

Reputation-aware Revenue Allocation for Auction-based Federated Learning

Xiaoli Tang, Han Yu

College of Computing and Data Science, Nanyang Technological University, Singapore
{xiaoli001, han.yu}@ntu.edu.sg

Abstract

Auction-based Federated Learning (AFL) has gained significant research interest due to its ability to incentivize data owners (DOs) to participate in FL model training tasks of data consumers (DCs) through economic mechanisms via the auctioneer. One of the critical research issues in AFL is decision support for the auctioneer. Existing approaches are based on the simplified assumption of a single, monopolistic AFL marketplace, which is unrealistic in real-world scenarios where multiple AFL marketplaces can co-exist and compete for the same pool of DOs. In this paper, we relax this assumption and frame the AFL auctioneer decision support problem from the perspective of helping them attract participants in a competitive AFL marketplace environment while safeguarding profit. To achieve this objective, we propose the Auctioneer Revenue Allocation Strategy for AFL (ARAS-AFL). We design the concepts of the attractiveness and competitiveness from the perspective of auctioneer reputation. Based on the Lyapunov optimization, ARAS-AFL helps individual AFL auctioneer achieve the dual objective of balancing the reputation management costs and its own profit by designing a dynamic revenue allocation strategy. It takes into account both the auctioneer’s revenue and the changes in the number of participants on the AFL marketplace. Through extensive experiments on widely used benchmark datasets, ARAS-AFL demonstrates superior performance compared to state-of-the-art approaches. It outperforms the best baseline by 49.06%, 98.69%, 10.32%, and 4.77% in terms of total revenue, number of data owners, public reputation and accuracy of federated learning models, respectively.

Introduction

With growing concerns over data privacy and user confidentiality, federated learning (FL) (Yang et al. 2019; Hu et al. 2024; Guo et al. 2021; Tang et al. 2024c; Li et al. 2021) has gained significant interest in recent years due to its versatile applications (Sun et al. 2024). In FL settings, self-interested *data owners (DOs)* meticulously evaluate various factors, including costs, potential gains, and utility, when deciding which FL *data consumer (DC)* to collaborate with. This necessity has spurred the development of FL incentive mechanisms (Tang 2024; Khan et al. 2020), designed to incentivize DO participation through strategic reward systems.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Auction-based federated learning (AFL) has emerged as a pivotal area within FL incentive mechanism research, primarily due to its ability to achieve both efficiency and fairness. AFL research revolves around three main areas based on the primary stakeholders involved: (Tang and Yu 2023d):

- *DC-Centric Approaches*: These methods assist DCs from two main perspectives: 1) Selecting appropriate DOs in monopolistic market scenarios, and 2) Strategically placing bids to acquire data from DOs in competitive market environments, with the aim of optimizing specific key performance indicators (KPIs) while adhering to budgetary constraints.
- *Auctioneer-Centric Approaches*: These approaches have two key focuses: 1) Internally, optimizing the DC-DO matching process, pricing mechanisms, and auction mechanism designs to achieve desired outcomes, such as maximizing social welfare or minimizing costs for the entire AFL ecosystem. 2) Externally, enhancing the competitiveness and attractiveness of the AFL ecosystem to retain existing participants, especially DOs who own valuable data resources while attracting new participants while optimizing its own profit.
- *DO-Centric Approaches*: These methodologies concentrate on determining the optimal allocation of local resources for federated learning tasks and customizing reserve prices to maximize the profitability of DOs.

In this paper, we focus on auctioneer-centric approaches in AFL. However, existing methods addressing auctioneer-centric challenges primarily concentrate on internal aspects, i.e., assisting the auctioneer in optimizing the matching between DCs and DOs, designing pricing mechanisms and auction mechanisms to achieve desired outcomes. To the best of our knowledge, no work has addressed external aspects, i.e., attracting more participants, especially DOs. Existing works for the internal aspect are based on the simplified assumption of a single, monopolistic AFL marketplace. However, this assumption does not hold true in practice.

Due to the self-interested nature of both the AFL marketplaces and DOs, data trading processes under AFL settings are a two-way interactive dynamic between them. On the one hand, the AFL marketplace, represented by the auctioneer, has the ability to selectively admit participants and determine the rewards for DOs. These strategies are designed

to maximize the auctioneer's profit as the coordinator and manager. Conversely, DOs also possess the autonomy to choose which AFL marketplace to join, given the availability of multiple options. Their decision is influenced by the potential payoffs they can obtain. This interdependent relationship creates a strategic interplay between the auctioneer's objectives of maximizing profit and retaining participants, and the DOs' goals of maximizing their individual payoffs. This two-way interactive process requires decision support support for the auctioneer that balance the interests of both AFL marketplaces and DOs, fostering a sustainable and mutually beneficial data trading environment within the AFL ecosystem (Tang and Yu 2024d).

To this end, we relax the monopolistic AFL marketplace assumption and assist the auctioneer in making decisions for external aspects by proposing the Auctioneer Revenue Allocation Strategy for AFL (ARAS-AFL). ARAS-AFL frames the objective of enhancing the competitiveness in the AFL ecosystem to attract DOs as a reputation management problem. Focused on individual-level decision-making for AFL marketplaces, i.e., achieving the dual objective of balancing the reputation management costs and profit of the AFL marketplace, ARAS-AFL designs a dynamic revenue allocation strategy under the reputation management framework for the AFL marketplace based on Lyapunov optimization. This strategy takes into consideration the marketplace's revenue and the changes in the number of DOs participating in the marketplace.

To the best of our knowledge, ARAS-AFL is the first revenue allocation decision support approach to help each AFL auctioneer manage its competitiveness. Extensive experiments on widely used benchmark datasets demonstrates its superiority over state-of-the-art approaches, outperforming the best baseline by 49.06%, 98.69%, 10.32%, and 4.77% in terms of total revenue, number of data owners, public reputation and accuracy of federated learning models, respectively.

