} \mathbb{E}_{(\mathbf{v}, r, \mathbf{x}) \sim D} \left[\sum_{i=1}^n [\lambda v_i \pi_i(\mathbf{v}, r; \beta) + (1 - \lambda) p_i(\mathbf{v}, r; \beta)] \right. \\ & \quad \left. - \beta(\mathbf{v}) \cdot \mathbb{I}\{\beta(\mathbf{v}) \geq r\} \right]. \quad (P1) \\ & \text{s.t. } \begin{cases} \sum_{i=1}^n \pi_i(\mathbf{v}, r; \beta) \leq \mathbb{I}\{\beta(\mathbf{v}) \geq r\}, \\ (\pi, p) \text{ is ex-post IC and ex-post IR.} \end{cases} \end{aligned}$$

Here D denotes the dataset where v_i is sampled from F_i and r is drawn from $W(\cdot|\mathbf{x})$, and λ is used to balance the surplus and profit.

Comparison with related work. Firstly, most related works focus on either the optimization of bid shading strategy or the design of IC auction for various settings, while the bid shading and the internal auction are optimized simultaneously in this paper. For example, the optimization problem (P1) is similar with that in the work by Bachrach et al. (2014), the difference is that we also need to optimize the bid shading strategy β , which jointly influences the outcome of mechanism (π, p) with the feedback r . Secondly, although some works (Grigas et al. 2017) also consider both the bidding strategy and the internal auction among the advertiser clients, the ad exchanges in these works are assumed to adopt second-price auction, whose truthfulness simplifies the optimization of bidding strategy.

Analysis

In this section, we derive a reduction from joint optimization problem (P1) of bid shading and ex-post IC auction to optimization of monotone bid functions and show its merits.

Reduction

In the optimization problem (P1), there are three functions to be determined, that is, the bid shading strategy β , the allocation rule π and the payment rule p . And the bid shading strategy will jointly influence the allocation and payment rule with the feedback r , making it difficult to solve the corresponding problem. However, motivated by the celebrated work (Myerson 1981) which designs the revenue-maximization auction via ranking the buyers by *virtual value* $\varphi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$, we use monotone bid functions to derive the bid shading strategy and construct a family of ex-post IC auctions, and then convert the original problem (P1) to direct optimization of bid functions.

To be specific, for each advertiser i , given a monotone bid function $\rho_i(b_i)$ and its inverse transform $\rho_i^{-1}(V) = \inf \{y | \rho_i(y) \geq V\}$, we can sort the advertisers in a non-increasing order of bid scores, i.e., $\rho_1(b_1) \geq \rho_2(b_2) \geq \dots \geq \rho_n(b_n)$. Then we can define a bidding strategy as $\beta^\rho(\mathbf{b}) = \rho_1(b_1)$. After submitting bid $\beta^\rho(\mathbf{b})$ to the ad exchange and getting the minimum bid to win as feedback r , we can further construct an auction as follows.

Mechanism 4 (ρ -Induced Auction). Given the bid scores $\rho_1(b_1) \geq \rho_2(b_2) \geq \dots \geq \rho_n(b_n)$ and the feedback r :

- Allocate the impression to advertiser 1 when $\rho_1(b_1) \geq r$, i.e., $\pi_1^\rho = \mathbb{I}\{\rho_1(b_1) \geq r\}$.
- Charge advertiser 1 the minimum bid to win the impression if allocated, i.e., $p_1^\rho = \pi_1^{-1}(\max\{\rho_2(b_2), r\})$.

It is straightforward to verify that the ρ -induced auction (Mechanism 4) is ex-post IC and ex-post IR. Thus the auctions induced by the monotone bid functions $\rho = \{\rho_i\}_{i=1}^n$ form a class of ex-post IC and IR auctions. For convenience, we use $\pi(\cdot; \beta^\rho)$ and $\pi^\rho(\cdot)$, $p(\cdot; \beta^\rho)$ and $p^\rho(\cdot)$ interchangeably. Besides, we can incorporate the features into ρ_i (i.e., $\rho_i(v_i; \mathbf{x})$) as long as it is partially monotone in the value. Therefore, we can formulate the optimization problem with

respect to ρ as follows.

$$\max_{\{\rho_i\}} \mathbb{E}_{(\mathbf{v}, r, \mathbf{x}) \sim D} \left[\sum_{i=1}^n \lambda v_i \pi_i^\rho(\mathbf{v}, r) + (1 - \lambda) p_i^\rho(\mathbf{v}, r) - \max_j \rho_j(v_j; \mathbf{x}) \sum_{i=1}^n \pi_i^\rho(\mathbf{v}, r) \right]. \quad (\text{P2})$$

s.t. $\rho_i(v_i; \mathbf{x})$ is monotone increasing in $v_i, \forall i$.

