A Bias-Free Revenue-Maximizing Bidding Strategy for Data Consumers in Auction-based Federated Learning

Xiaoli Tang[1] , **Han Yu**[1] , **Zengxiang Li**[2] and **Xiaoxiao Li**[3]
[1]College of Computing and Data Science, Nanyang Technological University, Singapore
[2]ENN Group, Beijing, China
[3]Department of Electrical and Computer Engineering, The University of British Columbia, Canada
{xiaoli001, han.yu}@ntu.edu.sg, lizengxiang@enn.cn, xiaoxiao.li@ece.ubc.ca

Abstract

Auction-based Federated Learning (AFL) is a burgeoning research area. However, existing bidding strategies for AFL data consumers (DCs) primarily focus on maximizing expected accumulated utility, disregarding the more complex goal of revenue maximization. They also only consider winning bids, leading to biased estimates by overlooking information from losing bids. To address these issues, we propose a Bias-free Revenue-maximizing Federated bidding strategy for DCs in AFL (BR-FEDBIDDER). Our theoretical exploration of the relationships between Return on Investment (ROI), bid costs, and utility, and their impact on overall revenue underscores the complexity of maximizing revenue solely by prioritizing ROI enhancement. Leveraging these insights, BR-FEDBIDDER optimizes bid costs with any given ROI constraint. In addition, we incorporate an auxiliary task of winning probability estimation into the framework to achieve bias-free learning by leveraging bid records from historical bid requests, including both winning and losing ones. Extensive experiments on six widely used benchmark datasets show that BR-FEDBIDDER outperforms eight state-of-the-art methods, surpassing the best-performing baseline by 5.66%, 6.08% and 2.44% in terms of the total revenue, ROI, and test accuracy of the resulting FL models, respectively.

1 Introduction

Driven by stringent user privacy and data confidentiality requirements, Federated Learning (FL) has recently garnered substantial attention from both academic and industrial domains [Liu *et al.*, 2022; Liu *et al.*, 2020; Yang *et al.*, 2019]. With data owners (DOs), a.k.a. FL clients, being self-interested entities that weigh a myriad of factors, ranging from costs to potential utility gains, when deciding which FL data consumer (DC) to collaborate with, the design of FL incentive mechanisms [Zhan *et al.*, 2021; Khan *et al.*, 2020] has taken center stage. These mechanisms aim to incentivize DOs to participate in FL through various rewarding strategies.

Auction-based federated learning (AFL) is a pivotal domain in the field of FL incentive mechanism design, primarily due to its potential to achieve both efficiency and fairness. AFL methods can be classified into three categories based on their primary emphasis: 1) DC-side methods, 2) auctioneer-side methods, and 3) DO-side methods [Tang *et al.*, 2024b; Tang *et al.*, 2024a]. The DC-side methods focus on how DCs select and bid for DOs, with the aim of optimizing key performance indicators (KPIs) while adhering to budgetary constraints. The auctioneer-side methods explore the optimization of DC-DO matching and pricing strategies, along with the design of effective auction mechanisms. The overarching goal is to achieve specific operational objectives, such as maximizing social welfare or minimizing social costs for the entire AFL ecosystem. Issues pertaining to DOs revolve around determining the allocation of resources and setting reserve prices to maximize their profits.

Recently, increasing research attention [Tang and Yu, 2023c; Tang and Yu, 2023a; Tang and Yu, 2023b] has been paid to DC-side issues, with optimal bidding methods emerging. Existing bidding strategies for DCs mainly focus on utility maximization, while revenue maximization remains an open research problem. This poses a challenge for DCs aiming to maximize their revenues when bidding for DOs in a competitive AFL marketplace.

In addition, optimal bidding in AFL involves a number of subtasks, including utility estimation, bid cost estimation, and winning price prediction. However, due to systemic differences in data distributions between the inference and training spaces, existing AFL bidding methods suffer from sample selection bias (SSB) [Yang *et al.*, 2021]. They require labeled data for supervised learning, but in the widely adopted generalized second-price (GSP) auction mechanism, only the winning DC can observe the actual utility of bid requests and market price. As a result, existing methods are only trained on winning records, while making inferences on the full-volume space with all bid records, leading to issues of SSB that significantly weaken model performance and generalization.

To address this challenge, we propose the Bias-free Revenue-maximizing Federated bidding strategy for DCs in AFL (BR-FEDBIDDER). We start by exploring the theoretical relationships between return on investment (ROI), bid costs and utility, and how they collectively impact overall revenue. Our analysis reveals that solely focusing on increasing

ROI without considering bid costs might lead to sub-optimal revenue. Based on these findings, we design a novel revenue-maximizing bidding strategy that focuses on bid cost maximization with an acceptable ROI constraint. In addition, we incorporate winning probability estimation as an auxiliary task into BR-FEDBIDDER to achieve bias-free learning by leveraging the abundant unsuccessful bid records in the full-volume historical bid request data, including both winning and losing ones.

Extensive experiments on six benchmark datasets show that BR-FEDBIDDER outperforms eight state-of-the-art approaches. On average, it achieves 5.66% higher total revenue, 6.08% higher ROI, and 2.44% higher FL model accuracy compared to the best baseline.

2 Related Works

Our research primarily aligns with the domain of bidding strategy design for forward auction-based FL DCs. [Tang and Yu, 2023c] proposed the first AFL DC bidding strategy, FedBidder. It considers the constraints of DCs' budgets, the relevance of DOs, and incorporates prior auction-related knowledge (e.g., DO distribution, probability of the DC winning the ongoing auction) into the design of bidding functions. It demonstrates the critical roles played by accurate estimation of DO utility, and the selection of an appropriate winning function in determining the optimal bidding strategy. To deal with the complex relationships among DCs during the auction process, [Tang and Yu, 2023a] frames the AFL ecosystem as a multi-agent system designed to nudge DCs to strategically bid, steering the AFL ecosystem towards equilibria with desirable overall system characteristics.

Nevertheless, existing methods focus on the utility, ignoring how to optimize revenue. In addition, these methods are faced with SSB issues. This paper bridges these important research gaps.

3 Preliminaries

A typical AFL ecosystem involves three primary stakeholders: 1) data owners (DOs), possessing potentially sensitive but valuable data; 2) data consumers (DCs), seeking data for collaboratively training AI models; and 3) an FL auctioneer who is responsible for matching DOs with DCs. The auctioneer initiates an auction, prompting interested DCs to bid for a specific DO's resources. Each DC determines its bid value b based on its objectives, limitations and the perceived utility of the DO's resources. Subsequently, by employing an auction mechanism (e.g., the generalised second-price auction (GSP)), the auctioneer determines the winning DC and the amount it needs to pay. Upon depletion of eligible DOs, the process concludes. Then, each DC commences FL model training with the recruited DOs.

