Intelligent Agents for Auction-based Federated Learning: A Survey

Xiaoli Tang[1] , **Han Yu**[1] , **Xiaoxiao Li**[2] and **Sarit Kraus**[1,3]

[1]College of Computing and Data Science, Nanyang Technological University, Singapore
[2]Department of Electrical and Computer Engineering, The University of British Columbia, Canada
[3]Department of Computer Science, Bar-Ilan University, Israel
{xiaoli001, han.yu}@ntu.edu.sg, xiaoxiao.li@ece.ubc.ca, sarit@cs.biu.ac.il

Abstract

Auction-based federated learning (AFL) is an important emerging category of FL incentive mechanism design, due to its ability to fairly and efficiently motivate high-quality data owners to join data consumers' (i.e., servers') FL training tasks. To enhance the efficiency in AFL decision support for stakeholders (i.e., data consumers, data owners, and the auctioneer), intelligent agent-based techniques have emerged. However, due to the highly interdisciplinary nature of this field and the lack of a comprehensive survey providing an accessible perspective, it is a challenge for researchers to enter and contribute to this field. This paper bridges this important gap by providing a first-of-its-kind survey on the Intelligent Agents for AFL (IA-AFL) literature. We propose a unique multi-tiered taxonomy that organises existing IA-AFL works according to 1) the stakeholders served, 2) the auction mechanism adopted, and 3) the goals of the agents, to provide readers with a multi-perspective view into this field. In addition, we analyse the limitations of existing approaches, summarise the commonly adopted performance evaluation metrics, and discuss promising future directions leading towards effective and efficient stakeholder-oriented decision support in IA-AFL ecosystems.

1 Introduction

Federated Learning (FL) is a collaborative machine learning (ML) paradigm that is able to train useful models while respecting user privacy and data confidentiality [Yang *et al.*, 2019; Konečný *et al.*, 2016; Zhang *et al.*, 2021]. FL has gained significant attention from academia [Yang *et al.*, 2019] and industry [Liu and others, 2020; Liu and others, 2022a] alike, leading to a diverse range of techniques [Kairouz *et al.*, 2021]. In FL, there are two types of participants: *data consumers* (DCs, who often perform the role of FL servers), overseeing the distribution and aggregation of global FL models, and *data owners* (DOs, who often play the role of FL clients), responsible for training the FL model using their local data. FL follows a distributed approach where each DO trains a local model on its private dataset, and shares it with the corresponding DC. The DC then aggregates the received local models following an aggregation algorithm (e.g., FedAvg [McMahan and others, 2017]) to obtain the global model, which is then distributed back to the DOs for further training until convergence criteria are met. This design ensures that private local data are not exposed to any party other than the original owner, thus reducing privacy risk.

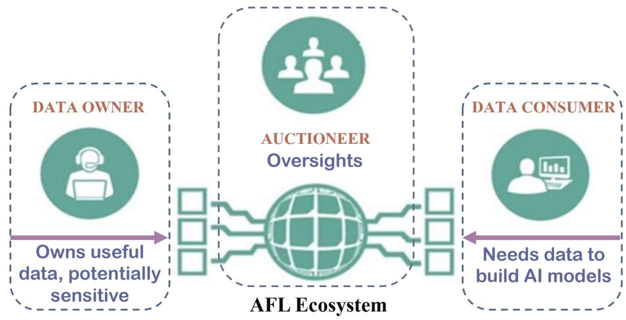

Figure 1: An overview of the AFL ecosystem.

Despite these advantages, existing FL works generally assume that all DOs agree to participate in the FL training process when requested [Thi Le and others, 2021]. However, in practice, DOs are self-interested entities who consider a complex set of factors (e.g., costs, potential risks of privacy exposure, expected utility gains) before deciding to join an FL task. This has motivated the study of FL incentive mechanisms [Khan and others, 2020], which aims to develop effective mechanisms that align the interests of DOs with the goals of DCs. They play a crucial role in encouraging DOs to actively participate in FL and make valuable contributions, ultimately leading to improved performance and broader adoption of FL in real-world applications.

Auction-based approaches have gained significant attention recently as an effective way to design FL incentive mechanisms. They offer a promising approach to motivating DOs to participate in FL in a fair and efficient manner. Under the typical auction-based FL (AFL) setting[1], three key stakeholders are involved: 1) DCs, 2) DOs, and 3) an auctioneer (as illustrated in Figure 1). The auctioneer plays a crucial role

[1]A possible example open AFL marketplace can be the Hierarchical Auctioning in Crowd-based Federated Learning system [Gao *et al.*, 2023]: https://hacfl.federated-learning.org/.

in coordinating the auction process, while DOs and DCs provide the auctioneer with their available data resources and bid values, respectively. The auction process as well as the entire AFL ecosystem center around the decision-making process of each stakeholder. The decisions made by each stakeholder impact the outcomes of AFL. To deal with the complexity, dynamism and personal nature of the context and the decision-making process, intelligent agents are often adopted to provide these stakeholders with AFL decision support, thereby inspiring the field of Intelligent Agents for AFL (IA-AFL).

IA-AFL is highly interdisciplinary in nature. It requires expertise from machine learning, multi-agent systems, game theory and auction theory, etc. This makes it challenging for researchers new to the field to grasp the latest developments. Currently, there is no survey paper on this important and rapidly developing field. To bridge this gap, we conduct a comprehensive survey of research works focusing on IA-AFL in this paper.[2] We analyse the AFL ecosystem in detail, with a focus on the diverse stakeholders involved and their decision-making priorities. Based on this analysis, we propose a unique multi-tiered taxonomy of IA-AFL that organises existing works according to 1) the stakeholders served, 2) the auction mechanism adopted, and 3) the goals of the agents to provide readers with a multi-perspective view into this field. In addition, we analyse the limitations of existing approaches, summarise the commonly adopted performance evaluation metrics, and discuss promising future directions towards effective and efficient stakeholder-oriented decision support in IA-AFL ecosystems.