We then propose our main theorem which derives a reduction from problem (P1) to the optimization problems of ρ under standard assumption about $\psi_i(v_i) = \lambda v_i + (1 - \lambda)\varphi_i(v_i)$. The intuition of reduction is that the optimal bidding strategy is monotone increasing in the winning value.

Theorem 5. If v_i is independent of each other given \mathbf{x} and $\psi_i(v_i)$ is monotone increasing in v_i for all advertisers, then the induced bidding strategy and ex-post IC and IR auction $(\beta^{\rho^*}, \pi^{\rho^*}, p^{\rho^*})$ of the optimal solution $\rho^* = \{\rho_i^*\}_{i=1}^n$ for problem (P2) is also the optimal solution for problem (P1).

Special Case and Relation with Myerson Auction

To help interpret Theorem 5, we provide analysis of a special case, where by the following assumption on distribution W , the optimal solution for (P1) admits a close-form expression. The correspondence between ρ and optimal solution (β, π, p) can thus be clearly identified.

Assumption 6. The feedback r and each advertiser's value v_i are defined on a continuous interval $[\underline{r}, \bar{r}]$. Moreover, the probability density function $W'(\cdot)$ is positive on the interval and $W(\xi)/W'(\xi) + \xi$ is strictly monotone increasing in ξ .

It is worth noting that many common used distributions in the field of auction satisfy Assumption 6, such as truncated-normal, gamma and log-normal distribution. Let $\bar{s}(V)$ denote the value that maximizes $(V - \xi)W(\xi)$, i.e., $\bar{s}(V) = \arg \max_{\xi \in [\underline{r}, \bar{r}]} (V - \xi)W(\xi)$, then we can verify that $\bar{s}(V)$ is well-defined and is monotone increasing in V according to Assumption 6. Therefore, the optimal bid shading and ex-post IC auction for problem (P1) can be derived naturally via payment identity (2) under this assumption.

Proposition 7. Suppose W satisfies Assumption 6, v_i is independent of each other and $\psi_i(v_i)$ is monotone increasing in v_i for all advertisers. Let $i^* \in \arg \max_i \psi_i(v_i)$, then the optimal bidding strategy $\beta^*(\mathbf{v})$ for (P1) only depends on $\psi_{i^*}(v_{i^*})$, i.e., $\beta^*(\mathbf{v}) = \bar{s}(\psi_{i^*}(v_{i^*}))$. And the optimal auction (π^*, p^*) for (P1) can be written as

$$\begin{aligned} \pi_i^*(\mathbf{v}, r; \beta^*) &= \mathbb{I}\{\beta^*(\psi_i(v_i)) \geq r\} \cdot \mathbb{I}\{i = i^*\}, \\ p_i^*(\mathbf{v}, r; \beta^*) &= \pi_i^*(\mathbf{v}, r; \beta^*) \cdot \hat{p}_i(\mathbf{v}, r), \end{aligned}$$

where $\hat{p}_i(\mathbf{v}, r) = \psi_i^{-1}(\max\{\max_{j \neq i} \psi_j(v_j), \frac{W(r)}{W'(r)} + r\})$.

In Proposition 7, $\hat{p}_i(\mathbf{v}, r)$ denotes the minimum bid for advertiser i^* to win the impression under the optimal bidding strategy β^* after the realization of r . Compared with Myerson auction (Myerson 1981), the challenge in optimization problem (P1) is that the unknown r causes randomness of whether there exists an impression to allocate, and an ex-post IC mechanism has to take such randomness into consideration. To be specific, under Assumption 6, the reserve price is now $\psi_i^{-1}(W(r)/W'(r) + r)$ instead of $\varphi_i^{-1}(0)$ in Myerson auction.

Proposition 8. *Suppose W satisfies Assumption 6, v_i is independent of each other and $\psi_i(v_i)$ is monotone increasing in v_i for all advertisers. Let $\rho_i^*(v_i) = \bar{s}(\psi_i(v_i))$ denote the monotone bid function for advertiser i , then the induced bidding strategy and ex-post IC and IR auction (β^*, π^*, p^*) is the optimal solution for problem (P1).*

Moreover, as stated in Proposition 8, we can also obtain the optimal formula of bid shading and auction for problem (P1) via constructing the monotone bid functions. And it is worth noting that the optimal monotone bid function $\rho_i^*(v_i)$ in Proposition 8 not only depends on $\psi_i(v_i)$ (ranking in Myerson auction), but also relies on the distribution W .