Let N denote the number of qualified bid requests in an AFL ecosystem. Each bid request is a high-dimensional feature vector $\boldsymbol{q}$ following the identical independent distribution (IID) from the prior distribution $p_{\mathcal{Q}(\boldsymbol{q})}$, with entries comprising features related to the DO (e.g., the quantity of the auctioned data samples, the identity of the DO). Then, each DC estimates the potential utility of bid request $\boldsymbol{q}$ if it wins the corresponding auction, upon receiving the transmitted bid request from the auctioneer. Specifically, we use $s(\boldsymbol{q})$ to denote the estimated utility of bid request $\boldsymbol{q}$ to the target DC. Similar to $\boldsymbol{q} \sim p_{\mathcal{Q}(\boldsymbol{q})}$, we denote the prior probability distribution of $s(\boldsymbol{q})$ as $p_{\mathcal{S}}(s)$. Then, based on the estimated utility $s(\boldsymbol{q})$, the target DC calculates its bid price b_i for bid request $\boldsymbol{q}$ according to its adopted bidding strategy (i.e., bidding function $b(s(\boldsymbol{q}))$). After collecting all the bid prices from DCs, the auctioneer decides the corresponding market price z following the GSP auction mechanism, which follows a prior distribution $p_Z(z)$. Given the market price z and the corresponding bid price $b(s(\boldsymbol{q}))$, the probability of the target DC winning the auction can be formulated as $W_Z(b(s(\boldsymbol{q}))))$. The estimated cost of bid request $\boldsymbol{q}$ with bid price $b(s(\boldsymbol{q}))$ is denoted as $c(b(s(\boldsymbol{q})))$.

To ensure the equations in the upcoming sections are solvable, we make the following assumptions: 1) The estimated utility of bid request $\boldsymbol{q}$ to the target DC is solely determined by $\boldsymbol{q}$. 2) The bidding strategy of the target DC is solely determined by the estimated utility $s(\boldsymbol{q})$. 3) The winning probability function is monotonically increasing with respect to (w.r.t.) the bidding price, and is only influenced by the bidding price. Then, we can conclude that bid request $\boldsymbol{q}$ influences the winning probability through the following chain: $\boldsymbol{q} \rightarrow s(\boldsymbol{q}) \rightarrow b(s(\boldsymbol{q})) \rightarrow W_Z(b(s(\boldsymbol{q}))))$.

Under the GSP setting, only winning DCs can observe the utility of $\boldsymbol{q}$ and the corresponding market price, while other DCs can observe their own respective bid prices. Therefore, in the winning set $\mathbb{H}_w$ of each DC, each bid request record is denoted as a quadruple $(\boldsymbol{q}, y, z, b) \in \mathbb{H}_w$, where y is the actual utility of bid request $\boldsymbol{q}$. In the losing set $\mathbb{H}_l$, each record is represented as a tuple $(\boldsymbol{q}, b) \in \mathbb{H}_l$.

4 The Proposed Approach

In this section, we begin by formulating the overall revenue, ROI and bid cost, and analysing their relationships. Building on the analysis, we propose BR-FEDBIDDER. It is based on the multi-task learning framework with winning probability estimation as an auxiliary task to jointly model the utility estimation function and market price function, while avoiding bias by employing both the winning set $\mathbb{H}_w$ and losing set $\mathbb{H}_l$.

4.1 Relationship Analysis

The goal of the target DC is to maximize its overall revenue, which can be reformulated as:

$$Overall\_Revenue = Overall\_ROI \times Overall\_Cost, \quad (1)$$

which means that the overall revenue of the target DC depends on both the overall ROI and the overall bid cost. The bid cost and estimated ROI of the target DC on bid request $\boldsymbol{q}$ can be denoted as:

$$Cost(\boldsymbol{q}) = \mathbb{E}[Cost(b(s(\boldsymbol{q})))], ROI(\boldsymbol{q}) = \frac{W_Z(b(s(\boldsymbol{q})))s(\boldsymbol{q})\zeta}{Cost(\boldsymbol{q})} \quad (2)$$

respectively. Here, ζ represents the value of marginal utility, signifying the unit price of auctioned data resources.

As shown in Eq. (2), the bid cost $Cost(\boldsymbol{q})$ is affected by the bid request $\boldsymbol{q}$ through the corresponding utility $s(\boldsymbol{q})$, which means $\boldsymbol{q} \rightarrow s(\boldsymbol{q}) \rightarrow Cost(\boldsymbol{q})$. As a result, Eq. (2) can be rewritten as:

$$Cost(s) = \mathbb{E}[Cost(b(s))], \quad ROI(s) = \frac{W_Z(b(s))s\zeta}{\mathbb{E}[Cost(s)]}. \quad (3)$$

As shown in Eq. (3), both the bid cost and ROI are affected by the estimated utility s. Therefore, it is essential to analyse the relationship between the ROI and estimated utility, and the bid cost and estimated utility.

1) **Relationship between ROI and Utility**: Suppose the bid price is fixed and does not change with the estimated utility, we obtain $\frac{\partial b(s)}{\partial s} = 0$, $\frac{\partial W_Z(b(s(\boldsymbol{q})))}{\partial s} = 0$, and $\frac{\partial \mathbb{E}[Cost(b(s(\boldsymbol{q})))]}{\partial s} = 0$. Taking the derivation of ROI w.r.t. the estimated utility based on Eq. (3), we obtain:

$$\frac{\partial ROI(s)}{\partial s} = W_Z(b(s))\mathbb{E}^{-1}[Cost(s)]\zeta. \quad (4)$$

Taking each part on the right side of Eq. (4) into consideration, we obtain $\frac{\partial ROI(s)}{\partial s} \geq 0$, which means that the ROI of the bid request increases monotonically w.r.t. the corresponding estimated utility. Thus, the bid request with higher utility increases the ROI of the DC.

2) **Relationship between Bid Cost and Utility**: Let $\bar{s}$ be the minimum acceptable estimated utility of bid requests. Then, we formulate the abridged bidding strategy as

$$b^{abr}(s) = \begin{cases} b(s), & if\ s \geq \bar{s}, \\ 0, & if\ s < \bar{s}, \end{cases}. \quad (5)$$

where $b(s)$ is the bid price for bid requests with the estimated utility $s \geq \bar{s}$; 0 for the case $s < \bar{s}$. This implies the bidding function opts not to bid for bid requests below the specified utility threshold. Given N available bid requests, substituting this abridged bidding function into Eq. (3) yields:

$$\begin{aligned} Overall\_Cost(\bar{s}) &= N \int_s \mathbb{I}[s \geq \bar{s}]\mathbb{E}[Cost(b(s))]p_S(s)ds \\ &= N \int_{\bar{s}}^1 \mathbb{E}[Cost(b(s))]p_S(s)ds. \end{aligned} \quad (6)$$

$\mathbb{I}[\cdot]$ is an indicator function. Based on Eq. (6), we obtain:

$$\frac{\partial Overall\_Cost(\bar{s})}{\partial \bar{s}} = -N\mathbb{E}[Cost(b(s))]p_S(s). \quad (7)$$

According to the definitions of each part on the right-hand side of Eq. (7), we obtain $\frac{\partial Overall\_Cost(\bar{s})}{\partial \bar{s}} \leq 0$. Thus, the overall bid cost monotonically decreases with the increase of the estimated utility of bid requests, implying a cost reduction with bid requests of higher estimated utility.