2 Preliminaries

2.1 A Typical AFL Ecosystem

As shown in Fig. 1, a typical AFL ecosystem involves three primary stakeholders [Tang and Yu, 2023c]: 1) DOs, who act as the sellers possessing potentially sensitive but valuable data and training resources; 2) DCs, who act as buyers of such data and training resources to build ML models; and 3) an auctioneer, overseeing the matching of DOs with DCs and providing essential governance oversight for the ecosystem.

DCs submit their bidding profiles (including the bidding prices and their FL tasks) to the auctioneer. DOs submit their asking profiles (including the FL tasks they are able to join and their asking prices) to the auctioneer. The auctioneer determines the winners, and the corresponding market prices based on the submitted asking profiles and the bidding profiles under a predefined auction mechanism, and informs the winners. The winning DCs then pay the DOs. Through such an auction process, each DC recruits DOs to join its FL task. Afterward, each DC orchestrates the FL model training process with its recruited DOs following an adopted FL protocol.

2.2 AFL Stakeholder Concerns

In the AFL ecosystem, the stakeholders play distinct roles with different interests and concerns.

[2]Although some of the papers included in this survey do not explicitly mention agents, their focus on providing decision support for stakeholders in AFL reflects their potential as useful building blocks for realizing an agent-based AFL system.

The **auctioneer**'s role is pivotal, overseeing the auction process and facilitating information flow between participating DOs and DCs. Its main focus is to maintain the sustainable operation of the AFL ecosystem by attracting and retaining more participants, optimizing key performance indicators for the entire ecosystem, and providing governance oversight.

Data consumers, acting as buyers in the auction market, are primarily concerned with effective selection or bidding for DOs to meet their key performance indicators (KPIs), while staying within budget constraints.

Data owners, acting as sellers in the auction market, prioritize maximizing their monetary rewards. They are also keen on safeguarding data privacy by optimizing data resource allocation and the setting of reserve prices (i.e., the minimum acceptable price for selling the corresponding data resources).

2.3 Terminology

For ease of understanding, we provide a brief overview of key terminology adopted by the AFL field:

Commodity / data resources: In AFL, the term commodity refers to the object being exchanged between DCs and DOs, denoting a specific value for buying or selling. It can represent a unit of data (e.g., a training sample), communication bandwidth committed by a DO, or a unit of compute resource. In this paper, we use the terms data resources and commodity interchangeably unless a specific distinction is necessary.

Valuation: Valuation in AFL involves the assessment of the monetary value of data resources. Different DCs and DOs may assign value to data resources differently based on their individual preferences. Valuation can be either private, undisclosed to others, or public.

Utility: For DCs, utility is defined as the difference between their valuation of the auctioned data resources and the eventual payment made for those resources. For DOs, utility is defined as the difference between the total payments received from DCs and the costs incurred for the data resources, including communication and computation costs.

Social welfare (SW): SW is the sum of utilities for some or all participants in an AFL ecosystem. It provides a measure of the collective benefit derived from all transactions.

2.4 Types of Auction

AFL ecosystems can adopt various auction mechanisms based on their specific application scenarios [Qiu and others, 2022], including 1) double auction, 2) combinatorial auction, 3) reverse auction, and 4) forward auction. Double auctions [Friedman, 2018] accommodate multiple DOs and DCs, with both sides submitting asks and bids to the auctioneer. Combinatorial auctions [De Vries and Vohra, 2003] are effective when DCs bid for data resource bundles, ideal for acquiring complementary data types. Reverse auctions [Parsons and others, 2011] involve DOs competing for FL tasks, while forward auctions involve DCs competing for data resources.

Winner determination and pricing methods in AFL auctions fall into three categories [Tu and others, 2022]: 1) first-price sealed-bid (FPSB) , 2) second-price sealed-bid (SPSB) , and 3) Vickrey Clarke-Groves (VCG). Under FPSB, the highest bidder wins the auction and pays the bid price. The simplicity of FPSB might lead to inefficiencies and overpayment. Under SPSB, the highest bidder wins the auction, but pays the second-highest bid price. SPSB encourages truthful bidding to reveal true item valuation. Under VCG, winners are determined by maximizing the total benefit, considering externalities. Payments are determined based on the value contributed by other bidders for efficient and accurate price discovery [Vickrey, 1961].

3 The Proposed `IA-AFL` Taxonomy

Based on the stakeholders, the types of auctions involved in AFL and their respective goals, we propose a taxonomy for the `IA-AFL` literature as shown in Figure 2. Specifically, it first separates `IA-AFL` literature into data consumer-oriented, data owner-oriented, and auctioneer-oriented methods. Since all auction mechanisms introduced in the last section can be adopted by the AFL process, we further classify `IA-AFL` works based on their respective adopted auction mechanisms. Then, as stakeholders can have different goals, we further divide `IA-AFL` works based on their objectives. This hierarchical taxonomy provides a clear overview of the current `IA-AFL` landscape.

3.1 Intelligent Agents for Data Consumers

Based on the adopted auction mechanism, DC-oriented `IA-AFL` works can be broadly categorized into two distinct groups: 1) those designed for reverse auctions, and 2) those designed for forward auctions. These agents are instrumental in facilitating strategic decision-making for DCs, ensuring effective participation in the AFL market while maximizing key performance indicators (KPIs) derived from the collaborative FL model training process.