Comparison with Related Work in Bid Shading

As far as we know, a line of works (Zhou et al. 2021; Pan et al. 2020) which optimizes bid shading in open first-price auctions relies on relatively strong hypothesis about the feedback distribution W . To be specific, these studies can be summarized into a two-stage solution framework. In Stage 1, $W(\xi|\mathbf{x})$ is estimated via bid landscape forecasting technique (Cui et al. 2011; Ren et al. 2019; Wu, Yeh, and Chen 2018; Ghosh et al. 2019) and $\hat{W}(\xi|\mathbf{x})$ is obtained to approximate $W(\xi|\mathbf{x})$. In Stage 2, the optimal solution is derived with respect to $\hat{W}(\xi|\mathbf{x})$ based on Proposition 7.

We argue that these work have limited performance in two aspects. Firstly, in order to admit an efficient computation of $\beta^*(\mathbf{v})$ in Stage 2, the hypothesis class for $\hat{W}(\xi|\mathbf{x})$ is usually formulated as a family of special distributions (e.g., gamma and log-normal distribution), which is characterized by Assumption 6 and can not fully capture the pattern of real distribution W . Secondly, the objective of Stage 1 is often to maximize the expected log-likelihood $\mathbb{E}_{r,\mathbf{x}}[\log \hat{W}(r|\mathbf{x})]$, while the objective of Stage 2 is to optimize the expected trade-offs between surplus and profit. And there are no direct relationship between these two objectives in two stages, which is illustrated in the full version.

However, note that Theorem 5 does not rely on any assumptions on the distribution W (e.g., Assumption 6), thus we can further reduce the negative effect caused by distribution estimation through optimizing the monotone bid functions directly, which will be shown in the next section.

Learning Framework

In this section, we first propose a neural network which is partially monotone and invertible to implement the monotone bid functions. Then we meticulously design a surrogate loss function to train the objective in an end-to-end manner.

Parameterized Bidding

Following the construction of ρ -induced auction, all we need to design are the monotone bid functions $\rho = \{\rho_i\}_{i=1}^n$, which need to meet the following requirements: (i) Each ρ_i is monotone increasing in v_i and the inverse transform should be computed efficiently. (ii) The monotone bid function should implicitly capture the pattern of the feedback distribution W and the value distribution F_i .

The similarity between Mechanism 4 and Myerson auction motivates us to adopt MyersonNet (Duetting et al. 2021)

to implement monotone bid functions. However, due to the diversity of the advertiser clients and the impressions from ad exchanges, we also incorporate the features from both advertisers and impressions within the design of monotone bid functions. Besides, it is impossible to design dedicated function for each advertiser. Instead, every advertisers share the same parametric function and we use $\bar{\mathbf{x}}_i = (\mathbf{x}, \text{ID}_i)$ to denote the input features to calculate ρ_i , where ID_i is the ID embedding for advertiser i . In light of the above characterizations, we extend the design of partially monotone network in (Daniels and Velikova 2010; Liu et al. 2021) and rewrite the bid function ρ_i with respect to v_i and \mathbf{x} as:

$$\hat{\rho}(v_i; \bar{\mathbf{x}}_i) = \min_{q \in [Q]} \max_{z \in [Z]} \{e^{w_{qz} + \theta \cdot \bar{\mathbf{x}}_i} v_i + \mathbf{W}_{qz} \cdot \bar{\mathbf{x}}_i + \alpha_{qz}\}, \quad (3)$$

where $\hat{\rho}$ is mainly implemented via a two-layer feed-forward network (illustrated in the full version), taking the value v_i and feature $\bar{\mathbf{x}}_i$ as input and outputting ρ_i . The monotonicity is then guaranteed via min and max operators over these partially monotone linear functions in v_i . Moreover, the inverse transform of $\hat{\rho}(v_i; \bar{\mathbf{x}}_i)$ can be computed efficiently in a closed-form expression, that is,

$$\hat{\rho}^{-1}(V; \bar{\mathbf{x}}_i) = \max_{q \in [Q]} \min_{z \in [Z]} \{e^{-w_{qz} - \theta \cdot \bar{\mathbf{x}}_i} (V - \mathbf{W}_{qz} \cdot \bar{\mathbf{x}}_i - \alpha_{qz})\}.$$

It is worth noting that the difference from the network architecture in (Liu et al. 2021) is that we adopt an additional unit $\theta \cdot \bar{\mathbf{x}}_i$ in the slope in Equation (3). In fact, as argued by Daniels and Velikova (2010), without $\theta \cdot \bar{\mathbf{x}}_i$, this partially monotone MIN-MAX network has the capacity to approximate any partially monotone functions arbitrarily well, given sufficiently large Q and Z . However, both Q and Z are pre-defined, and the neural network can not capture all the diversity of advertiser clients and impressions with limited Q and Z . Therefore, an additional term $\theta \cdot \bar{\mathbf{x}}_i$, which depends on the features, can further enhance the capability of the neural network to approximate the partially monotone functions.