As shown above, ROI monotonically increases w.r.t. the utility of the corresponding bid request when the bid price is fixed. In this sense, higher ROI will generate better overall revenue of the corresponding DC according to Eq. (1) when the overall cost is fixed. However, as shown in Section 4.1, the overall bid cost reduces with the increase of estimated utility. The major reason for this is the number of higher-utility bid requests tends to be limited in general, leading to severe competition among the DCs. Such severe competition tends to limit the total cost of bid requests with higher utility.

Thus, as the utility of bid requests increases, the overall bid cost and overall ROI exhibit opposing trends. Consequently, further examination is necessary to understand the relationship between estimated utility and overall revenue, which is the product of overall bid cost and overall ROI.

3) **Relationship between Revenue and Utility**: Substituting the $b^{abr}(s)$ in Eq. (5) into Eq. (3) yields:

$$Overall\_ROI(\bar{s}) = \frac{N \int_{\bar{s}}^1 \zeta W_Z(b(s))sp_S(s)ds}{N \int_{\bar{s}}^1 \mathbb{E}[Cost(b(s))]p_S(s)ds}. \quad (8)$$

Substituting $Overall\_ROI(\bar{s})$ defined in Eq. (8) and $Overall\_Cost(\bar{s})$ defined in Eq. (6) into Eq. (1) yields:

$$Overall\_Revenue(\bar{s}) = N \int_{\bar{s}}^1 \zeta W_Z(b(s))sp_S(s)ds. \quad (9)$$

The derivation of $Overall\_Revenue(\bar{s})$ w.r.t. $\bar{s}$ is:

$$\frac{\partial Overall\_Revenue(\bar{s})}{\partial \bar{s}} = -N\zeta W_Z(b(s))\bar{s}p_S(s)ds. \quad (10)$$

According to the definition of each part on the right-hand side of Eq. (10), we can get $\frac{\partial Overall\_Revenue(\bar{s})}{\partial \bar{s}} \leq 0$. This means that the overall revenue decreases monotonically with $\bar{s}$. Therefore, a focus on higher-utility bid requests may restrict revenue improvement.

The above analysis shows that it is difficult to boost revenue by focusing just on higher-utility bid requests and ignoring the balance between the bid cost and ROI. To address this issue, we propose the following revenue-maximizing bidding strategy for the target DC.

4.2 Revenue-maximizing Bidding Strategy

The revenue-maximizing bidding strategy for the target DC is formulated as:

$$b^O() = \text{argmax}\ Overall\_Cost, \quad s.t.\ Overall\_ROI \geq \mathcal{R}, \quad (11)$$

where $b^O()$ is the bidding function, which transforms the original revenue maximization problem into overall bid cost maximization under a given ROI $\mathcal{R}$. To obtain the optimal bidding function $b^O()$ for the target DC, we suppose that if each bid request $\boldsymbol{q}$'s ROI is greater than $\mathcal{R}$, the target DC's overall ROI is no less than $\mathcal{R}$. Based on this assumption, substituting the overall bid cost defined in Eq. (6) and ROI defined in Eq. (3) into Eq. (11) yields:

$$\begin{aligned} b^O() = \underset{b()}{\text{argmax}}\, N \int_s \mathbb{E}[Cost(b(s))]p_S(s)ds \\ s.t.\ \frac{\zeta}{\mathcal{R}} \times \frac{W_Z(b(s))s}{\mathbb{E}[Cost(b(s))]} \geq 1,\ i \in [1, N], \end{aligned} \quad (12)$$

where s is the estimated utility of $\boldsymbol{q}$. However, the ROI condition defined in Eq. (12) is too stringent and cannot be met in practice. Thus, we adopt $\epsilon_s = \max(0, 1 - \frac{\mathcal{R}_s}{\mathcal{R}})$ to relax this requirement, where $\mathcal{R}_s = \frac{W_Z(b(s))s}{\mathbb{E}[Cost(b(s))]}$ is the ROI associated with s. Then, Eq. (12) is reformulated as:

$$\begin{aligned} b^O() = \underset{b()}{\text{argmax}}\, N \int_s \mathbb{E}[Cost(b(s))]p_S(s)ds \\ s.t.\ \frac{\zeta}{\mathcal{R}} \times \frac{W_Z(b(s))s}{\mathbb{E}[Cost(b(s))]} \geq 1 - \epsilon_s. \end{aligned} \quad (13)$$

Let $\mathcal{C}$ denote the hyperparameter for regularization. Then, the loss function of $b^O()$ defined in Eq. (13) is formulated as:

$$\mathcal{L}(b(s)) = -N \int_s (1 - \mathcal{C}\epsilon_s)\mathbb{E}[Cost(b(s))]p_S(s)ds, \quad (14)$$

$$\epsilon_s = \max(0, 1 - \frac{\zeta}{\mathcal{R}} \times \frac{W_{\mathcal{Z}}(b(s))s}{\mathbb{E}[Cost(b(s))]}). \quad (15)$$

Eq. (14) and Eq. (15) indicate that the expected bid cost $\mathbb{E}[Cost(b(s))]$, corresponding winning function $W_{\mathcal{Z}}(b(s))$ and estimated utility s of the bid request by the target DC determine the optimal bidding function $b^O()$. We discuss how to obtain these three functions in the following sections.

Bid Cost & Winning Probability. Under GSP, the target DC wins only if its bid price $b(s)$ exceeds the corresponding market price z. Moreover, if it wins, it pays z as the actual cost; otherwise, it pays nothing. Therefore, the bid cost $c(b(s))$ can be expressed as:

$$Cost(b(s)) = \begin{cases} z, & \text{if } b(s) > z, \\ 0, & \text{otherwise} \end{cases} \quad (16)$$

Then, as $z \sim p_Z(z)$, we formulate $c(b(s))$ as:

$$\begin{aligned} \mathbb{E}[Cost(b(s))] &= \int_0^{b(s)} z p_Z(z)dz + \int_{b(s)}^{+\infty} 0 * p_Z(z)dz \\ &= \int_0^{b(s)} z p_Z(z)dz. \end{aligned} \quad (17)$$

Given $\boldsymbol{q}$, the probability of the target DC winning the auction (i.e., $W_Z(b(s))$) is defined as the probability that its bid price $b(s)$ exceeds the corresponding market price z:

$$W_Z(b(s)) = \int_0^{b(s)} p_Z(z)dz. \quad (18)$$

Then, the cost estimation and winning function modeling are transformed into the market price modeling. Therefore, the optimal bidding function depends on the utility estimation and market price modeling. It is intuitive to address these two issues independently. However, they are interdependent. On the one hand, as shown in [Tang and Yu, 2023c], estimated utility and bid price are positively related. On the other hand, under the GSP auction mechanism, the market price is the highest bid price submitted by competing DCs. Thus, the market price has a significant positive correlation with the estimated utility of data resources in theory. This makes it sub-optimal to perform these estimation tasks separately.