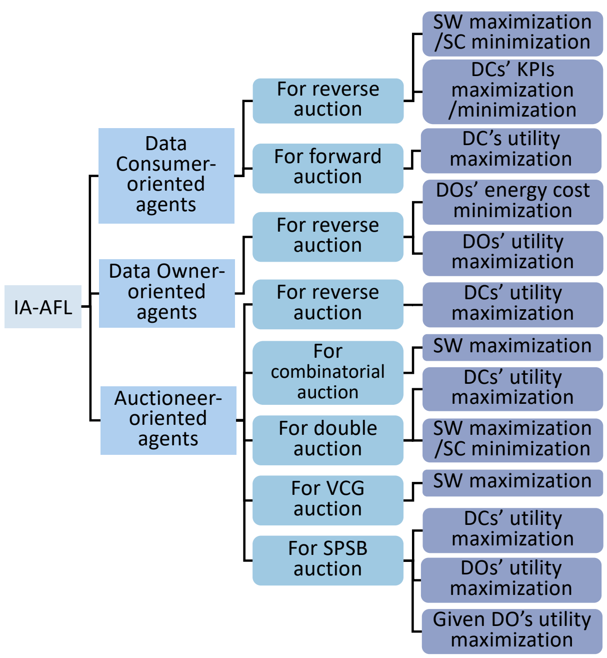

Figure 2: The `IA-AFL` taxonomy. DC, DO, SW and SC denote data consumer, data owner, social welfare and social cost, respectively.

For Reverse Auction

Under reverse auction, existing methods assume that there is only one DC and multiple DOs in the AFL marketplace. The intelligent agent for the DC plays a crucial role in selecting DOs. It makes decisions by evaluating DOs' asking profiles, assessing their potential contributions to the model, and aligning with the DC's objectives. Existing `IA-AFL` works for DCs under reverse auction can be broadly classified into two categories based on their designed objectives: 1) social welfare / social cost optimization approaches, and 2) DC KPI optimization approaches.

Social welfare / social cost optimization: To optimize the social welfare objective, [Jiao and others, 2020] first groups DOs based on Earth Mover's Distances (EMD) [Zhao and others, 2018]. The DC then greedily selects DOs from each group, determining payments based on marginal virtual social welfare density. To 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. In [Le and others, 2020], the workflow is similar, with a key distinction in the formulation of the DO selection process as a social cost minimization problem.

However, these works primarily focus on DO selection and payment determination over a single FL communication round. In [Zhou *et al.*, 2021], the DC is assisted in selecting and paying DOs for different FL communication rounds. The work decomposes the social cost minimization problem into a series of winner determination problems (WDPs) based on the number of global FL iterations. Each WDP is solved using a greedy algorithm to determine winning DOs, and a payment algorithm for computing remuneration to the winners. In [Yuan and others, 2021], the focus is on social cost minimization over the long run. The proposed FLORA method utilizes multiple polynomial-time online algorithms, including a fractional online algorithm and a randomized rounding algorithm, to select winning DOs and control the training accuracy of the global FL model. It also includes a payment algorithm to assist the DC in decision-making regarding DO selection and payment determination.

Different from the above two methods, which are designed for social cost minimization, [Wu and others, 2023] focuses on social welfare maximization. To achieve this goal, the proposed method follows deep reinforcement learning to select DOs and determine their payments under the VCG auction.

Data consumer KPI optimization: In [Fan *et al.*, 2020], the proposed method DQDRA maximizes the DC's valuation by determining winning DOs and the corresponding payments with a monotone greedy algorithm after receiving asking profiles from all DOs. Unlike DQDRA, which requires knowledge about the global distribution of all data for winning DO determination, RRAFL proposed in [Zhang *et al.*, 2021] leverages blockchain and reputation mechanisms instead. Winning DOs are selected based on their respective reputation, which are evaluated through historical contributions to the global FL model stored on the blockchain. Expanding on this, [Zhang *et al.*, 2022a] enhances RRAFL by introducing a novel contribution evaluation method using weighted samples. This adds nuance to the evaluation process, potentially offering a more sophisticated understanding of individual DOs' contributions. In [Zhang *et al.*, 2022b], RRAFL is extended by segmenting FL training tasks into multiple time steps based on global iterations, allowing adaptation to online learning applications. In [Zeng *et al.*, 2020], the proposed method FMORE helps the DC select the top K DOs with the highest score using the Lagrange multiplier method. [Batool *et al.*, 2022; Batool *et al.*, 2023] follow a similar method by incorporating blockchains [Kang and others, 2020; Kim *et al.*, 2019] and contract theory [Kang and others, 2019] to select DOs.

The aforementioned reputation-based DO selection methods do not explicitly consider the quality of the DOs, which is crucial for FL model performance. To address this limitation, [Deng and others, 2021] proposed FAIR, which integrates a quality-aware model aggregation algorithm with the reverse auction mechanism. FAIR determines winning DOs using a greedy algorithm based on Myerson's theorem [Myerson, 1998] to maximize the DC's valuation.

Unlike methods determining winning DOs and the corresponding payments in one communication round with a given budget, [Yang and others, 2023; Tan *et al.*, 2023; Tan and Yu, 2023] study how to allocate the DC's budget across multiple global FL communication rounds. [Yang and others, 2023] proposed BARA, an online reward budget allocation algorithm based on Bayesian optimization. Considering the urgency of recruitment, [Tan *et al.*, 2023; Tan and Yu, 2023] help the DC determine time-averaged optimal budget allocation for DOs.

Limitations: Existing works in this area often operate under the assumption of a monopolistic AFL market, where multiple DOs vie to join the FL training tasks of a single DC. However, this assumption diverges from the reality of practical AFL marketplaces, where numerous DCs may compete to attract multiple DOs for their respective FL training tasks.

For Forward Auction

Works in this field focus on maximizing the utility of a given DC within an AFL marketplace, which often involves multiple DCs. In [Tang and Yu, 2023c], a utility-maximizing bidding strategy, FedBidder, is designed for the DCs. It leverages various auction-related insights (e.g., DOs' data distributions, suitability to the task, DCs' bidding success probabilities, and budget constraints). The study emphasizes the crucial roles played by the estimation of DOs' utility and the appropriate winning function design in determining the optimal bidding function. To solve the optimal bidding function effectively, a utility estimation algorithm was proposed with two representative winning functions introduced, deriving two forms of optimal bidding functions for the DCs.