Related Work

According to the target stakeholder, existing AFL methods could be roughly grouped into three main categories (Tang et al. 2024a): 1) DC-centric approaches, 2) DO-centric approaches, and 3) auctioneer-centric approaches.

DC-centric approaches

Based on auction mechanisms applied, methods in this category could be further divided into two main groups: 1) methods for reverse auction and 2) methods for forward auction.

Methods for reverse auction These methods (Jiao et al. 2020; Zhou et al. 2021; Yuan et al. 2021; Wu et al. 2023; Fan et al. 2020; Zhang, Wu, and Pan 2021; Zhang, Jingwen and Wu, Yuezhou and Pan, Rong 2022; Zhang, Wu, and Pan 2022; Zeng et al. 2020; Batool, Zhang, and Toews 2022; Batool, Zahra and Zhang, Kaiwen and Toews, Matthew 2023) select and recruit data owners (DOs) in a monopolistic marketplace, aiming to maximize the DC's key performance indicators under certain constraints based on the reverse auction mechanism. For instance, in (Jiao et al. 2020), the authors first group DOs based on Earth Mover's Distances

(EMD) (Zhao et al. 2018) to optimize social welfare. The DC then greedily selects DOs from each group, determining payments based on marginal virtual social welfare density. To further enhance social welfare, the authors incorporate a graph neural network to manage relationships among DOs and use deep reinforcement learning to determine the winning DOs and their payments. On the other hand, (Zhang, Wu, and Pan 2021) focuses on maximizing the utility of the DC by integrating the reputation mechanism and blockchain technology with the reverse auction mechanism. The proposed approach, RRAFL, assists the DC in selecting DOs and determining payments to selected DOs.

Methods for forward auction Methods (Tang and Yu 2023d,b,c; Tang et al. 2024b; Tang and Yu 2024a,b,c) in this category aim to assist the DC in bidding for DOs in the competitive AFL marketplace by designing bidding strategies to maximize key performance indicators while adhering to budget constraints. Specifically, in (Tang and Yu 2023d), the authors demonstrate that the optimal bidding strategy for DCs to maximize utility depends on utility estimation and winning probability estimation. Subsequently, the optimal bidding strategy design is translated into the utility estimation function design and winning probability function design. By formulating these two functions, the authors derive the final form of the bidding function. In contrast to (Tang and Yu 2023d), which solely considers the competitive relationship among DCs in the AFL marketplace, (Tang and Yu 2023b) also accounts for potential cooperative relationships among DCs. Based on this consideration, the authors propose MARL-AFL by framing the AFL marketplace from the perspective of a multi-agent bidding system.

DO-centric approaches methods

These methods (Thi Le et al. 2021; Lu et al. 2023; Le et al. 2020; Zeng et al. 2020) are designed to assist the DOs in determining the optimal allocation of their data resources and formulating their asking profiles, with the aim of maximizing key performance indicators for the DOs. For instance, in (Thi Le et al. 2021), upon receiving FL training task profiles from the DC, which include the maximum tolerable time for FL training, the proposed method helps each DO optimize their asking profiles. On the other hand, some methods, such as (Tang and Yu 2023a), focus on assisting DOs in determining the number of FL training tasks to accept, the number of FL training tasks to subdelegate, and how to price these tasks. To achieve this multifaceted objective, the authors propose PAS-AFL by leveraging Lyapunov optimization and taking into account the DO's current reputation, pending FL tasks, willingness to engage in FL model training, and trust relationships with other DOs.

Auctioneer-centric approaches

Auctioneer-centric approaches (Seo, Niyato, and Elmroth 2021; Seo, Eunil and Niyato, Dusit and Elmroth, Erik 2022; Xu et al. 2023; Roy et al. 2021; Mai et al. 2022; Wang et al. 2023) are generally designed to assist the auctioneer in making optimal DC-DO matching decisions to maximize key performance indicators from the perspective of the entire

AFL marketplace. These approaches are based on mechanisms such as double auctions, combinatorial auctions, and reverse auctions. In (Xu et al. 2023), the proposed method aims to maximize social welfare while protecting the auctioneer’s utility. The approach involves two main stages: 1) combinatorial auction stage, where the platform selects winners whose total utility, combined with the platform’s utility, is greater than zero, and 2) bargaining stage, where winners are classified into two categories with different payment methods after completing the training model. The goal is to ensure that the auctioneer’s utility remains positive.

Our work falls under the auctioneer-centric approaches category. However, in contrast to existing methods in this category, which are generally based on the simplified assumption of a single AFL marketplace, the proposed ARAS-AFL method relaxes this unrealistic assumption. Instead, it frames the decision support for the auctioneer from the perspective of helping it manage its reputation in the competitive AFL ecosystem and attract more participants to obtain higher revenue.

Preliminaries

In this section, we introduce some background information related to the proposed method ARAS-AFL.

AFL Ecosystem

We delve into a dynamic and competitive AFL ecosystem wherein $|\mathcal{N}|$ marketplaces (referred to as the marketplace set \mathcal{N}) actively vie for the attention of DOs, enticing them to join their respective platforms at any given moment. Within this ecosystem, AFL marketplaces harness the data resources contributed by DOs to attract DCs in need of FL model training services. These DCs release their training tasks and offer monetary compensation for the services provided. Nevertheless, DOs retain the autonomy to assess and choose the AFL marketplace they prefer to join, enabling them to engage in FL training tasks facilitated by their selected platform. Due to privacy concerns, AFL marketplaces may possess limited access to comprehensive attribute information about the DOs. Simultaneously, DOs who have not previously participated in FL training tasks hosted by a particular AFL marketplace may lack insight into the potential rewards and incentives offered by that platform. It is crucial to acknowledge that AFL marketplaces, acting as intermediaries, aim to generate revenue by overseeing FL training tasks. They accomplish this by strategically distributing the payoffs received from DCs among themselves and the participating DOs, driven by their self-interested motive to maximize profitability.