To deal with this issue, we design a multi-task learning-based framework to jointly optimize the utility estimation and market price modeling simultaneously, which is explained in detail in the following section.

4.3 Multi-Task Learning-based Utility Estimation & Market Price Modeling

The framework for jointly optimizing the utility estimation and market price modeling consists of four main modules: 1) the interaction modeling module captures the high-order interactions among features of bid requests for better-shared bid request representations for the following modules; 2) the utility estimation module performs the utility estimation task based on the shared feature representation; 3) the market price modeling module executes the task of modeling market prices by leveraging the shared feature; and 4) the winning probability prediction module predicts the winning probability of a specific bid price based on the output of the market price modeling module.

1) **Interaction Modeling Module**: To capture the non-linear interaction patterns among features of bid requests, we then utilize the hard parameter sharing method widely adopted in multi-task learning approaches to further map the representation $\boldsymbol{q}$ of each bid request into the shared representation vector $\boldsymbol{q}'$, i.e.:

$$\boldsymbol{q}' = f_{srm}(\boldsymbol{q}), \quad (19)$$

where $f_{srm}(\cdot)$ represents the mapping function (i.e., hard-parameter sharing layers), and consists of a number of hidden layers, each of whose parameters are jointly trained by the three tasks that come after.

2) **Utility Estimation Module**: Based on the shared feature representation $\boldsymbol{q}'$ output by the upstream module, this module formulates the utility estimation as a regression task, which is defined as $s(\boldsymbol{q}) = f_{ue}(\boldsymbol{q}') = s$. $f_{ue}(\cdot)$ represents the mapping function of the utility estimation module. Then, the widely used squared error (SE) loss is adopted to train the model:

$$L_{ue} = -\frac{1}{2} \sum_{(\boldsymbol{q},y)\in\mathbb{H}_w} (y - s(\boldsymbol{q}))^2, \quad (20)$$

where y is the actual utility of the bid request $\boldsymbol{q}$.

3) **Market Price Modeling Module**: As bid prices are discrete integers within a specific range under AFL, we incorporate a Softmax layer into this module to take the predicted probability density of each possible market price value as the market price distribution output result [Yang *et al.*, 2021]. Specifically, given a market price $z = k$, the mapping function $f^k_{mp}(\cdot)$ first maps the representation $\boldsymbol{q}'$ generated by the interaction module to a score. Then, the normalized probability density $p(z = k|\boldsymbol{q})$ is given by: $p(z = k|\boldsymbol{q}) = \frac{\exp(f^k_{mp}(\boldsymbol{q}'))}{\sum_{z=0}^{z=b_{max}} \exp(f^z_{mp}(\boldsymbol{q}'))}$, where b_{max} represents the maximum potential bid price. To train this module, the following loss function is adopted:

$$L_{mp} = -\sum_{(\boldsymbol{q},z)\in\mathbb{H}_w} \log p(z|(\boldsymbol{q}))). \quad (21)$$

4) **Winning Probability Prediction Module**: Eq. (20) and Eq. (21) show that only the winning cases are used for the training of both the utility estimation task and the market price modeling task. This might result in sample selection biases. To deal with this problem, we take advantage of the bid information generated under both the winning cases and the losing cases in the multi-task learning framework, and incorporate the winning probability prediction as the auxiliary task. As described above, the winning probability is defined as the likelihood of the bid price b exceeding the corresponding market price z. Given each market price probability density function $p(z|(\boldsymbol{q})))$ obtained above, a specific bid price b's winning probability $W_Z(b(s))$ is derived as:

$$W_Z(b(s)|\boldsymbol{q}) \triangleq p(z < b|\boldsymbol{q}) = \sum_{k=0}^{b-1} p(z = k|\boldsymbol{q}). \quad (22)$$

Then, we can directly obtain the winning probability of any bid price b on the basis of the output of market price distribution $p(z|\boldsymbol{q})$ in the market price modeling module according to Eq. (22) without additional parameter.

However, under the GSP auction mechanism, the market price is observable only when the DC successfully wins the corresponding bid request by exceeding the market price. Conversely, if the DC loses the auction, it is only aware that the market price is equal to or greater than its bid price, represented as $z \geq b$. To ensure unbiased estimation, the proposed winning probability prediction module is designed to maximize the likelihood of both winning and losing scenarios, effectively utilizing the complete bid information available in the entire inference space. Consequently, the loss function for this module is formulated as:

$$\begin{aligned} L_{wp} &= -\log \sum_{(\boldsymbol{q},b)\in\mathbb{H}_w} p(b > z|(\boldsymbol{q}))) \\ &\quad - \log \sum_{(\boldsymbol{q},b)\in\mathbb{H}_l} p(b \leq z|(\boldsymbol{q}))) \\ &= -\log \sum_{(\boldsymbol{q},b)\in\mathbb{H}_{train}} [W_Z(b(s)|\boldsymbol{q})]^w - [1 - W_Z(b(s)|\boldsymbol{q})]^{1-w} \\ &= -\sum_{(\boldsymbol{q},b)\in\mathbb{H}_t} (w \log(W_Z(b(s)|\boldsymbol{q})) \\ &\quad + (1-w)\log[1 - W_Z(b(s)|\boldsymbol{q})]), \end{aligned} \tag{23}$$

where $w = 1$ if the DC wins $\boldsymbol{q}$; otherwise, $w = 0$. $\mathbb{H}_t = \mathbb{H}_w \cap \mathbb{H}_l$. Eq. (23) shows that L_{wp} is used to predict all training samples' winning probability.

The market price modeling module is trained across the entire inference space to directly rectify the bias by introducing the winning probability prediction as an auxiliary task into the multi-task learning framework. Moreover, the utility estimation module is also able to execute bias-free training indirectly via the shared representation, benefiting from multi-task learning. In this sense, the sample selection bias issue can be adequately mitigated by winning probability prediction (i.e., the auxiliary task).

Overall Training Objective. The objective of the proposed multi-task learning-based framework is the joint optimization of both the utility estimation and market price modeling, taking advantage of the auxiliary winning probability prediction. The overall training goal to be optimized is formulated as the combination of all these losses described above: $L(\theta) = w_{ue} \cdot L_{ue} + w_{mp} \cdot L_{mp} + w_{wp} \cdot L_{wp}$, where θ represents all the parameters involved in the multi-task learning framework. w_{ue}, w_{mp}, and w_{wp} are hyperparameters used to control the importance of different tasks.