However, this approach overlooks the intricate relationships among DCs, which can be simultaneously competitive and cooperative. To address this issue, researchers have explored incorporating more than one agent for each DC. In [Tang and Yu, 2023a], the AFL ecosystem is modeled as a multi-agent system to guide DCs in strategically bidding towards an equilibrium with desirable overall system characteristics. The proposed approach, MARL-AFL, assigns two agents to each DC: 1) a bidding agent for determining bid prices, and 2) a bar agent for setting the bidding lower bound for the corresponding bidding agent. The bar agent is introduced to address potential collusive behaviors among bidding agents, such as bidding with an extremely low bid price, which can be detrimental to the health of the entire ecosystem. Both the bidding agents and the bar agents are designed based on deep Q-networks (DQN) [Mnih and others, 2015].

In [Tang and Yu, 2023b], MultiBOS-AFL is proposed to assist the DC in bidding for DOs in competitive AFL marketplaces. Unlike FedBidder and MARL-AFL, which assume that the entire team of DOs required for an FL task must be assembled before training can commence, MultiBOS-AFL helps the DC bid for DOs gradually over multiple FL model training sessions. To achieve this goal, each DC is assigned two agents: one for optimizing inter-session budget pacing, and the other for optimizing intra-session bidding.

Limitations: In this area, existing studies often assume that DOs arrive sequentially before the auction begins. However, real-world scenarios frequently involve DOs arriving in diverse orders, either before or during FL training tasks. The current body of research lacks robust solutions to navigating these dynamic and evolving situations effectively.

3.2 Intelligent Agents for Data Owners

In AFL, DOs function as the sellers, offering their valuable data resources to DCs. This transaction leads them to eventually become participants in the FL training processes initiated by various DCs, with the prospect of receiving monetary rewards in return. Consequently, intelligent agents tailored for DOs play a crucial role in providing guidance on strategic decision-making related to the allocation of their data resources and determining the asking profiles for these resources. Their final objective is to optimize the monetary profits derived from their involvement in AFL.

For Reverse Auction

Data owner energy cost minimization: In [Thi Le and others, 2021], the data resource trading process between a data consumer and multiple data owners is modeled as a reverse auction. Upon receiving FL training task profiles from the data consumer, which include the maximum tolerable time for FL training, each data owner optimizes asking profiles. These profiles, encompassing parameters like uplink transmission power, local accuracy level, and CPU cycle frequency, are fine-tuned iteratively to minimize energy costs.

Data owner utility maximization: In [Lu *et al.*, 2023], a within-cluster DO selection scheme was proposed for reverse auction to address the problem of uneven data resource consumption in a given cluster. DOs determine bid prices by maximizing their total utility. Similarly, [Le and others, 2020] also focuses on maximizing DO utility. However, unlike [Lu *et al.*, 2023] which solves the utility maximization problem to obtain bid prices, [Le and others, 2020] aims to derive asking profiles including CPU cycle frequency, uplink transmission power and training costs, in order to maximize utility. In [Zeng *et al.*, 2020], when a DO receives an FL training task and a scoring function from the DC, the proposed method assists it in deciding whether to bid based on its available data resources. If the DO chooses to bid, decisions regarding the number of resources to allocate and the corresponding charges to the DC are made using Euler's method.

Limitations: To the best of our knowledge, only these four studies currently address the issue of agent-based DO decision support. However, each of these works only concentrates on a single aspect of a DO's concerns. In practice, each decision made by a DO should encompass multiple facets simultaneously to meet its KPIs. Focusing solely on one aspect may lead to sub-optimal solutions.

3.3 Intelligent Agents for the Auctioneer

In an AFL ecosystem, the auctioneer serves as the coordinator and administrator, overseeing the flow of information between DOs and DCs, and facilitating the trading processes. Therefore, intelligent agents designed for the auctioneer are pivotal in offering strategic guidance for matching DOs and DCs. The ultimate goal is to optimize the monetary profits derived from their engagement within the AFL ecosystem. Existing methods in this domain are designed for four main auction mechanisms: 1) reverse auction, 2) combinatorial auction, 3) double auction, and 4) VCG/SPSB auction.

For Reverse Auction

Data consumer utility maximization: In [Seo and others, 2021], the auctioneer, represented by the software-defined network controller, facilitates decision-making between the DC and DOs. It determines the minimum number of global communication rounds required to meet the quality requirements of the FL model. This decision-making process occurs within the context of a reverse auction-based data trading system. Similarly, in [Seo and others, 2022], a software-defined network controller serves as the auctioneer, positioned between the DC and DOs. The proposed method in this paper assists the auctioneer in making decisions during the selection of winning DOs. The objective is to maximize the utility of the DC, via a greedy method.

Limitations: Like the `IA-AFL` approaches designed for the DC under reverse auction, these methods also operate under the assumption of a monopolistic AFL market. This assumption might constrain the practical applicability of these methods in real-world scenarios.

For Combinatorial Auction

Social welfare maximization: [Xu and others, 2023] aims to maximize social welfare and protect the utility of the auctioneer. The approach involves two main stages: 1) the combinatorial auction stage, where the platform selects winners who make the total utility of the platform and themselves greater than zero, and 2) the bargaining stage, where winners are classified into two categories with different payment methods after completing the training model. The goal is to ensure the utility of the auctioneer remains positive.

Limitations: [Xu and others, 2023] operates under the premise of a monopoly AFL market, assuming a single platform orchestrating the auction processes. While this setting provides a basis for understanding, a critical challenge lies in expanding participation, particularly attracting more DOs to engage in AFL. Enticing a diverse range of participants and optimizing the platform's functionality under more realistic, competitive scenarios remains an open area for exploration.

For Double Auction

Under double auction settings, the auctioneer agent ultimately coordinates agents serving DOs and DCs. Therefore, they are treated as auctioneer agents by extension.

Data consumer utility maximization: FEST [Roy and others, 2021] matches DOs and DCs with the goal of maximizing DC utility. This utility is a composite function involving the DC's valuation for data resources, the DO's asking price, and the corresponding execution time and reputation value. FEST assist DOs in determining winning candidate DCs using a greedy approach, followed by helping DCs select DOs to maximize their utility.