Surrogate Loss Function

With the design of monotone bid functions $\rho = \{\rho_i\}_{i=1}^n$, we can specify the optimization problem (P2). Without loss of generality, we assume that $\rho_1 \geq \rho_2 \geq \dots \geq \rho_n$. Then the bidding strategy is $\beta^p(\mathbf{v}) = \rho_1$, and the allocation and payment rule for advertiser 1 can be written as $\pi_1^p(\mathbf{v}, r) = \mathbb{I}\{\rho_1 \geq r\}$ and $p_1^p(\mathbf{v}, r) = \pi_1^p(\mathbf{v}, r) \cdot \hat{\rho}^{-1}(\max\{\rho_2, r\}; \bar{\mathbf{x}}_1)$. For convenience, we use $l^p(\mathbf{v}, r, \mathbf{x})$ to denote the linear combination of surplus and profit, that is,

$$l^p(\mathbf{v}, r, \mathbf{x}) = \lambda v_1 + (1 - \lambda) \hat{\rho}^{-1}(\max\{\rho_2, r\}; \bar{\mathbf{x}}_1) - \rho_1.$$

Then for problem (P2), we aim to minimize the following loss function:

$$\mathcal{L}(w, \mathbf{W}, \theta, \alpha) = -\mathbb{E}_{(\mathbf{v}, r, \mathbf{x}) \sim D} [l^p(\mathbf{v}, r, \mathbf{x}) \cdot \mathbb{I}\{\rho_1 \geq r\}], \quad (4)$$

which is the opposite number of the expected objective. However, note that the indicator function $\mathbb{I}\{\rho_1 \geq r\}$ in Equation (4) is not continuous at $\rho_1 = r$, we replace the indicator function by a sigmoid function to obtain a fully differentiable *surrogate loss*, that is,

$$\hat{\mathcal{L}}(w, \mathbf{W}, \theta, \alpha) = -\mathbb{E}_{(\mathbf{v}, r, \mathbf{x}) \sim D} [l^p(\mathbf{v}, r, \mathbf{x}) \cdot \sigma(\eta \frac{\rho_1 - r}{r})], \quad (5)$$

where $\sigma(\xi) = \frac{1}{1 + \exp(-\xi)}$ and $\eta > 0$ is a temperature parameter which denotes the degree of the approximation with the

indicator function. We gradually increase η during the training process to approach $\mathbb{I}\{\rho_1 \geq r\}$ and when $\eta \rightarrow +\infty$, $\sigma(\eta \frac{\rho_1 - r}{r}) \rightarrow \mathbb{I}\{\rho_1 \geq r\}$, $\forall \rho_1 \neq r$. Besides, $\frac{\rho_1 - r}{r}$ in the sigmoid function is used to reduce the effects caused by the wide range of r , and thus normalize the input to the sigmoid function. Therefore, we can optimize the objective in an end-to-end manner via minimizing the surrogate loss $\hat{\mathcal{L}}(\mathbf{w}, \mathbf{W}, \boldsymbol{\theta}, \boldsymbol{\alpha})$, and name the method as **MINMAX**.

Experiments

In this section, we first exhibit the representation capability of our partially monotone network to approach the optimum on various distributions. Then, based on both offline data and online A/B test, we show the effectiveness and superiority of our learning method against state-of-the-art algorithms.

Experiment Setup

Evaluation metrics. Metrics in our offline and online experiments depend on whether the optimal solution can be fully characterized. In offline experiments, when the values and the feedback are sampled from predefined distributions, we can derive the optimal solution (π^*, p^*, β^*) . However, on offline industrial data and online A/B test, we can not fully characterize the optimal solution. Instead, we can replace the performance of optimal solution by the upper bound $\sum_{t=1}^T (\max_i v_i^t - r)^+$ where $\xi^+ = \max\{\xi, 0\}$. Suppose the output of the methods to be verified is (π, p, β) , then we consider the following metrics: bid deviation (BD), achieved objective (AO) and achieved post-objective (APO).