In this intricate landscape, each AFL marketplace $i \in \mathcal{N}$ must adopt an effective FL revenue allocation strategy, denoted as α_i , to enhance its attractiveness and recruit a substantial pool of DOs to its platform. The revenue allocation strategy plays a pivotal role in enhancing the competitiveness of the AFL marketplace, helping it attract more DOs, and maximizing long-term revenue. This approach benefits the marketplace while fostering a sustainable environment for all stakeholders in the AFL ecosystem.

We frame AFL marketplace attractiveness from the perspective of reputation. We then orient the objectives of AFL marketplaces toward devising effective revenue allocation strategies under the reputation management framework.

The Data Trading Process under ARAS-AFL Framework

The data trading process among DCs and DOs coordinated and managed by the AFL marketplace under the proposed ARAS-AFL in a competitive AFL ecosystem operates in the following six steps within a given time slot t , which can represent a day, a week, or any other suitable time frame depending on the specific application scenario.

1. DCs select an AFL marketplace to release their corresponding FL training tasks based on their observation of the number of DOs currently participating in each marketplace. As many new DCs may have limited information about the available AFL marketplaces, they intuitively gravitate towards marketplaces with a larger number of DOs, as this implies a higher probability of completing the FL training tasks with high quality and in a timely manner.
2. DOs who wish to participate in FL choose an AFL marketplace based on its reputation. The intuition behind this decision is that an AFL marketplace with a good reputation is more likely to provide higher payoffs to participating DOs. Each AFL marketplace has a perceived reputation based on satisfaction ratings provided by past DOs, reflecting their experiences with the respective marketplace. This reputation information, jointly maintained by the DOs, serves as a valuable guide for future DOs in choosing which AFL marketplace to join.
3. The AFL marketplaces allocate the FL model training tasks released by DCs to available DOs on the marketplace, considering factors such as their availability and capabilities, followed by DOs training the FL models following the traditional FL training algorithm like FedAvg (McMahan et al. 2017). After obtaining the trained FL models, the AFL marketplace transfers the trained models to the corresponding DCs.
4. Each AFL marketplace determines the revenue-sharing ratio based on its revenue allocation strategy between the marketplace itself and the participating DOs.
5. After the DOs complete the assigned FL training tasks, they evaluate the AFL marketplace based on their earned income. At the end of the time slot t , the reputation of each marketplace is updated based on the reputation evaluation ratings provided by the participating DOs, as detailed in Section .
6. Finally, each AFL marketplace dynamically adjusts its revenue allocation strategy based on the changes in the DO population on its platform and its own income generated during the time slot t . The details of the FL training task allocation strategy will be discussed in Section .

The Proposed ARAS-AFL Approach

In this section, we delineate the primary mechanisms governing the data trading process within the proposed

ARAS-AFL. This encompasses the procedures by which DCs issue their FL model training tasks to AFL marketplaces, the methods by which DOs assess the reputation of these marketplaces, and the strategies employed by AFL marketplaces to enhance their appeal to DOs while ensuring profitability. Central to this discussion is the development of effective revenue allocation strategies within the reputation management framework, which serves to optimize the attractiveness of AFL marketplaces to DOs while safeguarding their financial viability.

DCs Release FL Training Tasks

For DCs, selecting a high-quality AFL marketplace can significantly improve their chances of obtaining high quality FL models in a timely manner. However, in the absence of crucial information regarding AFL marketplaces, such as FL model training capacity and the quality of DOs, DCs may tend to base their decisions on the marketplace's DO availability. Typically, a higher number of DOs on a marketplace correlates with a greater likelihood of completing FL model training tasks swiftly, albeit disregarding other factors. Consequently, the probability of a DC releasing an FL model training task on AFL marketplace i during time period t can be expressed as:

$$p_i^{release}(t) = \frac{M'_i(t)}{\sum_{n \in \mathcal{N}} M'_n(t)}, \quad (1)$$

where $M'_i(t) = M_i(t) + M_{base}$ denote the initial number of DOs available on AFL marketplace i at the beginning of time period t , where $M_i(t)$ represents the number of DOs on marketplace i during time period t . Here, M_{base} denotes a small positive value, signifying the loyal core or intrinsic DOs of marketplace i . This inclusion ensures that even weaker AFL marketplaces are afforded an opportunity to rejuvenate and expand.

AFL Marketplace Reputation Evaluation

The reputation of an AFL marketplace encompasses two facets: firstly, the private reputation value logged by each DO, and secondly, the public reputation value, both derived from the Beta Reputation System (BRS) (Josang and Ismail 2002). In this context, we assume DOs provide honest assessments of their satisfaction with each AFL marketplace, with evaluations conducted after receiving payment for completing FL training tasks on the respective marketplace.

For AFL marketplace i , the satisfaction level of DO j regarding i , based on its earnings $profit_i^j$ from the FL model training task, falls into two categories: satisfaction, indicated by $profit_i^j \geq \bar{profit}_i^j$, where \bar{profit}_i^j denotes DO j 's average earnings across various AFL marketplaces in the past; and dissatisfaction, denoted by $profit_i^j < \bar{profit}_i^j$. We employ the variables sc_i^j and uc_i^j to tally the occurrences of DO j expressing satisfaction and dissatisfaction with AFL marketplace i , respectively.