Social welfare / social cost optimization: [Mai *et al.*, 2022] assists the auctioneer in matching DCs and DOs, with the aim of maximizing social welfare. DOs submit asking profiles, and DCs submit bidding profiles to the auctioneer, which, in turn, uses the Lagrangian function to perform DO-DC matching. In [Wang and others, 2023], the focus is on social cost minimization under double auction. The authors formulate a nonlinear mixed-integer program for long-term social cost minimization. They propose an algorithmic approach to generate candidate training schedules and solve the problem using an online primal-dual-based algorithm [Buchbinder and others, 2009] with a carefully embedded payment design.

Limitations: Current methods predominantly operate under a centralized framework, where a central server continuously aggregates global system information and computes optimal decisions for the auctioneer. While the merits of a centralized architecture, such as rapid convergence and global optimality, are evident, they come at the cost of significant communication and computation overhead, especially in large-scale AFL ecosystems. Whenever there are shifts in the requirements of DCs, the auctioneer must collect extensive information across the entire ecosystem and recompute decisions. Moreover, in the event of hardware failures or attacks on the auctioneer, the entire ecosystem can be compromised.

For VCG Auction

Social welfare maximization: FVCG [Cong and others, 2020b] helps the auctioneer determine the amount of acceptable data to maximize its utility, factoring in data quality and privacy cost from DOs. It adopts a composite neural network-based payment function to derive payments for each DO, aiming to maximize social welfare and ensure fairness among DOs. Extending FVCG, [Cong and others, 2020a] introduced PVCG, which incorporates a game-theoretical model for the co-creation of virtual goods. PVCG helps the auctioneer determine the acceptance of input resources from each DO based on its asking profile, and imposes penalties if it fails to deliver the claimed resources. The objective is to maximize social welfare and mitigate information asymmetry.

Limitations: As the number of DOs increases, the need for more effective and efficient models to learn how to compensate DOs effectively becomes apparent for both FVCG and PVCG. Furthermore, it is crucial to evaluate the effectiveness of FVCG and PVCG in comparison to other sharing rules, such as Shapley value [Liu and others, 2022b] and labour union [Gollapudi and others, 2017].

For SPSB Auction

Data consumer utility maximization: In [Xu and others, 2021], a multi-bid auction mechanism is introduced to address bandwidth allocation challenges for self-interested DCs. The primary objective is to maximize the utility of DCs. Under this method, DCs submit bidding profiles specifying their requested bandwidth and unit price to the auctioneer. The auctioneer then allocates the bandwidth to DCs based on the market clearing price, and each DC incurs charges according to the SPSB auction mechanism.

Data owner utility maximization: In [Lim and others, 2020], the focus is on multiple DCs engaging in competitive bidding for data resources from a specific DO. The bids from DCs undergo a transformation, and the winning DCs are selected, with payments determined using the SPSB auction mechanism. The overarching objective is to maximize the utility of the DO. [Ng and others, 2020a; Ng and others, 2020b] incorporate Unmanned Aerial Vehicles (UAVs) as wireless relays to enhance communication between DOs and DCs. The optimal coalitional structure between UAV coalitions and DO coalitions is determined through the SPSB auction, aiming to maximize the utility of the UAV coalitions.

Limitations: Existing works in this area operate under the assumption that a DO can participate in at most one FL training task at any given time. In practice, DOs may have spare capacities to engage in multiple FL tasks concurrently. In such cases, resource allocation strategies should consider both the bandwidth and computing resources of the DOs. Exploring and adapting auction mechanisms to address the complexities arising from DOs' simultaneous involvement in multiple FL tasks is an open research question.

4 Evaluation Methodology

To assess `IA-AFL` methods, a combination of theoretical analysis and experimental evaluation is commonly adopted.

4.1 Theoretical Analysis

Given the nature of the auction and the emphasis on incentive mechanisms in FL, `IA-AFL` methods are expected to attain certain desirable properties [Zeng and others, 2021; Qiu and others, 2022; Ali and others, 2021].

1. *Budget Balance (BB)*: The budget balance property should hold, i.e., the total payments for DOs must not surpass the budget allocated by the DCs.
2. *Collusion Resistant (CR)*: This property imposes that no subgroups of participants can achieve higher profits through collusion or unethical conduct.
3. *Pareto Efficiency (PE)*: `IA-AFL` methods must meet the PE requirement when maximizing the social welfare of the entire AFL ecosystem.
4. *Fairness*: This property means that the entire AFL ecosystem should achieve a predefined fairness notion, such as contribution fairness, regret distribution fairness, or expectation fairness [Shi and Yu, 2023].
5. *Individual Rationality (IR)*: An `IA-AFL` method is deemed IR only if the profits for all participants are non-negative.
6. *Incentive Compatibility (IC) / Truthfulness*: Achieving IC/Truthfulness indicates that it is optimal for all participants to truthfully declare their contributions and cost types. Reporting untruthful information does not yield additional gain.
7. *Computational efficiency (CE)*: This property demands that the incorporated agents must guarantee the completion of the auction process and payment within polynomial time for operational efficiency in AFL.

4.2 Experimental Evaluation Metrics

Experimental evaluation plays a pivotal role in assessing and validating the efficacy of `IA-AFL` methods. It is instrumental in gauging the performance of these agents under complex settings. The following experimental evaluation metrics are commonly adopted by existing literature to quantitatively measure the effectiveness and impact of `IA-AFL`:

1. *Quality-of-Experience (QoE)*. QoE is expressed as the ratio between FL task completion time to the deadline of the task. It measures the speed at which a DC receives service from a DO, providing insights into the responsiveness and efficiency of the `IA-AFL` method.
2. *Utility*. It reflects the utility attained by DCs or DOs during the successful execution of FL tasks. A higher value indicates greater satisfaction with the received results, offering insights into the effectiveness of decisions made by the `IA-AFL` method. It can be expressed in various forms (e.g., the averaged form or the summation form).
3. *Task Completion Ratio*. This metric is expressed as the number of successful trades by DCs and is calculated as the ratio of the total number of winning DCs to the total number of DCs in the AFL marketplace. A higher task completion ratio indicates that more FL tasks are successfully allocated to DOs, providing a measure of the efficiency of the `IA-AFL` method.
4. *Payment*. Payment for DOs quantifies the financial compensation they received for the successful completion of FL tasks. This metric reflects the economic incentive and compensation provided to DOs, highlighting their contributions to the AFL marketplace under the given `IA-AFL` method.
5. *Social welfare*: Social welfare is a comprehensive metric that considers the collective well-being or total utility of all participants in the AFL marketplace, including both DCs and DOs. It provides a holistic measure of the overall effectiveness and fairness of the AFL ecosystem by considering the welfare of all stakeholders.