4.4 Optimal Bidding Strategy for Data Consumers

After estimating the utility and modeling the market price based on the multi-task learning framework, we can get the predicted utility of each bid request, denoted as $H = \{s_1, \cdots, s_m, \cdots, \}$, where s_m denotes the estimated utility of bid request $\boldsymbol{q}_m$. Then, the loss function defined in Eq. (14) and (15) can be re-expressed as: $\mathcal{L}(b(s)) = -\sum_{s\in H}(1 - \mathcal{C}\epsilon_s)\mathbb{E}[Cost(b(s))]$, where ϵ_s is denoted as: $\epsilon_s = \max(0, 1 - \frac{\zeta}{\mathcal{R}} \times \frac{W_{\mathcal{Z}}(b(s))s}{\mathbb{E}[Cost(b(s))]})$. Based on Eq. (17), we get $\frac{\partial \mathbb{E}[Cost(b(s))]}{\partial b(s)} = b(s)p_{\mathcal{Z}}(b(s))$. Then, $\frac{\partial \mathcal{L}(b(s))}{\partial b(s)}$ is:

$$\begin{aligned} \frac{\partial \mathcal{L}(b(s))}{\partial b(s)} &= -\sum_{s\in H}\{b(s)p_{\mathcal{Z}}(b(s)) \\ &\times (1 - \mathbb{I}[\epsilon_s > 0]\mathbb{C}(1 - \zeta\frac{W_{\mathcal{Z}}(b(s))s}{\mathcal{R}\mathbb{E}[Cost(b(s))]}) \\ &+ \mathbb{I}[\epsilon_s > 0]\mathbb{C}\frac{\zeta s}{\mathcal{R}} \times (p_{\mathcal{Z}}(b(s)) - \frac{W_{\mathcal{Z}}(b(s))}{\mathbb{E}[Cost(b(s))]}b(s)p_{\mathcal{Z}}(b(s)))\} \end{aligned} \tag{24}$$

As shown in Eq. (24), we need to define the detailed formulation of the bidding strategy $b(s)$. Following [Zhang *et al.*, 2014; Perlich *et al.*, 2012; Zhang and Wang, 2015; Ren *et al.*, 2017], we adopt the following bidding function: 1) The linear bidding function $b_1(s)$ (BR-FEDBIDDER1):

$$b_1(s) = \omega \times s. \tag{25}$$

2) The concave shape bidding function $b_2(s)$ (BR-FEDBIDDER2):

$$b_2(s) = \omega_1\sqrt{s + \omega_2^2} - \omega_2. \tag{26}$$

Alg. 1 provides the learning process for optimizing ω in the bidding function. The complete learning procedure of BR-FEDBIDDER is as follows: 1) Utilize the multi-task learning-based framework to model the utility and market price of the bid request being auctioned. 2) Estimate the bid cost and winning probability using Eqs. (17) and (18). 3) Optimize the parameter of the bidding function following Alg. 1.

Algorithm 1 The learning procedure of parameter ω in the bidding function of BR-FEDBIDDER.

INPUT: Training Data H, the multi-task learning framework, Learning rate η_ω
OUTPUT: Optimal bidding function $b^O()$

```
Select the bidding function defined in Eq. (25) or (26)
Initialize ω of the selected bidding function
for training rounds do
    for s ∈ H do
        Get ∂L(b(s))/∂b(s) according to Eq. (24)
        Get ∂b(s)/∂ω according to the selected bidding function
        ω ← ω − η_ω (∂L(b(s))/∂b(s)) (∂b(s)/∂ω)
    end for
end for
return b^O()
```

5 Experimental Evaluation

5.1 Experiment Setup

Datasets. We conduct experimental evaluation using commonly adopted datasets in FL studies, including MNIST (http://yann.lecun.com/exdb/mnist/), CIFAR-10 (https://www.cs.toronto.edu/kriz/cifar.html), Fashion-MNIST (i.e., FMNIST) [Xiao *et al.*, 2017], EMNIST-digits (i.e., EMNIST-D), EMNIST-letters (i.e., EMNIST-L) [Cohen *et al.*, 2017] and Kuzushiji-MNIST (i.e., KMNIST) [Clanuwat *et al.*, 2018]. The FL models used are the same as those employed in [Tang and Yu, 2023c].

