
PRIOR-INDEPENDENT AUCTIONS FOR THE DEMAND SIDE OF FEDERATED LEARNING

Andreas Haupt
LIDS
MIT
Cambridge, MA 02446, USA
haupt@mit.edu

Vaikkunth Mugunthan
CSAIL
MIT
Cambridge, MA, 02446
vaik@mit.edu

ABSTRACT

Federated learning (FL) is a paradigm that allows distributed clients to learn a shared machine learning model without sharing their sensitive training data. While largely decentralized, FL requires resources to fund a central orchestrator or to reimburse contributors of datasets to incentivize participation. Inspired by insights from prior-independent auction design, we propose a mechanism, FIPIA (Federated Incentive Payments via Prior-Independent Auctions), to collect monetary contributions from self-interested clients. The mechanism operates in the semi-honest trust model and works even if clients have a heterogeneous interest in receiving high-quality models, and the server does not know the clients' level of interest. We run experiments on the MNIST, FashionMNIST, and CIFAR-10 datasets to test clients' model quality under FIPIA and FIPIA's incentive properties.

1 INTRODUCTION

Federated learning (FL) (Bonawitz et al., 2019; McMahan et al., 2017) is a mechanism that allows decentralized clients to collaboratively learn a machine learning model while preserving individual clients' privacy. That is, clients use datasets that share the same feature space but differ in sample space. However, without incentives, participating clients may provide obsolete information or decide to opt out of participation. Also, clients can submit inferior models yet end up receiving updated high-quality models. Hence, causing a free-rider problem (Feldman & Chuang, 2005).

There are numerous system designs (Chen et al., 2020; Jiao et al., 2019; Li et al., 2019; Yu et al., 2020) that target this problem by having systems that reimburse clients that submit high-quality models. Such systems hence design incentives for the supply of models to FL.

Correspondingly, we consider the clients receiving models and their incentives, the *demand side of FL*. We assume (an assumption to be relaxed in future work) that clients submit models trained on their entire dataset—with no need to incentivize them to do so. Under this assumption, we design a system to collect resources from clients.

To do so, we adapt auction designs to allocate freely-replicable (digital) goods (Hartline, 2013) to the FL setting and propose a system that works without central-entity evaluation. Our proposed system shares higher quality models with clients that pay more, as we demonstrate using the MNIST (LeCun et al., 1998), FashionMNIST (Xiao et al., 2017), and CIFAR-10 (Krizhevsky et al., 2009) datasets. We show experimentally that non-truthful bidding is only limitedly profitable.

2 LITERATURE REVIEW

First, our model connects to models incentivizing parity in test accuracy in FL. (Li et al., 2019) proposes a variant of a global objective criterion to induce parity in test accuracy between clients holding non-i.i.d data. While this literature also studies the supply side, i.e., who receives models, the paper implicitly assumes that the interest in high-quality models is homogeneous across clients and tries to induce parity without collecting any monetary contributions.

Our approach also connects to the literature on incentive design for FL (Chen et al., 2020; Jiao et al., 2019; Yu et al., 2020). These papers propose systems that distribute a fixed *incentive budget* to clients. Participating clients are promised payments depending on their model quality. These payments incentivize clients to *contribute* high-quality models, but do not generate resources from clients that *receive* models. Closest to our paper, (Chen et al., 2020) studies how to incentivize clients to submit their full datasets to an FL system. Crucially, in their paper, the clients’ private information is given by their maximal dataset size only, their interest in models they receive is assumed to be known to the system. One of the main properties of FIPIA is that it works without such knowledge.

This paper uses tools from the honest-but-curious or semi-honest trust model (Beaver, 1991). Our federated evaluation method is a variant of (Mugunthan et al., 2020), which operates in the same trust model.

Finally, we use an incentive design, the circuit auction, from prior-independent mechanism design (Hartline, 2013).

3 METHOD

We consider a horizontal FL setting. There are k clients $i \in [k], k := \{1, 2, \dots, k\}$ that can evaluate a model w on their test sets to get a measure of quality such as accuracy, $\text{Score}(i, w)$. In usual FL settings, all clients receive the same models w . In our approach, we relax this assumption. That is, each client i receives model w_i after every federated round. To determine who will get which model, we also introduce payments $t_i, i = 1, 2, \dots, k$, which clients pay to the server. These payments can be re-distributed amongst clients at a later stage using an existing incentive mechanism (Chen et al., 2020; Jiao et al., 2019; Yu et al., 2020).

Assumption 1. *Clients maximize a linear combination of payments they make and the quality of the model they receive,*

$$\theta_i \text{Score}(i, w_i) - t_i. \tag{1}$$

θ_i represent client i ’s valuation per unit of score. The valuations θ_i are not known to the server.

For example, Assumption 1 holds when clients incur a constant cost for each misclassified data point in their test set in each federated round. Implicit in this assumption is that clients are uninterested in other clients’ model qualities, i.e., there are no externalities between clients.

Assumption 2. *Clients submit models trained on their entire dataset.*

This assumption mirrors our focus on the *demand side* of FL. Our system leaves incentivizing high-quality models to other incentive mechanisms.

Assumption 3. *The datasets used by clients are drawn from the same distribution (i.i.d. data).*

An extension of the proposed system to the non-i.i.d. setting will involve additional complexities, compare (Hsu et al., 2019). We leave this for further work.