Subsequently, we calculate the reputation value r_i^j for DO j concerning AFL marketplace i using the BRS as follows:

$$r_i^j = \mathbb{E}[Beta(sc_i^j + 1, uc_i^j + 1)] = \frac{sc_i^j + 1}{sc_i^j + uc_i^j + 2}. \quad (2)$$

Regarding the public reputation value of each AFL marketplace, the proposed ARAS-AFL upholds a comprehensive reputation table. This table consolidates satisfaction evaluations provided by DOs engaged in FL model training tasks within the respective marketplace. It serves as a public attribute of the AFL marketplace, accessible to all DOs. The public reputation value r_i^* is derived through the following calculation:

$$\begin{aligned} r_i^* &= \mathbb{E}[Beta(\sum_{j \in \mathcal{K}} (sc_i^j + 1), \sum_{j \in \mathcal{K}} (uc_i^j + 1))] \\ &= \frac{\sum_{j \in \mathcal{K}} (sc_i^j + 1)}{\sum_{j \in \mathcal{K}} (sc_i^j + uc_i^j + 2)}, \end{aligned} \quad (3)$$

where \mathcal{K} is the DO set.

By integrating the public reputation r_i^* and the private reputation r_i^j logged by DO j , a holistic reputation value $r_i^{j'}$ for AFL marketplace i can be derived. This comprehensive measure encapsulates both the collective perception and the individual experiences of DO j within the marketplace. The formulation is as follows:

$$r_i^{j'} = (1 - \gamma)r_i^* + \gamma r_i^j, \quad (4)$$

where γ signifies the confidence level of DO j in assessing the reputation of AFL marketplace i . If DO j has not previously engaged in any FL model training tasks on AFL marketplace i , its decision-making relies solely on the public reputation of the marketplace. In this case, setting γ to 0 results in $r_i^{j'} = r_i^*$. Conversely, if DO j places full trust in its own evaluation of AFL marketplace i , γ is set to 1, leading to $r_i^{j'} = r_i^j$. This approach empowers both new and existing DOs with a flexible means to select AFL marketplaces based on their individual preferences and experiences.

Reputation Management Model

AFL marketplaces often allocate a larger portion of their profits (denoted as α) to DOs in an effort to bolster their reputation and, consequently, attract a greater number of DOs. However, despite increasing the proportion of rewards for FL model training tasks, AFL marketplace i may not consistently draw more DOs due to uncertainties regarding competitors' strategies. This scenario can lead to regret for the marketplace.

To address this, we introduce a virtual queue $q_i(t)$ for each AFL marketplace i to quantify its regret. This virtual queue encapsulates the dynamics of marketplace i 's regret and can be formulated as follows:

$$q_i(t + 1) = \max[q_i(t) + \Delta p_i(t) - \Delta M_i(t), 0]. \quad (5)$$

$\Delta M_i(t)$ denotes the change in the number of DOs on AFL marketplace i from $(t - 1)$ -th time period to t -th time period. $\Delta p_i(t)$ denotes the amount of variation in the payoffs paid to DOs by platform i during $(t - 1)$ -th time period and t -th time period and is calculated as

$$\Delta p_i(t) = (\alpha_i(t)P_i(t) - \alpha_i(t - 1)P_i(t - 1))\mathbb{1}_{\Delta M_i(t) < 0}, \quad (6)$$

where $P_i(t)$ represents the aggregate payoff offered by DCs to AFL marketplace i for completed FL model training tasks

at time t . The function $\mathbb{I}_{(\text{condition})}$ serves as an indicator, taking the value 1 if the condition holds true, and 0 otherwise.

The virtual queue $q_i(t)$ in Eq. (5) serves as a cumulative measure of regret for AFL marketplace i , reflecting its profit-sharing strategy’s impact on DO attraction. Its dynamics evolve as follows:

- The regret queue $q_i(t)$ for AFL marketplace i grows under the condition that $\Delta M_i(t) < 0$ upon its update at time t . In this scenario, $q_i(t)$ expands by $\alpha_i(t)P_i(t) - \alpha_i(t-1)P_i(t-1) - M_i(t)$. This mechanism ensures that $q_i(t)$, representing regret, continues to escalate if marketplace i fails to draw more DOs despite increased costs incurred in time period t .
- Conversely, $q_i(t)$ diminishes by $\Delta M_i(t)$ when $\Delta M_i(t) \geq 0$.

The queuing dynamics in Eq. (5) elucidate that when AFL marketplace i allocates a large proportion of revenue to DOs ($\alpha_i(t)$) yet transacts few FL model training tasks, i.e., attracting a limited number of new DOs, the marketplace’s regret queue grows.

Moving forward, we address the two primary objectives of the proposed ARAS-AFL. Firstly, it aims to assist each AFL marketplace i in dynamically optimizing its revenue-sharing strategy $\alpha_i(t)$. This optimization ensures that all AFL marketplaces maintain a relatively stable regret value within the competitive ecosystem. By mitigating fluctuations in regret, the ecosystem remains resilient, averting undue volatility. To achieve this, we redefine the time-coupling constraint as a new objective, imposing an upper limit on queue growth to prevent unbounded expansion. Subsequently, we employ the Lyapunov optimization technique (Neely 2010) to fulfill this objective:

$$\mathcal{L}(\Theta(t)) = \frac{1}{2} \sum_{i \in \mathcal{N}} q_i^2(t) \geq 0. \quad (7)$$

In assessing the anticipated growth of the Lyapunov function $\mathcal{L}(\Theta(t))$, we define a Lyapunov drift as:

$$\Delta(\Theta(t)) = \mathbb{E}[\mathcal{L}(\Theta(t+1)) - \mathcal{L}(\Theta(t)) | \Theta(t)]. \quad (8)$$

It describes the anticipated fluctuation in the function over a single time slot, considering the current state, with expectations drawn from queue evolution statistics. According to the Lyapunov drift theorem, the minimization of $\Delta(\Theta(t))$ at each time t ensures the stabilization of all queues.