Bud	Method	MNIST			CIFAR			FMNIST			EMNISTD			EMNISTL			KMNIST		
		Rev	ROI	Acc	Rev	ROI	Acc	Rev	ROI	Acc	Rev	ROI	Acc	Rev	ROI	Acc	Rev	ROI	Acc
50	Const	9.92	0.22	84.77	9.53	0.21	45.00	11.38	0.25	73.57	11.44	0.25	82.72	7.73	0.17	74.32	7.98	0.18	73.11
	Rand	11.03	0.23	84.65	12.10	0.25	44.95	11.60	0.23	74.78	12.15	0.25	82.34	10.50	0.22	74.62	10.57	0.22	71.66
	Bmub	10.40	0.21	85.06	11.04	0.22	45.79	11.94	0.25	74.74	11.61	0.24	82.78	9.06	0.19	74.57	9.63	0.20	75.61
	Lin	11.96	0.24	85.05	13.34	0.27	46.66	14.02	0.29	75.29	12.93	0.27	83.03	11.37	0.24	74.87	13.24	0.27	73.77
	BM	14.12	0.29	85.01	14.63	0.30	46.58	14.92	0.31	76.27	13.95	0.29	82.78	14.62	0.30	74.83	13.19	0.27	76.77
	FBs	16.42	0.33	85.32	16.87	0.35	47.05	16.80	0.34	77.50	15.41	0.32	83.55	16.42	0.34	75.17	16.19	0.34	76.37
	FBc	16.33	0.33	84.98	16.47	0.35	47.62	17.01	0.35	77.17	15.99	0.33	83.37	16.26	0.33	74.54	16.24	0.34	76.43
	RLB	17.31	0.35	85.83	16.97	0.34	47.86	17.62	0.36	77.77	16.42	0.34	83.60	16.80	0.34	74.91	17.81	0.36	77.38
	FMRB1	17.61	**0.37**	87.49	**17.26**	**0.36**	49.51	17.63	0.36	78.91	**16.91**	0.35	**84.86**	**17.13**	**0.35**	75.42	17.87	**0.38**	78.68
	FMRB2	**17.65**	0.36	**87.86**	17.20	0.35	**49.98**	**17.67**	**0.37**	**79.60**	16.77	**0.36**	84.78	17.12	**0.35**	**75.66**	**18.08**	**0.38**	**78.78**
150	Const	14.13	0.10	85.27	14.51	0.10	45.53	14.07	0.10	74.27	12.20	0.08	82.80	12.05	0.08	74.42	13.87	0.10	73.56
	Rand	13.88	0.10	85.04	13.41	0.09	45.50	15.54	0.11	74.89	12.82	0.09	82.61	10.23	0.07	74.86	16.36	0.11	72.68
	Bmub	17.05	0.11	85.53	14.02	0.09	46.84	16.18	0.11	75.14	14.97	0.10	83.26	15.00	0.10	74.60	15.09	0.10	75.25
	Lin	18.99	0.13	85.58	18.33	0.12	47.01	18.46	0.12	76.73	17.67	0.12	83.33	18.46	0.12	75.11	18.60	0.12	74.38
	BM	24.35	0.16	85.59	23.08	0.16	47.82	24.81	0.17	77.37	22.86	0.15	83.26	23.32	0.16	75.23	24.65	0.17	77.27
	FBs	37.81	0.25	85.88	35.73	0.24	48.17	37.75	0.25	77.83	35.28	0.24	83.93	35.80	0.24	75.33	36.63	0.25	77.29
	FBc	41.78	0.28	85.76	39.94	0.27	48.86	42.89	0.29	77.46	39.07	0.26	83.72	40.00	0.27	75.22	40.19	0.27	76.55
	RLB	46.08	0.31	86.55	45.37	0.31	49.09	49.25	0.33	77.48	45.98	0.31	83.94	44.28	0.30	75.58	48.47	0.33	78.01
	FMRB1	**50.31**	**0.34**	87.81	**47.64**	**0.32**	50.50	51.10	**0.35**	78.92	48.42	**0.32**	84.71	**49.37**	**0.33**	**76.30**	50.88	**0.34**	**78.97**
	FMRB2	50.06	**0.34**	**88.71**	47.33	**0.32**	**51.85**	**51.17**	**0.35**	**78.97**	**49.75**	**0.32**	**85.22**	48.58	**0.33**	75.97	**51.00**	**0.34**	78.36
300	Const	15.69	0.05	85.38	16.70	0.06	46.11	17.74	0.06	74.50	17.87	0.06	82.96	14.66	0.05	74.58	17.13	0.06	74.13
	Rand	25.20	0.09	85.12	30.62	0.10	45.88	30.39	0.10	75.10	26.03	0.09	82.69	27.06	0.09	74.95	26.60	0.09	72.02
	Bmub	25.22	0.08	85.73	27.61	0.09	47.44	28.47	0.10	75.13	26.12	0.09	83.13	25.73	0.09	74.79	27.18	0.09	74.69
	Lin	29.99	0.10	85.82	29.43	0.10	47.69	30.47	0.10	76.91	29.92	0.10	83.36	29.29	0.10	75.11	29.90	0.10	75.30
	BM	39.03	0.13	86.25	36.77	0.12	48.63	37.28	0.12	78.19	36.61	0.12	83.13	36.61	0.12	75.41	36.61	0.12	77.80
	FBs	69.76	0.23	86.32	64.73	0.22	49.02	67.31	0.23	78.30	64.56	0.22	84.06	64.91	0.22	75.46	64.81	0.22	76.93
	FBc	72.08	0.24	86.19	68.07	0.23	48.91	72.88	0.25	78.07	68.23	0.23	83.98	67.00	0.22	75.67	67.31	0.23	76.18
	RLB	86.08	0.29	86.86	78.83	0.26	49.69	85.64	0.29	78.58	79.53	0.27	84.13	79.53	0.27	75.60	79.53	0.27	78.31
	FMRB1	**93.05**	**0.31**	88.08	85.18	**0.29**	51.75	92.61	**0.31**	**79.99**	85.95	**0.29**	85.39	87.44	**0.30**	76.73	**88.17**	**0.30**	79.96
	FMRB2	92.87	**0.31**	**89.98**	**85.30**	**0.29**	**52.05**	**92.82**	**0.31**	79.74	**86.06**	**0.29**	**85.54**	**87.69**	0.29	**76.76**	87.69	0.29	**80.57**

Table 1: Comparison results across different budget settings and datasets under the IID settings.

Bud	Method	MNIST	CIFAR	FMNIST	EMNISTD	EMNISTL	KMNIST
50	Const	75.04	21.86	65.51	78.10	66.44	69.31
	Rand	74.80	20.15	65.17	77.95	65.77	68.32
	Bmub	74.96	22.15	65.18	78.05	67.04	68.50
	Lin	75.39	24.45	65.72	78.29	68.35	68.78
	BM	75.79	27.00	65.98	79.09	69.55	69.70
	FBs	76.51	27.51	66.14	79.35	69.43	70.06
	FBc	75.48	26.56	65.79	79.14	69.28	69.67
	RLB	76.20	28.02	66.43	79.66	70.15	70.11
	FMRB1	**77.92**	**29.95**	67.19	**80.62**	70.80	71.83
	FMRB2	77.78	29.70	**67.59**	79.96	**71.11**	**72.90**
150	Const	75.03	23.24	65.16	78.23	66.59	69.38
	Rand	75.01	22.80	65.29	78.12	65.79	68.55
	Bmub	75.49	23.75	65.68	78.34	67.30	69.16
	Lin	75.70	27.08	65.89	78.65	68.63	69.05
	BM	76.23	28.33	66.44	79.23	70.16	70.22
	FBs	76.52	29.48	66.46	79.64	70.49	70.78
	FBc	75.94	28.34	66.57	79.06	69.96	69.74
	RLB	76.73	30.42	66.78	80.07	70.98	70.82
	FMRB1	78.20	32.22	**68.24**	81.24	71.25	72.18
	FMRB2	**78.37**	**32.31**	68.07	**81.46**	**71.90**	**73.32**
300	Const	75.02	24.67	65.68	78.28	67.32	69.21
	Rand	74.95	24.66	65.40	78.10	65.95	68.69
	Bmub	75.47	25.99	65.83	78.50	67.79	69.85
	Lin	76.03	28.13	66.26	78.99	68.69	69.32
	BM	76.20	28.22	66.51	79.48	70.43	70.51
	FBs	76.57	30.03	66.93	79.68	70.29	70.86
	FBc	76.25	28.95	66.67	79.36	69.98	70.05
	RLB	76.43	30.41	66.82	80.21	71.05	70.96
	FMRB1	79.50	**32.89**	68.34	**81.73**	71.68	**73.59**
	FMRB2	**80.27**	32.51	**69.40**	81.99	**72.93**	72.52

Table 2: Accuracy (%) comparison under Non-IID scenarios.