5 Promising Future Research Directions

Through our survey, it can be observed that AFL is still in its early stages of development, with various challenges yet to be addressed. This section delves into potential future directions for this nascent and interdisciplinary field.

5.1 Dynamic Decision Update

Existing `IA-AFL` methods are generally static approaches, represented by linear or non-linear functions. These functions derive their parameters from historical auction data through heuristic techniques. However, these static methods face a challenge when applied to new auctions, as the dynamics of these auctions may differ significantly from historical data. The inherent dynamism of the AFL market poses a considerable obstacle for static bidding methods to achieve desired outcomes in novel auction scenarios consistently.

To address this challenge, incorporating dynamic decision updates for both DOs and DCs, in accordance with the principles of demand-supply economics [Nedelec and others, 2022], is a promising direction. Such dynamic pricing approaches extend the auctioneer's role as well. A promising avenue for future exploration involves utilizing deep learning approaches to comprehend and model the behaviors of both DOs and DCs. Integrating these learned behaviors into various decision-making processes holds the potential to significantly enhance their utilities, adapting to the evolving dynamics of AFL marketplaces.

5.2 Multi-Agent Systems

AFL involves diverse stakeholders, each assuming distinct roles and harboring varied concerns. AFL, at its core, constitutes a multi-agent system (MAS), where intelligent entities interact dynamically within a complex framework. As illustrated in [Tang and Yu, 2023b], the relationships among DCs add a layer of intricacy, characterized by the simultaneous existence of both competition and cooperation. Moreover, within this ecosystem, the decision-making process of each participant carries direct or indirect repercussions on the choices made by other involved parties. Hence, adopting a MAS perspective to conceptualize AFL to provide a holistic understanding of the intricate interplay among diverse entities is a promising research direction [Kraus and others, 2023].

5.3 Preserving Privacy and Improving Security

Most existing auction-based mechanisms involve third-party entities, such as edge servers acting as auctioneers to manage each auction process. However, relying on third-party entities raises concerns about security and potential privacy breaches [Tang and Yu, 2022]. To address these challenges, several studies, including [Batool *et al.*, 2023; Zhang *et al.*, 2021; Batool *et al.*, 2022], utilize blockchain technology to safeguard trading information against tampering by malicious entities. However, implementing an auction algorithm within a blockchain network necessitates sharing private information among stakeholders, potentially giving rise to privacy concerns [Tang and Yu, 2022]. Moreover, in most existing works, DOs participate in the auction process without directly disclosing their private information, potentially dampening the enthusiasm of DOs. Therefore, a critical challenge arises in ensuring the security and reliability of auction mechanisms, while minimizing the risk of privacy leakage. In addition, it is essential to develop strategies to prevent malicious edge servers from launching attacks on DOs [Lyu *et al.*, 2020].

5.4 Online Auction Mechanisms

The current paradigm of `IA-AFL`, rooted in traditional auction methods, predominantly operates in an offline mode. This implies that the initiation of auctions relies on having a sufficient number of available bidders. For instance, in [Zeng *et al.*, 2020], the model aggregator initiates the process of determining winners once a satisfactory number of bids from DOs is received. In such offline auctions, both the DOs and the DCs may experience prolonged waiting times, even if they do not emerge as the eventual auction winners. This can discourage potential participants from actively engaging in the AFL marketplace. In contrast, online auction [Zhang and others, 2020] empowers the auctioneer, DCs and DOs to make real-time decisions, such as selecting winners and determining payments, as soon as a participant joins the auction. Online auctions offer the advantage of overcoming time and space constraints, ultimately resulting in cost savings. Therefore, online auction is a promising research direction for designing stronger incentive mechanisms in AFL.

5.5 Efficient Contribution Evaluation Methods

A crucial phase in the auction process involves the selection of the winning DOs, which heavily relies on evaluating the contributions of each DO. The prevailing approach employed by existing `IA-AFL` methods centers on contribution evaluation methods based on Shapley values. However, as highlighted in [Liu and others, 2022b], methods grounded in Shapley values are often time-consuming, posing a challenge to the computational efficiency when the system is scaled up. Furthermore, these methods operate under the assumption that DCs and other participants will truthfully assess the contribution of each DO, introducing a potential limitation in scenarios where honesty cannot be guaranteed. Hence, exploring alternative, more efficient contribution evaluation methods is a promising research direction to enhance the efficacy of `IA-AFL` methods.

5.6 Explainable AFL

As indicated by [Tang and Yu, 2022], explainability is an important aspect for auctions. Therefore, in the realm of AFL, an intriguing future direction is the advancement of Explainable AFL. This forward-looking approach entails the integration of mechanisms geared towards augmenting the transparency and interpretability of both the auction processes and federated training processes [Li *et al.*, 2023]. The implementation of explainability in AFL holds the potential to foster heightened levels of trust, accountability, comprehensibility and auditability regarding the decision-making processes involved in both the auction and the federated training phases.