Optimal Decision Strategy

In addition to the goal of maintaining stability in regret among AFL marketplaces, the proposed ARAS-AFL also targets the enhancement of revenue for all involved marketplaces, ensuring sustainable long-term growth. The payoff $U_i(t)$ for AFL marketplace i at time step t is defined as:

$$U_i(t) = (1 - \alpha_i(t))P_i(t). \quad (9)$$

Then, we can formulate the total revenue $U(t)$ generated by all AFL marketplaces at time step t as follows:

$$U(t) = \sum_{i \in \mathcal{N}} [(1 - \alpha_i(t))P_i(t)]. \quad (10)$$

Taking into account both the regret and profit dynamics of AFL marketplaces, devising revenue allocation strategies that cater to all involved DOs transforms into a dual objective optimization challenge. Consequently, we formalize the task of selecting a revenue allocation strategy for each AFL marketplace to maximize $\{\textit{profit-drift}\}$ objective function:

$$\begin{aligned} \max \frac{1}{T} \sum_{t=0}^{T-1} \{ \rho \mathbb{E}[U(t) | \Theta(t)] - \Delta(\Theta(t)) \} \\ \text{s.t.} \quad 0 < \alpha_i(t) (i \in \mathcal{N}) < b \leq 1. \end{aligned} \quad (11)$$

The constraint $0 < \alpha_i(t) (i \in \mathcal{N}) < b \leq 1$ in Eq. (11) stipulates that, regardless of the circumstances, once a DO completes an FL model training task, the AFL marketplace must compensate the DO with a certain amount of remuneration. This constraint ensures that the payment ratio of the remuneration never exceeds the total reward for the FL model training task. The parameter ρ serves as a weighting factor, allowing AFL marketplaces to express their relative preferences between regret and profit.

By substituting Eqs. (7), (8), and (10) into Eq. (11), we obtain:

$$\begin{aligned} \max \frac{1}{T} \sum_{t=0}^{T-1} \sum_{i \in \mathcal{N}} \{ -\frac{1}{2} P_i^2(t) \alpha_i^2(t) + [P_i(t)P_i(t-1)\alpha_i(t-1) \\ + P_i(t)\Delta M_i(t) - \rho P_i(t) - q_i(t)P_i(t)]\alpha_i(t) \\ + q_i(t)\alpha_i(t-1)P_i(t-1) + \rho P_i(t) + q_i(t)\Delta M_i(t) \\ - \frac{1}{2} [\alpha_i(t-1)P_i(t-1) + \Delta M_i(t)]^2 \}. \end{aligned} \quad (12)$$

Directing our attention to the components reliant on the decision variable $\alpha_i(t)$, we reinterpret Eq. (12) as follows:

$$\begin{aligned} \max \frac{1}{T} \sum_{t=0}^{T-1} \sum_{i \in \mathcal{N}} \{ -\frac{1}{2} P_i^2(t) \alpha_i^2(t) + [P_i(t)P_i(t-1)\alpha_i(t-1) \\ + P_i(t)\Delta M_i(t) - \rho P_i(t) - q_i(t)P_i(t)]\alpha_i(t) \}. \end{aligned} \quad (13)$$

To make it simple, we set F as:

$$\begin{aligned} F = -\frac{1}{2} P_i^2(t) \alpha_i^2(t) + [P_i(t)P_i(t-1)\alpha_i(t-1) \\ + P_i(t)\Delta M_i(t) - \rho P_i(t) - q_i(t)P_i(t)]\alpha_i(t). \end{aligned} \quad (14)$$

By computing the first-order partial derivative of F with respect to $\alpha_i(t)$ and setting it equal to zero, denoted as $\frac{\partial F}{\partial \alpha_i(t)} = 0$, we derive the optimal revenue allocation strategy $\alpha_i^*(t)$ for AFL marketplace i at time step t as follows:

$$\alpha_i^*(t) = \frac{P_i(t-1)\alpha_i(t-1) + \Delta M_i(t) - \rho - q_i(t)}{P_i(t)}. \quad (15)$$

Experimental Evaluation

Experiment Setup

Datasets Our experiments are based on six commonly adopted datasets in FL: MNIST¹, CIFAR-10², Fashion-MNIST (FMNIST) (Xiao, Rasul, and Vollgraf 2017),

¹<http://yann.lecun.com/exdb/mnist/>

²<https://www.cs.toronto.edu/~kriz/cifar.html>

EMNIST-digits (EMNISTD) / letters (EMNISTL) (Cohen et al. 2017), Kuzushiji-MNIST (KMNIST) (Clanuwat et al. 2018). The FL models used are identical to those employed in (Tang and Yu 2023d).

Experimental scenarios As there is a lack of real-time FL training tasks releasing data shared across multiple AFL marketplaces, we consider the following two scenarios:

- **FL training task overrelease market.** In this scenario, the number of FL training tasks released by DCs in the AFL ecosystem exceeds the number of available DOs. We set up the AFL ecosystem with a total of 600 DOs, and the DCs on the ecosystem released 2,000 FL training tasks simultaneously. All the AFL marketplaces in the ecosystem competed for the 600 DOs.
- **FL training task underrelease market.** In this scenario, at each time step t , the number of available DOs in the ecosystem was greater than the number of FL training tasks released by DCs. Specifically, we set the number of available DOs to 1,000 in the AFL ecosystem, and the DCs released 80 FL training tasks.

Under either of these two experimental scenarios, we generated 400,000 FL training tasks with payments following a normal distribution with a mean of 80 and a standard deviation of 10, i.e., $\mathcal{N}(80, 10^2)$. We released the FL training tasks in batches over time. During our experiments, we set the entire process of each FL training task, from release to completion and subsequent evaluation by DOs on the corresponding AFL marketplace, within a time period t .

We set the confidence degree γ in Equation (3) to 0.5 for each DO and the weighting factor ρ in Equation (11) to 0.1. Additionally, we set $0.5 > \alpha_i(t) \geq 0.3$ to ensure that the basic cost of DOs and the basic operating cost of the AFL marketplace were covered.