Comparison Methods. We compare the two variants of BR-FEDBIDDER: **fmrb1** based on Eq. (25) and **fmrb2** based on Eq. (26) with eight established bidding approaches: **Constant Bid (Const)** [Zhang *et al.*, 2014], **Randomly Generated Bid (Rand)** [Zhang *et al.*, 2021; Zhang *et al.*, 2022], **Below Max Utility Bid (Bmub)**, **Linear-Form Bid (Lin)** [Perlich *et al.*, 2012], **Bidding Machine (BM)** [Ren *et al.*, 2017], and **Reinforcement Learning-based Bid (RLB)** [Cai *et al.*, 2017]. Details of these methods can be found in [Tang and Yu, 2023a]. Moreover, we also include **Fed-Bidder** [Tang and Yu, 2023c] for comparison. Tailored for data consumers in AFL settings, Fed-Bidder aims to competitively bid for data owners to maximize the utility. It offers two variants: one with a simple winning function, **FedBidder-sim (FBs)**, and the other with a complex winning function, **Fed-Bidder-com (FBc)**.

Experiment Scenarios. We compare BR-FEDBIDDER with baselines in two scenarios, each containing 10,000 DOs: 1) **IID data**: Following [Tang and Yu, 2023c], we assign sequential IDs to DOs and introduce Gaussian noise to the data resources of the first half of the population, creating low-quality resources, while the second half remains unchanged with high-quality data. Dataset sizes of various DOs vary randomly from 500 to 5,000 samples. 2) **Non-IID data**: Following [Shi and Yu, 2023], one class (for datasets other than EMNISTL) or six classes (for EMNISTL) is designated as the minority class, assigned to the first 100 DOs. These DOs possess images for all classes, while others exclusively have images for the remaining classes. Each DO has 3,000 images.

Implementation Details. To gather historical AFL bidding data, we employed a data accumulation approach, observing the behaviors of 160 DCs and 10,000 DOs in four distinct settings. In the first, each baseline bidding method was used by one-eighth of the DCs. In the second, four strategies, including BM, FBs, FBc, and RLB with AI techniques like BR-FEDBIDDER, were adopted by three-sixteenths of the population each, while the remaining baselines were chosen by one-sixteenth each. In tailored AFL settings, FBs and FBc were favored. In the third setting, 50 DCs opted for FBs and FBc, while 10 DCs selected each of the other six baselines. In the fourth, 65 DCs chose FBs and FBc, with 5 DCs each for the other six baselines. Using the GSP forward auction mechanism, we continuously monitored DC behaviors to gradually

Method	MNIST		CIFAR-10		FMNIST		EMNISTD		EMNISTL		KMNIST	
	IID	NIID	IID	NIID	IID	NIID	IID	NIID	IID	NIID	IID	NIID
BR-FEDBIDDER	87.08	77.50	50.75	32.89	79.98	67.34	84.39	80.26	75.72	71.68	79.14	71.59
BR-FEDBIDDER-W	85.98	77.27	49.14	32.53	79.74	65.40	82.54	77.27	73.76	70.93	77.56	71.52

Table 3: Ablation study results in terms of test accuracy (%) of the resulting FL models.

accumulate data, compensating for the absence of a dedicated AFL dataset.

To evaluate the effectiveness of BR-FEDBIDDER, we create a group of ten FL DCs, each utilizing one of the aforementioned bidding approaches to join the auction for each bid request from the available FL DOs. Following [Tang and Yu, 2023c], bid requests are delivered in chronological order. Upon receiving a bid request, each DC derives its bid price based on its adopted bidding strategy. Then, the auctioneer gathers the bid prices, identifies the winner, and determines the market price using the GSP auction mechanism. The winning DC remits the market price to the DO, and the data aligns with the winning DC for FL model training. The process concludes when there are no more bid requests or when the budget is depleted.

Evaluation Metrics. Our study aims to optimize total revenue while maintaining a satisfactory ROI for the data consumer with the objective of achieving satisfactory FL model performance. Therefore, we adopt the **revenue (Rev)**, **ROI** and the **accuracy (Acc)** of the FL models as the evaluation metrics.

5.2 Results and Discussion

To evaluate the performance of the proposed BR-FEDBIDDER, we perform experiments on six datasets by varying the budget among $\{50, 150, 300\}$ under the IID and Non-IID scenarios. The results are shown in Tables 1 and 2.

Table 1 shows the comparison results for revenue, ROI, and FL model accuracy under the IID scenario. Across all six datasets and three budget settings, our proposed BR-FEDBIDDER consistently outperforms baseline methods. Specifically, compared to the best-performing baseline, BR-FEDBIDDER achieves a 5.66% improvement in revenue. In addition, BR-FEDBIDDER consistently achieves satisfactory ROI, with a 6.08% improvement over the best-performing baseline. These results demonstrate the efficacy of our approach in optimizing the revenue of data consumers while ensuring a satisfactory ROI. The accuracy results in Table 1 align with the performance in revenue and ROI, with BR-FEDBIDDER achieving a 2.01% improvement under the IID scenario.

The comparative results under the Non-IID scenario can be found in Table 2. It can be observed that under the Non-IID scenario, the proposed method BR-FEDBIDDER consistently outperforms existing methods in terms of FL model accuracy. In particular, on average, BR-FEDBIDDER achieves 2.87% higher FL model accuracy compared to the best performance achieved by baselines under the Non-IID data scenario. All these results demonstrate the effectiveness of our approach BR-FEDBIDDER in helping DCs optimize their key performance indicators and bidding strategies for DOs under the emerging AFL scenarios.

Lin and Bmub typically outperform Const and Rand due to their use of utility in the bidding process. However, Bmub is less effective than Lin due to its inclusion of randomness. Meanwhile, the more advanced methods BM, FBs, FBc, RLB, and BR-FEDBIDDER perform significantly better than the simpler approaches. This is largely due to the inclusion of winning probability estimation and the use of machine learning and reinforcement learning frameworks.

RLB and BR-FEDBIDDER both outperform BM, FBs, and FBc, likely due to their direct consideration of bid cost and available budget. While BM does consider market price distribution, it derives this distribution by marginalizing the prediction of the market price density of each bid request, which may lead to overfitting. In contrast, BR-FEDBIDDER obtains the market price distribution by the predefined winning function, which helps predict the expectation of the bid cost more accurately. While RLB uses dynamic programming to optimize its bidding process, its immediate reward setting can lead to blindly bidding for data samples without considering cost, which is not an issue with BR-FEDBIDDER.

Ablation Study. To compare the impact of using full-volume bidding records versus only winning records, we compare the performance of the BR-FEDBIDDER approach with its variant, BR-FEDBIDDER-W. The key difference lies in the loss function used in the winning probability prediction module (Eq. (23)). BR-FEDBIDDER-W employs a modified version of the loss function considering only winning records, while BR-FEDBIDDER uses the entire bidding data. Table 3 displays the results of our experiments. We note that BR-FEDBIDDER outperforms BR-FEDBIDDER-W consistently across all six datasets, both under IID and NONIID settings. This indicates that the auxiliary task allows training the model on the entire inference space, thus effectively mitigating the SSB issue, leading to significantly improved test accuracy.