$$BD = \frac{\sum_{t=1}^T \frac{|\beta(\mathbf{v}^t) - \beta^*(\mathbf{v}^t)|}{\beta(\mathbf{v}^t) + \beta^*(\mathbf{v}^t)} \cdot \mathbb{I}\{\max\{\beta^*(\mathbf{v}^t), \beta(\mathbf{v}^t)\} \geq r^t\}}{\sum_{t=1}^T \mathbb{I}\{\max\{\beta^*(\mathbf{v}^t), \beta(\mathbf{v}^t)\} \geq r^t\}},$$

$$AO = \frac{\sum_{t=1}^T \lambda S(\mathbf{v}^t, r^t; \pi, p, \beta) + (1 - \lambda)R(\mathbf{v}^t, r^t; \pi, p, \beta)}{\sum_{t=1}^T \lambda S(\mathbf{v}^t, r^t; \pi^*, p^*, \beta^*) + (1 - \lambda)R(\mathbf{v}^t, r^t; \pi^*, p^*, \beta^*)},$$

$$APO = \frac{\sum_{t=1}^T \lambda S(\mathbf{v}^t, r^t; \pi, p, \beta) + (1 - \lambda)R(\mathbf{v}^t, r^t; \pi, p, \beta)}{\sum_{t=1}^T (\max_i v_i^t - r)^+}.$$

Here we use $(\mathbf{v}^t, r^t, \mathbf{x}^t)$ to denote the t -th instance in the test dataset. Besides, we mainly focus on two classical objectives in our experiments, *i.e.*, *surplus maximization* ($\lambda = 1$) and *profit maximization* ($\lambda = 0$).

Baseline methods. We first modify the network architecture for comparison. Same as that in (Liu et al. 2021), **Mod1** uses $\hat{\rho}(v_i; \mathbf{x}_i) = \min_{q \in [Q]} \max_{z \in [Z]} \{e^{w_{qz}} v_i + \mathbf{W}_{qz} \cdot \mathbf{x}_i + \alpha_{qz}\}$ to implement the monotone bid function and aim to minimize the surrogate loss. While for **Mod2**, it uses a different monotone bid function:

$$\hat{\rho}(v_i; \mathbf{x}_i) = \min_{q \in [Q]} \max_{z \in [Z]} \{e^{w_{qz} + \boldsymbol{\theta}_{qz} \cdot \mathbf{x}_i} v_i + \mathbf{W}_{qz} \cdot \mathbf{x}_i + \alpha_{qz}\},$$

where each group has its own weight $\boldsymbol{\theta}_{qz}$. Apart from the aforementioned methods which are variants of MINMAX, we consider the following methods for comparison, which are widely used for bid shading in open first-price auctions. Note that these methods only focus on the bid shading technique for surplus maximization ($\lambda = 1$), and ignore the auction design for the advertiser clients, we complete the design of an ex-post IC and IR auction by deriving the corresponding payment rule based on payment identity (2). More details are available in the full version.

Objective	Metrics	uniform	exponential
surplus ($\lambda = 1$)	AO	0.9997(0.0029)	1.0019(0.0057)
	BD	0.0069(0.0023)	0.0234(0.0014)
profit ($\lambda = 0$)	AO	0.9841(0.0375)	0.9771(0.0125)
	BD	0.0203(0.0042)	0.0460(0.0101)

Table 1: Performance on uniform and exponential distributions. Numbers in the brackets denote standard deviations.

- Estimation of bid shading ratio (**BSR**) (Gligorijevic et al. 2020). This method transforms the optimization problem of bid shading to the regression problem, where the label is given by the bid shading ratio $\frac{r}{\max v_i}$.
- Estimation of distribution with sigmoid function (**Dist-S**) (Pan et al. 2020). This method uses $\hat{W}(\xi|\mathbf{x}) = (1 + e^{-\alpha(\mathbf{x}) - k(\mathbf{x}) \log \xi})^{-1}$ to approximate $W(\xi|\mathbf{x})$ via maximum likelihood estimation.
- **Dist-G** and **Dist-L** (Zhou et al. 2021). These two methods resemble Dist-S. The only difference is that Dist-G and Dist-L use gamma distribution and log-normal distributions to approximate $W(\xi|\mathbf{x})$, respectively.

Offline Experiments

Data description. We use both synthetic data and industrial data to conduct the offline experiments. The synthetic data is generated from five different distributions while the industrial data is sampled from the open iPinYou RTB dataset (Zhang et al. 2014). More details for the generated offline data can be found in the full version.

Representation capability. As for synthetic data, we can compute the optimal solution, then we can exhibit the representation capability of our partially monotone network to approach the optimum. Firstly, we present the overall evaluated metrics in Table 1. We can see that for surplus maximization and profit maximization, both achieved objective and bid deviation of MINMAX are very close to the optimum on three types of distributions. Secondly, note that the optimal bidding strategy in Proposition 7 and the bidding strategy induced by ρ depend on only one of the values, we can compare these two bidding strategies via plotting the curve of $(v_i^*, \beta(\mathbf{v}))$. As Fig.2 shows, the solid lines denote the optimal bidding strategies while the scatter points form the bidding strategies of MINMAX. Besides, we use different colors to denote different bidding strategies for distributions $W(\xi|\mathbf{x})$ with various parameters. We can see that scatter points and solid lines coincide for every distribution, especially in the area with high probability, which means our partially monotone network uses piece-wise linear functions to approach the optimal bidding strategy and captures the pattern of both the feedback distribution W and the value distribution F_i implicitly.