Comparison methods To our best knowledge, no existing decision support method has been designed specifically for auctioneers in the AFL ecosystem. Therefore, we selected the following two commonly used strategies as benchmarks:

- **Random revenue allocation (Rand):** At each time step, each auctioneer randomly generates a revenue allocation ratio $\alpha_i(t)$ for the DOs.
- **Greedy revenue allocation (Greedy).** To ensure that the AFL marketplace retains as much revenue as possible from each completed FL training task released by DCs, the auctioneer adopts the smallest possible task commission allocation ratio $\alpha_i(t)$ for DOs.

Evaluation metrics We evaluate the proposed method and the benchmark strategies based on the following metrics:

- **Averaged total revenue (Rev).** Since one of the objectives of AFL marketplaces is to maximize revenue, we adopt the averaged total revenue over time for each AFL marketplace as an evaluation metric.
- **Averaged number of DOs (#DOs).** AFL marketplaces aim to attract and retain more participants, especially DOs, by allocating revenue to them. Therefore, we use the averaged number of recruited DOs as a metric to evaluate the effectiveness of the proposed method.

- **Public reputation (PR).** This metric evaluates the effectiveness of all the compared methods by recording the averaged reputation of each AFL marketplace.
- **Accuracy of the FL models (Acc).** To ensure the long-term sustainability and fairness of the AFL ecosystem, any integrated method should balance the interests of all stakeholders, including DCs. Consequently, we evaluate the effectiveness of the auctioneer decision support strategies by assessing the test accuracy of the FL models obtained by DCs within the AFL ecosystem. This metric serves as an indicator of whether the proposed methods foster an environment that supports DCs in achieving high-quality FL model performance, thereby aligning with their interests and promoting the overall development of the ecosystem.

Results and Discussion

In this section, we conduct an in-depth analysis of the performance of the comparison approaches concerning four crucial evaluation metrics: total revenue, number of DOs, public reputation, and FL model accuracy. Each experiment setting was rigorously executed five times to ensure statistical robustness. The averaged results are meticulously presented in Tables 1 and 2, complemented by visually illustrative representations in Figures 1 and 2.

Table 1 unveils the results of the comparison methods under the overrelease scenario. It is evident that across all six datasets, our proposed method, ARAS-AFL, consistently surpasses all baseline methods concerning all four evaluation metrics. Notably, when compared to the best performance achieved by the baselines, ARAS-AFL demonstrates remarkable improvements of 64.80%, 150.71%, 13.63%, and 4.34% in terms of the averaged total revenue, the averaged number of DOs, the public reputation, and the FL model accuracy, respectively.

Figure 1 vividly depicts the corresponding results during the 200-epoch experiment, at which point the number of DOs on each AFL marketplace has nearly stabilized. The results align seamlessly with the findings presented in Table 1, substantiating the efficacy of our proposed approach.

Furthermore, the comparative results under the under-release scenario are comprehensively detailed in Table 2 and Figure 2. It is noteworthy that under this scenario, ARAS-AFL demonstrates substantial improvements of 33.32%, 46.67%, 69.98%, and 5.20% in the averaged total revenue, the averaged number of DOs, the public reputation, and the FL model accuracy, respectively.

It is worth highlighting that, as depicted in the last subfigures of Figures 1 and 2, the public reputation obtained by the Rand method experiences less fluctuation during the experiments. Additionally, the public reputation achieved by the proposed method, ARAS-AFL, exhibits a decreasing trend over the course of the experiments. This phenomenon could be attributed to the fact that as the AFL marketplace attracts more DOs, there may not be enough high-payoff FL training tasks from DCs, resulting in some DOs being able to obtain FL training tasks with lower payoffs. Consequently, this may lead to an unsatisfactory evaluation of the corresponding AFL marketplace by the DOs. Despite this, the proposed

Dataset	Averaged total revenue			Averaged number of DOs			Public reputation			Accuracy of the FL models (%)		
	ARAS-AFL	Rand	Greedy	ARAS-AFL	Rand	Greedy	ARAS-AFL	Rand	Greedy	ARAS-AFL	Rand	Greedy
MNIST	9.41	5.73	1.40	414	158	31	0.482	0.424	0.281	84.82	81.45	83.16
CIFAR	8.31	5.05	1.31	416	171	46	0.485	0.427	0.282	42.37	40.03	40.84
FMNIST	9.29	5.65	1.40	415	160	31	0.482	0.424	0.281	75.10	73.34	73.58
EMNISTD	9.13	5.57	1.39	413	159	31	0.482	0.425	0.281	82.95	78.21	77.68
EMNISTL	9.39	5.57	2.24	369	156	41	0.517	0.445	0.297	74.56	67.44	68.93
KMNIST	9.28	5.69	1.88	373	154	39	0.497	0.446	0.281	73.14	70.31	68.93

Table 1: Performance comparison under the overrelease scenario.

Dataset	Averaged total revenue			Averaged number of DOs			Public reputation			Accuracy of the FL models (%)		
	ARAS-AFL	Rand	Greedy	ARAS-AFL	Rand	Greedy	ARAS-AFL	Rand	Greedy	ARAS-AFL	Rand	Greedy
MNIST	2.22	1.55	1.64	424	286	219	0.491	0.462	0.468	85.21	80.93	83.79
CIFAR	2.10	1.50	1.53	434	304	238	0.518	0.489	0.492	43.18	40.79	41.27
FMNIST	2.20	1.54	1.62	421	287	219	0.490	0.463	0.468	75.76	72.89	74.04
EMNISTD	2.18	1.53	1.60	421	283	217	0.490	0.460	0.466	84.52	79.72	79.63
EMNISTL	2.40	1.68	2.07	391	282	246	0.459	0.409	0.414	76.11	68.38	69.52
KMNIST	2.42	1.74	1.72	407	263	227	0.467	0.419	0.421	73.17	68.36	67.88

Table 2: Performance comparison under the underrelease scenario.