6 Conclusions

This paper focuses on providing the optimal bidding strategies for data consumers in the AFL ecosystem to guide them in bidding for data owners in the competitive market so as to maximize their revenue and improve FL model accuracy by proposing a novel revenue-maximizing bidding strategy (BR-FEDBIDDER). This paper first theoretically explores the relationships between Return on Investment (ROI) and utility, as well as bid cost and utility, and their impact on overall revenue, which highlights the challenge of revenue maximization when solely prioritizing ROI enhancement without taking into account the impact of bid cost. Based on these findings, BR-FEDBIDDER maximizes bid costs with a targeted ROI constraint. In this process, it incorporates an ancillary task of winning probability estimation into the framework to achieve bias-free learning by leveraging bid records from the entire volume of bid requests.

Acknowledgments

This research is supported, in part, by the National Research Foundation Singapore and DSO National Laboratories under the AI Singapore Programme (No. AISG2-RP-2020-019); Joint NTU-WeBank Research Centre on Fintech; and the RIE 2020 Advanced Manufacturing and Engineering (AME) Programmatic Fund (No. A20G8b0102), Singapore.

References

[Cai *et al.*, 2017] Han Cai, Kan Ren, Weinan Zhang, Kleanthis Malialis, Jun Wang, Yong Yu, and Defeng Guo. Real-time bidding by reinforcement learning in display advertising. In *Proceedings of the 10th ACM International Conference on Web Search and Data Mining (WSDM'17)*, pages 661–670, 2017.

[Clanuwat *et al.*, 2018] Tarin Clanuwat, Mikel Bober-Irizar, Asanobu Kitamoto, Alex Lamb, Kazuaki Yamamoto, and David Ha. Deep learning for classical japanese literature. *arXiv preprint*, page 1812.01718, 2018.

[Cohen *et al.*, 2017] Gregory Cohen, Saeed Afshar, Jonathan Tapson, and Andre Van Schaik. EMNIST: Extending MNIST to handwritten letters. In *Proceedings of the 2017 International Joint Conference on Neural Networks (IJCNN'17)*, pages 2921–2926, 2017.

[Khan *et al.*, 2020] Latif U Khan, Shashi Raj Pandey, Nguyen H Tran, Walid Saad, Zhu Han, Minh NH Nguyen, and Choong Seon Hong. Federated learning for edge networks: Resource optimization and incentive mechanism. *IEEE Communications Magazine*, 58(10):88–93, 2020.

[Liu *et al.*, 2020] Yang Liu, Anbu Huang, Yun Luo, He Huang, Youzhi Liu, Yuanyuan Chen, Lican Feng, Tianjian Chen, Han Yu, and Qiang Yang. Fedvision: An online visual object detection platform powered by federated learning. In *Proceedings of the 32nd Annual Conference on Innovative Applications of Artificial Intelligence (IAAI-20)*, pages 13172–13179, 2020.

[Liu *et al.*, 2022] Zelei Liu, Yuanyuan Chen, Yansong Zhao, Han Yu, Yang Liu, Renyi Bao, Jinpeng Jiang, Zaiqing Nie, Qian Xu, and Qiang Yang. Contribution-aware federated learning for smart healthcare. In *Proceedings of the 34th Annual Conference on Innovative Applications of Artificial Intelligence (IAAI-22)*, pages 12396–12404, 2022.

[Perlich *et al.*, 2012] Claudia Perlich, Brian Dalessandro, Rod Hook, Ori Stitelman, Troy Raeder, and Foster Provost. Bid optimizing and inventory scoring in targeted online advertising. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'12)*, pages 804–812, 2012.

[Ren *et al.*, 2017] Kan Ren, Weinan Zhang, Ke Chang, Yifei Rong, Yong Yu, and Jun Wang. Bidding machine: Learning to bid for directly optimizing profits in display advertising. *IEEE Transactions on Knowledge and Data Engineering*, 30(4):645–659, 2017.

[Shi and Yu, 2023] Yuxin Shi and Han Yu. Fairness-aware client selection for federated learning. In *Proceedings of IEEE International Conference on Multimedia and Expo 2023 (ICME'23)*, 2023.

[Tang and Yu, 2023a] Xiaoli Tang and Han Yu. Competitive-cooperative multi-agent reinforcement learning for auction-based federated learning. In *IJCAI*, 2023.

[Tang and Yu, 2023b] Xiaoli Tang and Han Yu. Multi-session budget optimization for forward auction-based federated learning. *arXiv preprint arXiv:2311.12548*, 2023.

[Tang and Yu, 2023c] Xiaoli Tang and Han Yu. Utility-maximizing bidding strategy for data consumers in auction-based federated learning. In *Proceedings of the 2023 IEEE International Conference on Multimedia and Expo (ICME'23)*, pages 330–335, 2023.

[Tang *et al.*, 2024a] Xiaoli Tang, Han Yu, and Xiaoxiao Li. Agent-oriented joint decision support for data owners in auction-based federated learning. In *ICME*, 2024.

[Tang *et al.*, 2024b] Xiaoli Tang, Han Yu, Xiaoxiao Li, and Sarit Kraus. Intelligent agents for auction-based federated learning: A survey. *arXiv preprint arXiv:2404.13244*, 2024.

[Xiao *et al.*, 2017] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms. *arXiv preprint*, page 1708.07747, 2017.

[Yang *et al.*, 2019] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2):1–19, 2019.

[Yang *et al.*, 2021] Haizhi Yang, Tengyun Wang, Xiaoli Tang, Qianyu Li, Yueyue Shi, Siyu Jiang, Han Yu, and Hengjie Song. Multi-task learning for bias-free joint ctr prediction and market price modeling in online advertising. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management (CIKM'21)*, pages 2291–2300, 2021.

[Zhan *et al.*, 2021] Yufeng Zhan, Jie Zhang, Zicong Hong, Leijie Wu, Peng Li, and Song Guo. A survey of incentive mechanism design for federated learning. *IEEE Transactions on Emerging Topics in Computing*, 10(2):1035–1044, 2021.

[Zhang and Wang, 2015] Weinan Zhang and Jun Wang. Statistical arbitrage mining for display advertising. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'15)*, pages 1465–1474, 2015.

[Zhang *et al.*, 2014] Weinan Zhang, Shuai Yuan, and Jun Wang. Optimal real-time bidding for display advertising. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'14)*, pages 1077–1086, 2014.

[Zhang *et al.*, 2021] Jingwen Zhang, Yuezhou Wu, and Rong Pan. Incentive mechanism for horizontal federated learning based on reputation and reverse auction. In *Proceed-*

ings of the Web Conference 2021 (WWW'21), pages 947–956, 2021.

[Zhang *et al.*, 2022] Jingwen Zhang, Yuezhou Wu, and Rong Pan. Online auction-based incentive mechanism design for horizontal federated learning with budget constraint. *arXiv preprint*, page 2201.09047, 2022.