Comparison with baseline methods. We compare our methods and baselines on two types of datasets. We first conduct experiments on synthetic data generated by log-normal and gamma distribution. The performance of all methods on synthetic data are listed in Table 2. Compared with Dist-G,

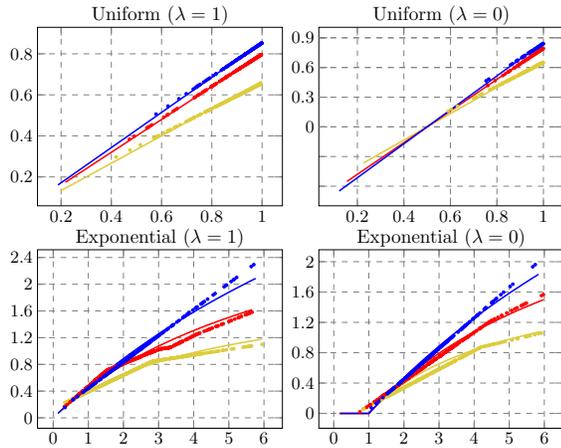


Figure 2: Comparison with optimal bidding strategy, where x-axis denotes the value while y-axis denotes the bid price.

Methods	log-normal		gamma	
	surplus	profit	surplus	profit
Dist-S	0.9894	0.7573	0.9091	0.8138
Dist-G	0.8242	0.5863	0.9693	0.8814
Dist-L	0.9954	0.8624	0.9601	0.8769
Mod1	0.9916	0.9242	0.9804	0.9274
Mod2	0.9929	0.9412	0.9807	0.9324
Simp	0.9882	0.9338	0.9793	0.9297
MINMAX	0.9928	0.9441	0.9815	0.9339

Table 2: Performance comparison (AO) on log-normal and gamma distributions, where the parameters of distributions are linear on the feature space \mathbb{R}^{20} .

Dist-L performs better on log-normal distribution but worse on gamma distribution, which justifies the discussion about the methods based on distribution estimation. We can see that MINMAX approaches the optimal objective in all cases, verifying the effectiveness of our methods for tackling the inconsistency problem caused by distribution estimation. We further examine the impact of hyper-parameters and architectures of our partially monotone network. The results in Table 2 show the superiority of our methods except for surplus maximization on log-normal distribution. Concretely, MINMAX has better representation capability than Mod1 and the optimization in MINMAX is more stable than that in Mod2. Besides, we set $Q = Z = 1$ in MINMAX and obtain Simp. We note that Simp even outperforms Mod1 for profit maximization, showing the effectiveness of the term $\theta \cdot \bar{x}_i$ in the slope.

We then compare the performance on industrial data. Here we only consider the baseline methods that are widely used for bid shading, i.e., BSR, Dist-S, Dist-G and Dist-L. In order to have a fair comparison among different methods, we use the same architecture based on DeepFM (Guo et al. 2017) to generate high-level representations of features and then feed them into the neural networks in different methods. As Table 3 shows, among the baseline methods, Dist-L has the best performance for both surplus and profit max-

Methods	surplus($\lambda = 1$)	profit($\lambda = 0$)
BSR	0.4138(0.0128)	0.2297(0.0092)
Dist-S	0.4921(0.0104)	0.2352(0.0118)
Dist-G	0.4978(0.0048)	0.2735(0.0018)
Dist-L	0.5159(0.0010)	0.2844(0.0038)
MINMAX	0.5542(0.0014)	0.3491(0.0020)

Table 3: Performance comparison (APO) on industrial data.

imization, and methods based on distribution estimation all perform better than point estimation method BSR because of robustness. Besides, MINMAX outperforms all the baseline methods for both surplus and profit maximization by a clear margin. Concretely, compared with Dist-L, MINMAX has 7.42% surplus lift and 22.75% profit lift. We argue that the benefits of our methods are twofold. Firstly, they are capable of approximating any partially monotone functions. Secondly, they incorporate the objective of optimization problem into the design of surrogate loss function, resulting in straightforward end-to-end training.

Online Experiments

In order to show that our proposed method can be incorporated into the complex advertising systems and improve the target performance, we deploy it in the real-time bid shading serving module of one of the world’s leading advertising platforms. This real-time bid shading serving module interacts with other modules of the advertising systems (e.g., auto-bidding) and serves billions of bid requests that are run via open first-price auctions per day. In practice, we compare MINMAX with the state-of-the-art method Dist-L for the commonly used surplus maximization.