Figure 1: Comparison of different method under the overrelease scenario.




Figure 2: Comparison of different method under the underrelease scenario.

ARAS-AFL still enables the AFL marketplace to maintain higher revenues and attract more DOs, as evidenced by the comprehensive results presented in Tables 1 and 2, as well as Figures 1 and 2.

Conclusions

This paper frames the issues in the AFL ecosystem from the perspective of auctioneers and investigates decision support strategies for them. We propose a novel method, termed ARAS-AFL, to assist auctioneers in attracting more Os within a competitive AFL environment while ensuring their profitability. ARAS-AFL approaches the challenges of attractiveness and competitiveness through the lens of reputation. By leveraging Lyapunov optimization techniques, ARAS-AFL enables each individual AFL marketplace to achieve the dual objective of balancing reputation management costs and profitability by dynamically adjusting its revenue allocation strategy within the reputation management

framework. The decision-making process considers both the marketplace’s revenue and the changes in the number of participants, allowing for informed and adaptive strategies. To the best of our knowledge, ARAS-AFL is the first approach designed specifically to assist auctioneers in making decisions to address external competitiveness challenges while ensuring the long-term sustainability of their marketplace.

Acknowledgments

This research is supported by the RIE2025 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) (Award I2301E0026), administered by A*STAR, as well as supported by Alibaba Group and NTU Singapore through Alibaba-NTU Global e-Sustainability CorpLab (ANGEL); and National Research Foundation, Singapore and DSO National Laboratories under the AI Singapore Programme (AISG Award No: AISG2-RP-2020-019).

References

- Batool, Z.; Zhang, K.; and Toews, M. 2022. Fl-mab: client selection and monetization for blockchain-based federated learning. In *Proceedings of the 37th ACM/SIGAPP symposium on applied computing*, 299–307.
- Batool, Zahra and Zhang, Kaiwen and Toews, Matthew. 2023. Block-RACS: Towards Reputation-Aware Client Selection and Monetization Mechanism for Federated Learning. *ACM SIGAPP Applied Computing Review*, 23(3): 49–66.
- Clanuwat, T.; Bober-Irizar, M.; Kitamoto, A.; Lamb, A.; Yamamoto, K.; and Ha, D. 2018. Deep Learning For Classical Japanese Literature. *arXiv preprint*, 1812.01718.
- Cohen, G.; Afshar, S.; Tapson, J.; and Van Schaik, A. 2017. EMNIST: Extending MNIST to handwritten letters. In *IJCNN*, 2921–2926.
- Fan, S.; Zhang, H.; Zeng, Y.; and Cai, W. 2020. Hybrid blockchain-based resource trading system for federated learning in edge computing. *IEEE Internet of Things Journal*, 8(4): 2252–2264.
- Guo, S.; Zhang, T.; Xu, G.; Yu, H.; Xiang, T.; and Liu, Y. 2021. Byzantine-resilient decentralized stochastic gradient descent. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(6): 4096–4106.
- Hu, M.; Cao, Y.; Li, A.; Li, Z.; Liu, C.; Li, T.; Chen, M.; and Liu, Y. 2024. FedMut: Generalized Federated Learning via Stochastic Mutation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 12528–12537.
- Jiao, Y.; Wang, P.; Niyato, D.; Lin, B.; and Kim, D. I. 2020. Toward an Automated Auction Framework For Wireless Federated Learning services market. *IEEE Transactions on Mobile Computing*, 20(10): 3034–3048.
- Josang, A.; and Ismail, R. 2002. The beta reputation system. In *Proceedings of the 15th bled electronic commerce conference*, volume 5, 2502–2511. Citeseer.
- Khan, L. U.; Pandey, S. R.; Tran, N. H.; Saad, W.; Han, Z.; Nguyen, M. N. H.; and Hong, C. S. 2020. Federated Learning for Edge Networks: Resource Optimization and Incentive Mechanism. *IEEE Communications Magazine*, 58(10): 88–93.
- Le, T. H. T.; Tran, N. H.; Tun, Y. K.; Han, Z.; and Hong, C. S. 2020. Auction based Incentive Design for Efficient Federated Learning in Cellular Wireless Networks. In *WCNC*, 1–6.
- Li, A.; Zhang, L.; Wang, J.; Han, F.; and Li, X.-Y. 2021. Privacy-preserving efficient federated-learning model debugging. *IEEE Transactions on Parallel and Distributed Systems*, 33(10): 2291–2303.
- Lu, R.; Zhang, W.; Wang, Y.; Li, Q.; Zhong, X.; Yang, H.; and Wang, D. 2023. Auction-based cluster federated learning in mobile edge computing systems. *IEEE Transactions on Parallel and Distributed Systems*, 34(4): 1145–1158.
- Mai, T.; Yao, H.; Xu, J.; Zhang, N.; Liu, Q.; and Guo, S. 2022. Automatic double-auction mechanism for federated learning service market in Internet of Things. *IEEE Transactions on Network Science and Engineering*, 9(5): 3123–3135.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS'17)*, 1273–1282.
- Neely, M. J. 2010. Stochastic network optimization with application to communication and queueing systems. *Synthesis Lectures on Communication Networks*, 3(1): 1–211.
- Roy, P.; Sarker, S.; Razzaque, M. A.; Mamun-or Rashid, M.; Hassan, M. M.; and Fortino, G. 2021. Distributed Task Allocation in Mobile Device Cloud Exploiting Federated Learning and Subjective Logic. *Journal of Systems Architecture*, 113(2): doi:10.1016/j.sysarc.2020.101972.
- Seo, E.; Niyato, D.; and Elmroth, E. 2021. Auction-based Federated Learning using Software-defined Networking for resource efficiency. In *2021 17th International Conference on Network and Service Management (CNSM)*, 42–48. IEEE.
- Seo, Eunil and Niyato, Dusit and Elmroth, Erik. 2022. Resource-efficient federated learning with non-iid data: An auction theoretic approach. *IEEE Internet of Things Journal*, 9(24): 25506–25524.
- Sun, H.; Tang, X.; Yang, C.; Yu, Z.; Wang, X.; Ding, Q.; Li, Z.; and Yu, H. 2024. HiFi-Gas: Hierarchical Federated Learning Incentive Mechanism Enhanced Gas Usage Estimation. In *the 36th Annual Conference on Innovative Applications of Artificial Intelligence (IAAI-24)*, 22824–22832.
- Tang, X. 2024. Stakeholder-oriented decision support for auction-based federated learning. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI'24)*, 8514–8515.
- Tang, X.; and Yu, H. 2023a. Agent-oriented Joint Decision Support for Data Owners in Auction-based Federated Learning. In *Proceedings of the 2024 IEEE International Conference on Multimedia and Expo (ICME'24)*.
- Tang, X.; and Yu, H. 2023b. Competitive-Cooperative Multi-Agent Reinforcement Learning for Auction-based Federated Learning. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI'23)*, 4262–4270.
- Tang, X.; and Yu, H. 2023c. Multi-Session Budget Optimization for Forward Auction-based Federated Learning. *arXiv preprint arXiv:2311.12548*.
- Tang, X.; and Yu, H. 2023d. 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)*, 330–335.
- Tang, X.; and Yu, H. 2024a. A Cost-Aware Utility-Maximizing Bidding Strategy for Auction-Based Federated Learning. *IEEE Transactions on Neural Networks and Learning Systems*.
- Tang, X.; and Yu, H. 2024b. Efficient Large-Scale Personalizable Bidding for Multi-Agent Auction-Based Federated Learning. *IEEE Internet of Things Journal*.
- Tang, X.; and Yu, H. 2024c. Fairness-Aware Reverse Auction-Based Federated Learning. *IEEE Internet of Things Journal*.