6 Concluding Remarks

In this paper, we conduct a comprehensive review of `IA-AFL` methods through a unique multi-tiered taxonomy that organises existing works according to 1) the stakeholders served, 2) the auction mechanism adopted, and 3) the goals of the agents. Furthermore, we critically analyze the limitations of current approaches, outline commonly utilized performance evaluation methodologies, and deliberate on promising future directions. To the best of our knowledge, it is the first survey on `IA-AFL`, providing researchers with an accessible guide into this interdisciplinary field.

Acknowledgements

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); the RIE 2020 Advanced Manufacturing and Engineering (AME) Programmatic Fund (No. A20G8b0102), Singapore; the Natural Sciences and Engineering Research Council of Canada (NSERC); Public Safety Canada; CIFAR Catalyst Grant; and Compute Canada Research Platform. Kraus has been supported in part by MOST under grant 0004668 and the EU Project TAILOR under grant 952215.

References

[Ali and others, 2021] Asad Ali et al. Incentive-driven federated learning and associated security challenges: A systematic review, 2021.

[Batool *et al.*, 2022] Zahra Batool, Kaiwen Zhang, and Matthew Toews. Fl-mab: client selection and monetization for blockchain-based federated learning. In *SAC*, pages 299–307, 2022.

[Batool *et al.*, 2023] Zahra Batool, Kaiwen Zhang, and Matthew Toews. Block-racs: Towards reputation-aware client selection and monetization mechanism for federated learning. *SAC*, 23(3):49–66, 2023.

[Buchbinder and others, 2009] Niv Buchbinder et al. The design of competitive online algorithms via a primal–dual approach. *FTTCS*, 3(2–3):93–263, 2009.

[Cong and others, 2020a] Mingshu Cong et al. Optimal procurement auction for cooperative production of virtual products: Vickrey-clarke-groves meet cremer-mclean. *arXiv preprint arXiv:2007.14780*, 2020.

[Cong and others, 2020b] Mingshu Cong et al. A vcg-based fair incentive mechanism for federated learning. *arXiv preprint arXiv:2008.06680*, 2020.

[De Vries and Vohra, 2003] Sven De Vries and Rakesh V Vohra. Combinatorial auctions: A survey. *JOC*, 15(3):284–309, 2003.

[Deng and others, 2021] Yongheng Deng et al. Fair: Quality-aware federated learning with precise user incentive and model aggregation. In *INFOCOM*, 2021.

[Fan *et al.*, 2020] Sizheng Fan, Hongbo Zhang, Yuchen Zeng, and Wei Cai. Hybrid blockchain-based resource trading system for federated learning in edge computing. *IOTJ*, 8(4):2252–2264, 2020.

[Friedman, 2018] Daniel Friedman. *The double auction market: institutions, theories, and evidence*. Routledge, 2018.

[Gao *et al.*, 2023] Yulan Gao, Yansong Zhao, and Han Yu. Multi-tier client selection for mobile federated learning networks. In *ICME*, pages 666–671, 2023.

[Gollapudi and others, 2017] Sreenivas Gollapudi et al. Profit sharing and efficiency in utility games. In *ESA*, 2017.

[Jiao and others, 2020] Yutao Jiao et al. Toward an automated auction framework for wireless federated learning services market. *TMC*, 20(10):3034–3048, 2020.

[Kairouz *et al.*, 2021] Peter Kairouz, H. Brendan McMahan, et al. Advances and open problems in federated learning. *FTML*, 14(1-2):1–210, 2021.

[Kang and others, 2019] Jiawen Kang et al. Incentive mechanism for reliable federated learning: A joint optimization approach to combining reputation and contract theory. *IOTJ*, 6(6):10700–10714, 2019.

[Kang and others, 2020] Jiawen Kang et al. Reliable federated learning for mobile networks. *Wireless Commun*, 27(2):72–80, 2020.

[Khan and others, 2020] Latif U. Khan et al. Federated learning for edge networks: Resource optimization and incentive mechanism. *Commun Mag*, 58(10):88–93, 2020.

[Kim *et al.*, 2019] Hyesung Kim, Jihong Park, Mehdi Bennis, and Seong-Lyun Kim. Blockchained on-device federated learning. *Commun Lett*, 24(6):1279–1283, 2019.

[Konečný *et al.*, 2016] Jakub Konečný, H Brendan McMahan, Daniel Ramage, and Peter Richtárik. Federated optimization: Distributed machine learning for on-device intelligence, 2016.

[Kraus and others, 2023] Sarit Kraus et al. Customer service combining human operators and virtual agents: A call for multidisciplinary ai research. 2023.

[Le and others, 2020] Tra Huong Thi Le et al. Auction based incentive design for efficient federated learning in cellular wireless networks. In *WCNC*, pages 1–6, 2020.

[Li *et al.*, 2023] Anran Li, Rui Liu, Ming Hu, Luu Anh Tuan, and Han Yu. Towards interpretable federated learning. *arXiv preprint arXiv:2302.13473*, 2023.

[Lim and others, 2020] Wei Yang Bryan Lim et al. Incentive mechanism design for resource sharing in collaborative edge learning. *arXiv preprint arXiv:2006.00511*, 2020.

[Liu and others, 2020] Yang Liu et al. FedVision: An online visual object detection platform powered by federated learning. In *IAAI*, pages 13172–13179, 2020.

[Liu and others, 2022a] Zelei Liu et al. Contribution-aware federated learning for smart healthcare. In *IAAI*, pages 12396–12404, 2022.

[Liu and others, 2022b] Zelei Liu et al. GTG-Shapley: Efficient and accurate participant contribution evaluation in federated learning. *TIST*, 13(4):1–21, 2022.

[Lu *et al.*, 2023] Renhao Lu, Weizhe Zhang, Yan Wang, Qiong Li, Xiaoxiong Zhong, Hongwei Yang, and Desheng Wang. Auction-based cluster federated learning in mobile edge computing systems. *TPDS*, 34(4):1145–1158, 2023.

[Lyu *et al.*, 2020] Lingjuan Lyu, Han Yu, and Qiang Yang. Threats to federated learning: A survey. *arXiv preprint arXiv:2003.02133*, 2020.