To demonstrate the performance of MINMAX compared to Dist-L, we conduct online A/B tests with a percentage of traffic volume, which was randomly selected with the same scale for both methods. Hence we compare obtained surplus $\sum S(v^t, r^t; \pi, p, \beta)$ and obtained impressions $\sum \mathbb{I}\{\beta(v^t) \geq r^t\}$ instead of achieved post-objective. To keep dataset information confidential, we only show the improvement of these two metrics. Specifically, MINMAX improves the surplus by 6.41% at a much lower impression increase (1.80%), verifying the effectiveness of our methods in online setting. Apart from the improvements on key metrics, we further argue that MINMAX has an additional advantage on more efficient computation in deployment, that is, there is no need to search optimal bid for surplus maximization based on the estimation of distribution W , all we need is the inference based on trained model.

Conclusion

In this paper, we consider both the optimization of bid shading and the design of ex-post IC auction for the management of a DSP. Firstly, we derive a reduction of the joint optimization problem. Then we propose a neural network to implement the monotone bid functions, and design surrogate loss function to enable end-to-end training of the objective. Finally, extensive experiments are conducted to demonstrate the effectiveness and superiority of our algorithm.

Acknowledgments

This work was supported by the National Key Research and Development Program of China under Grant No. 2020AAA0103401.

References

- Akbarpour, M.; and Li, S. 2018. Credible Mechanisms. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC '18, 371. New York, NY, USA: Association for Computing Machinery. ISBN 9781450358293.
- Babaioff, M.; Cole, R.; Hartline, J.; Immorlica, N.; and Lucier, B. 2020. Non-quasi-linear Agents in Quasi-linear Mechanisms. arXiv:2012.02893.
- Bachrach, Y.; Ceppi, S.; Kash, I. A.; Key, P.; and Kurokawa, D. 2014. Optimising trade-offs among stakeholders in ad auctions. In *Proceedings of the fifteenth ACM conference on Economics and computation*, 75–92.
- Balcan, M.-F.; Sandholm, T.; and Vitercik, E. 2018. A general theory of sample complexity for multi-item profit maximization. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 173–174.
- Chakravorti, B.; Sharkey, W. W.; Spiegel, Y.; and Wilkie, S. 1995. Auctioning the airwaves: the contest for broadband PCS spectrum. *Journal of Economics & Management Strategy*, 4(2): 345–373.
- Clarke, E. H. 1971. Multipart pricing of public goods. *Public choice*, 17–33.
- Cole, R.; and Roughgarden, T. 2014. The sample complexity of revenue maximization. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, 243–252.
- Conitzer, V.; and Sandholm, T. 2002. Complexity of Mechanism Design. In *UAI '02, Proceedings of the 18th Conference in Uncertainty in Artificial Intelligence*, 103–110.
- Cui, Y.; Zhang, R.; Li, W.; and Mao, J. 2011. Bid landscape forecasting in online ad exchange marketplace. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, 265–273.
- Daniels, H.; and Velikova, M. 2010. Monotone and partially monotone neural networks. *IEEE Transactions on Neural Networks*, 21(6): 906–917.
- Duan, Z.; Tang, J.; Yin, Y.; Feng, Z.; Yan, X.; Zaheer, M.; and Deng, X. 2022. A Context-Integrated Transformer-Based Neural Network for Auction Design. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvari, C.; Niu, G.; and Sabato, S., eds., *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 5609–5626. PMLR.
- Duetting, P.; Feng, Z.; Narasimhan, H.; Parkes, D. C.; and Ravindranath, S. S. 2021. Optimal Auctions through Deep Learning. *Communications of the ACM*, 64 (Issue 8): 109–116.
- Dütting, P.; Feng, Z.; Narasimhan, H.; Parkes, D.; and Ravindranath, S. S. 2019. Optimal auctions through deep learning. In *International Conference on Machine Learning*, 1706–1715. PMLR.
- Edelman, B.; Ostrovsky, M.; and Schwarz, M. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review*, 97(1): 242–259.
- Ghosh, A.; Mitra, S.; Sarkhel, S.; Xie, J.; Wu, G.; and Swaminathan, V. 2019. Scalable bid landscape forecasting in real-time bidding. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 451–466. Springer.
- Gligorijevic, D.; Zhou, T.; Shetty, B.; Kitts, B.; Pan, S.; Pan, J.; and Flores, A. 2020. Bid Shading in The Brave New World of First-Price Auctions. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, 2453–2460.
- Golrezaei, N.; Lin, M.; Mirrokni, V.; and Nazerzadeh, H. 2021. Boosted Second Price Auctions: Revenue Optimization for Heterogeneous Bidders. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 447–457.
- Google. 2021. Moving AdSense to a first-price auction. <https://blog.google/products/adsense/our-move-to-a-first-price-auction/>. Accessed: 2023-08-23.
- Grigas, P.; Lobos, A.; Wen, Z.; and Lee, K.-c. 2017. Profit Maximization for Online Advertising Demand-Side Platforms. arXiv:1706.01614.
- Groves, T.; et al. 1973. Incentives in teams. *Econometrica*, 41(4): 617–631.
- Guo, H.; Tang, R.; Ye, Y.; Li, Z.; and He, X. 2017. DeepFM: a factorization-machine based neural network for CTR prediction. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 1725–1731.
- Hortaçsu, A.; Kastl, J.; and Zhang, A. 2018. Bid shading and bidder surplus in the us treasury auction system. *American Economic Review*, 108(1): 147–69.
- Liu, X.; Yu, C.; Zhang, Z.; Zheng, Z.; Rong, Y.; Lv, H.; Huo, D.; Wang, Y.; Chen, D.; Xu, J.; et al. 2021. Neural auction: End-to-end learning of auction mechanisms for e-commerce advertising. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 3354–3364.
- Mohri, M.; and Medina, A. M. 2014. Learning theory and algorithms for revenue optimization in second price auctions with reserve. In *International conference on machine learning*, 262–270. PMLR.
- Morgenstern, J.; and Roughgarden, T. 2016. Learning simple auctions. In *Conference on Learning Theory*, 1298–1318. PMLR.
- Myerson, R. B. 1981. Optimal auction design. *Mathematics of operations research*, 6(1): 58–73.
- Nedelec, T.; El Karoui, N.; and Perchet, V. 2019. Learning to bid in revenue-maximizing auctions. In *International Conference on Machine Learning*, 4781–4789. PMLR.
- Ni, B.; and Tang, P. 2022. Characterization of Incentive Compatibility of an Ex-ante Constrained Player. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 5156–5163.