- Tang, X.; and Yu, H. 2024d. Towards Trustworthy AI-Empowered Real-Time Bidding for Online Advertisement Auctioning. *ACM Computing Surveys*.
- Tang, X.; Yu, H.; Li, X.; and Kraus, S. 2024a. Intelligent Agents for Auction-based Federated Learning: A Survey. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI'24)*, 8253–8261.
- Tang, X.; Yu, H.; Li, Z.; and Li, X. 2024b. A Bias-Free Revenue-Maximizing Bidding Strategy for Data Consumers in Auction-based Federated Learning. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI'24)*, 4991–4999.
- Tang, X.; Yu, H.; Tang, R.; Ren, C.; Li, A.; and Li, X. 2024c. Dual Calibration-based Personalised Federated Learning. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI'24)*, 4982–4990.
- Thi Le, T. H.; Tran, N. H.; Tun, Y. K.; Nguyen, M. N. H.; Pandey, S. R.; Han, Z.; and Hong, C. S. 2021. An Incentive Mechanism for Federated Learning in Wireless Cellular Networks: An Auction Approach. *IEEE Transactions on Wireless Communications*, 20(8): 4874–4887.
- Wang, F.; Jiao, L.; Zhu, K.; Lin, X.; and Li, L. 2023. Toward Sustainable AI: Federated Learning Demand Response in Cloud-Edge Systems via Auctions. In *IEEE INFOCOM 2023-IEEE Conference on Computer Communications*, 1–10. IEEE.
- Wu, L.; Guo, S.; Hong, Z.; Liu, Y.; Xu, W.; and Zhan, Y. 2023. Long-Term Adaptive VCG Auction Mechanism for Sustainable Federated Learning With Periodical Client Shifting. *TMC*.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-MNIST: A Novel Image Dataset for Benchmarking Machine Learning Algorithms. *arXiv preprint*, 1708.07747.
- Xu, B.; Zuo, L.; Jin, J.; Han, L.; and Hu, K. 2023. CAB: a combinatorial-auction-and-bargaining-based federated learning incentive mechanism. *WWW*, 1–22.
- Yang, Q.; Liu, Y.; Chen, T.; and Tong, Y. 2019. Federated Machine Learning: Concept and Applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2): 12:1–12:19.
- Yuan, Y.; Jiao, L.; Zhu, K.; and Zhang, L. 2021. Incentivizing federated learning under long-term energy constraint via online randomized auctions. *IEEE Transactions on Wireless Communications*, 21(7): 5129–5144.
- Zeng, R.; Zhang, S.; Wang, J.; and Chu, X. 2020. FMore: An Incentive Scheme of Multi-dimensional Auction for Federated Learning in MEC. In *ICDCS*, 278–288.
- Zhang, J.; Wu, Y.; and Pan, R. 2021. Incentive Mechanism for Horizontal Federated Learning Based on Reputation and Reverse Auction. In *WWW*, 947–956.
- Zhang, J.; Wu, Y.; and Pan, R. 2022. Online Auction-Based Incentive Mechanism Design for Horizontal Federated Learning with Budget Constraint. *arXiv preprint*, 2201.09047.
- Zhang, Jingwen and Wu, Yuezhou and Pan, Rong. 2022. Auction-Based Ex-Post-Payment Incentive Mechanism Design for Horizontal Federated Learning with Reputation and Contribution Measurement. *arXiv preprint arXiv:2201.02410*.
- Zhao, Y.; et al. 2018. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*.
- Zhou, R.; Pang, J.; Wang, Z.; Lui, J. C.; and Li, Z. 2021. A truthful procurement auction for incentivizing heterogeneous clients in federated learning. In *ICDCS*, 183–193.