[Mai *et al.*, 2022] Tianle Mai, Haipeng Yao, Jiaqi Xu, Ni Zhang, Qifeng Liu, and Song Guo. Automatic double-auction mechanism for federated learning service market in internet of things. *TNSE*, 9(5):3123–3135, 2022.

[McMahan and others, 2017] Brendan McMahan et al. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*, pages 1273–1282, 2017.

[Mnih and others, 2015] Volodymyr Mnih et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

[Myerson, 1998] Roger B Myerson. Population uncertainty and poisson games. *IJGT*, 27:375–392, 1998.

[Nedelec and others, 2022] Thomas Nedelec et al. Learning in repeated auctions. *FTML*, 15(3):176–334, 2022.

[Ng and others, 2020a] Jer Shyuan Ng et al. Communication-efficient federated learning in UAV-enabled iov: a joint auction-coalition approach. In *GLOBECOM*, pages 1–6, 2020.

[Ng and others, 2020b] Jer Shyuan Ng et al. Joint auction-coalition formation framework for communication-efficient federated learning in UAV-enabled internet of vehicles. *TIST*, 22(4):2326–2344, 2020.

[Parsons and others, 2011] Simon Parsons et al. Auctions and bidding: A guide for computer scientists. *CSUR*, 43(2):1–59, 2011.

[Qiu and others, 2022] Houming Qiu et al. Applications of auction and mechanism design in edge computing: A survey. *TCCN*, 8(2):1034–1058, 2022.

[Roy and others, 2021] Palash Roy et al. Distributed task allocation in mobile device cloud exploiting federated learning and subjective logic. *JSA*, 113(2), 2021.

[Seo and others, 2021] Eunil Seo et al. Auction-based federated learning using software-defined networking for resource efficiency. In *CNSM*, pages 42–48, 2021.

[Seo and others, 2022] Eunil Seo et al. Resource-efficient federated learning with non-iid data: An auction theoretic approach. *IOTJ*, 9(24):25506–25524, 2022.

[Shi and Yu, 2023] Yuxin Shi and Han Yu. Fairness-aware client selection for federated learning. In *ICME*, 2023.

[Tan and Yu, 2023] Xavier Tan and Han Yu. Hire when you need to: Gradual participant recruitment for auction-based federated learning. *arXiv preprint arXiv:2310.02651*, 2023.

[Tan *et al.*, 2023] Xavier Tan, Wei Yang Bryan Lim, Dusit Niyato, and Han Yu. Reputation-aware opportunistic budget optimization for auction-based federation learning. In *IJCNN*, pages 1–8, 2023.

[Tang and Yu, 2022] Xiaoli Tang and Han Yu. Towards trustworthy ai-empowered real-time bidding for online advertisement auctioning. *arXiv preprint arXiv:2210.07770*, 2022.

[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 *ICME*, 2023.

[Thi Le and others, 2021] Tra Huong Thi Le et al. An incentive mechanism for federated learning in wireless cellular networks: An auction approach. *TWC*, 20(8):4874–4887, 2021.

[Tu and others, 2022] Xuezhen Tu et al. Incentive mechanisms for federated learning: From economic and game theoretic perspective. *TCCN*, 8(3):1566–1593, 2022.

[Vickrey, 1961] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *JF*, 16(1):8–37, 1961.

[Wang and others, 2023] Fei Wang et al. Toward sustainable ai: Federated learning demand response in cloud-edge systems via auctions. In *INFOCOM*, pages 1–10, 2023.

[Wu and others, 2023] Leijie Wu et al. Long-term adaptive vcg auction mechanism for sustainable federated learning with periodical client shifting. *TMC*, 2023.

[Xu and others, 2021] Jie Xu et al. Bandwidth allocation for multiple federated learning services in wireless edge networks. *TWC*, 21(4):2534–2546, 2021.

[Xu and others, 2023] Bo Xu et al. Cab: a combinatorial-auction-and-bargaining-based federated learning incentive mechanism. *WWW*, pages 1–22, 2023.

[Yang and others, 2023] Yunchao Yang et al. Bara: Efficient incentive mechanism with online reward budget allocation in cross-silo federated learning. *arXiv preprint arXiv:2305.05221*, 2023.

[Yang *et al.*, 2019] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept and applications. *TIST*, 10(2):12:1–12:19, 2019.

[Yuan and others, 2021] Yulan Yuan et al. Incentivizing federated learning under long-term energy constraint via online randomized auctions. *TWC*, 21(7):5129–5144, 2021.

[Zeng and others, 2021] Rongfei Zeng et al. A comprehensive survey of incentive mechanism for federated learning. *arXiv preprint arXiv:2106.15406*, 2021.

[Zeng *et al.*, 2020] Rongfei Zeng, Shixun Zhang, Jiaqi Wang, and Xiaowen Chu. Fmore: An incentive scheme of multi-dimensional auction for federated learning in MEC. In *ICDCS*, pages 278–288, 2020.

[Zhang and others, 2020] Jixian Zhang et al. An online auction mechanism for time-varying multidimensional resource allocation in clouds. *FGCS*, 111:27–38, 2020.

[Zhang *et al.*, 2021] Jingwen Zhang, Yuezhou Wu, and Rong Pan. Incentive mechanism for horizontal federated learning based on reputation and reverse auction. In *WWW*, page 947–956, 2021.

[Zhang *et al.*, 2022a] Jingwen Zhang, Yuezhou Wu, and Rong Pan. Auction-based ex-post-payment incentive mechanism design for horizontal federated learning with reputation and contribution measurement. *arXiv preprint arXiv:2201.02410*, 2022.

[Zhang *et al.*, 2022b] 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.

[Zhao and others, 2018] Yue Zhao et al. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018.

[Zhou *et al.*, 2021] Ruiting Zhou, Jinlong Pang, Zhibo Wang, John CS Lui, and Zongpeng Li. A truthful procurement auction for incentivizing heterogeneous clients in federated learning. In *ICDCS*, pages 183–193, 2021.