- Niu, C.; Zhou, M.; Zheng, Z.; Wu, F.; and Chen, G. 2017. ERA: Towards privacy preservation and verifiability for on-line ad exchanges. *Journal of Network and Computer Applications*, 98: 1–10.
- Pan, S.; Kitts, B.; Zhou, T.; He, H.; Shetty, B.; Flores, A.; Gligorijevic, D.; Pan, J.; Mao, T.; Gultekin, S.; et al. 2020. Bid Shading by Win-Rate Estimation and Surplus Maximization. arXiv:2009.09259.
- Rahme, J.; Jelassi, S.; Bruna, J.; and Weinberg, S. 2021. A Permutation-Equivariant Neural Network Architecture For Auction Design. In *35th AAAI Conference on Artificial Intelligence, AAAI 2021, 35th AAAI Conference on Artificial Intelligence, AAAI 2021*, 5664–5672. Association for the Advancement of Artificial Intelligence. Publisher Copyright: Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.; 35th AAAI Conference on Artificial Intelligence, AAAI 2021 ; Conference date: 02-02-2021 Through 09-02-2021.
- Ren, K.; Qin, J.; Zheng, L.; Yang, Z.; Zhang, W.; and Yu, Y. 2019. Deep landscape forecasting for real-time bidding advertising. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 363–372.
- Shen, W.; Peng, B.; Liu, H.; Zhang, M.; Qian, R.; Hong, Y.; Guo, Z.; Ding, Z.; Lu, P.; and Tang, P. 2020. Reinforcement mechanism design: With applications to dynamic pricing in sponsored search auctions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 2236–2243.
- Shen, W.; Tang, P.; and Zuo, S. 2019. Automated Mechanism Design via Neural Networks. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 215–223.
- Tang, P. 2017. Reinforcement mechanism design. In *IJCAI*, 5146–5150.
- Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1): 8–37.
- Wu, W.; Yeh, M.-Y.; and Chen, M.-S. 2018. Deep censored learning of the winning price in the real time bidding. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2526–2535.
- Zhang, W.; Kitts, B.; Han, Y.; Zhou, Z.; Mao, T.; He, H.; Pan, S.; Flores, A.; Gultekin, S.; and Weissman, T. 2021. MEOW: A Space-Efficient Nonparametric Bid Shading Algorithm. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 3928–3936.
- Zhang, W.; Yuan, S.; Wang, J.; and Shen, X. 2014. Real-time bidding benchmarking with ipinyou dataset. arXiv:1407.7073.
- Zhou, T.; He, H.; Pan, S.; Karlsson, N.; Shetty, B.; Kitts, B.; Gligorijevic, D.; Gultekin, S.; Mao, T.; Pan, J.; et al. 2021. An Efficient Deep Distribution Network for Bid Shading in First-Price Auctions. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 3996–4004.