3.1 INSUFFICIENCY OF AVERAGE ACCURACY MAXIMIZATION

A standard FL system wishes to maximize the average model quality amongst all clients, i.e., it maximizes $\max_w \sum_{i=1}^k \lambda_i \text{Score}(i, w)$, where $\lambda_i \geq 0$, denotes the relative importance of client i . In the notation introduced above, this corresponds to $w_i = w$, for $i = 1, 2, \dots, k$.

Assuming that a central orchestrator or engineer for the model needs to be funded, a classical FL system might not raise a sufficient amount of money to make the system sustainable. For example, if funding is done through voluntary monetary contributions by clients, a free-rider problem arises (Mas-Colell et al., 1995, Section 11.C), and systems will likely be under-funded. Interestingly, the celebrated Vickrey-Clarke-Groves mechanism (Clarke, 1971) fails to collect *any* funds when applied on the demand side of FL. In this mechanism, each client that receives a model pays the system an amount corresponding to the negative effect it has on other clients. The receiving client pays the difference in valuation, according to (1), for all other clients when it participates in contrast to when it does not participate. As distributing the same model to an additional client does not have any effect on other clients in our model, *clients pay nothing*. Hence, a system that collects funds in environments with costly orchestration is needed.

Function	Description
$\text{Train}(i, w)$	Updated model after training w incrementally on client i 's data
$\text{Thresh}(b)$	Transmission probability based on bid b
$\text{Aggr}(W)$	Aggregated model of models $w \in W$
$\text{Punish}(d)$	(Monetary) punishment based on deviation d
$\text{Score}(i, w)$	Score of model w by client i

Table 1: Methods used in our algorithm.

3.2 FEDERATED INCENTIVE PAYMENTS VIA PRIOR-INDEPENDENT AUCTIONS (FIPIA)

Our main contribution is a system for collecting payments in the semi-honest trust model, presented in algorithm 1. Our system is inspired by auction formats that guarantee (up to a multiplicative constant) maximal revenue in a wide range of theoretical setting satisfying Assumptions 1 and 2 (Hartline, 2013; Goldberg & Hartline, 2001). Formalizing this optimality is part of follow-up work.

Algorithm 1: Federated Incentive Payments via Prior-Independent Auctions (FIPIA)

Input: Number of federated rounds T , Number of clients k

Client Initialization: Each client $i \in [k]$ initializes models w_i and sends bid b_i to server

for $t = 0, \dots, T$ **do**

foreach *Client* i **do**

$w_i \leftarrow \text{Train}(i, w_i)$

$(p, s, \hat{s}) \leftarrow \text{FedEval}(w, b)$

 Server draws a uniformly random permutation $\pi: [k] \rightarrow [k]$

foreach *Client* i **do**

if $b_i(\hat{s} - s_i) > b_{\pi(i)}(\hat{s} - s_{\pi(i)})$ **then**

 Client \hat{i} transfers model \hat{w} giving score \hat{s} to i

$t_i \leftarrow b_{\pi(i)}(\hat{s} - s_{\pi(i)}) + p_i$

else

$t_i \leftarrow p_i$

Function $\text{FedEval}(w, b)$:

Server Initialization: $a_{ij} \sim \text{Bernoulli}(\text{Thresh}(b_i)), i \neq j \in [k], a_{ii} = 1,$

$W_{i \rightarrow} := \{j \in [k] | a_{ij} = 1\}, W_{\rightarrow i} := \{j \in [k] | a_{ji} = 1\}$

foreach *Client* i **do**

 Send w_i to clients $W_{i \rightarrow}$

foreach *Client* j **do**

 Evaluate $\hat{w}_i \leftarrow \text{Aggr}(W_{\rightarrow j})$ and send \hat{w}_i to Clients $W_{j \rightarrow}$

foreach *Client* j **do**

 Evaluate $s_i^j \leftarrow \text{Score}(j, w_i)$ and $\hat{s}_i^j \leftarrow \text{Score}(j, \hat{w}_i), i \in [k]$ and send to Server

Server Evaluation: $s_i \leftarrow \text{med}((s_i^j)_{j \in W_{i \rightarrow}}), i \in [k]$

$\hat{s}_i \leftarrow \text{med}((\hat{s}_i^j)_{j \in W_{i \rightarrow}}), i \in [k]$

$\hat{s} \leftarrow \max_{i \in [k]} \hat{s}_i$

$p_j \leftarrow \sum_{i=1}^k \text{Punish}(s_i^j - s_i) + \text{Punish}(\hat{s}_i^j - \hat{s}_i)$

return p, s, \hat{s}

The mechanism starts with the submission of *bids* by clients, which are reports of their valuations θ_i . These reports might or might not be truthful. For each federated around T , clients first train on new data they receive. Then, our federated evaluation presented as a subroutine in algorithm 1 determines incentive payments for truthful evaluation, p , a vector of scores for each client's model s , the score \hat{s} of the model with the highest median score among all aggregated models. Then, the server matches each client randomly with another client. The server first computes

$$b_{\pi(i)}(\hat{s} - s_{\pi(i)})$$

from the median scores received in the federated evaluation. This value is assuming $\pi(i)$'s report was truthful ($b_{\pi(i)} = \theta_{\pi(i)}$) the gain in objective for $\pi(s)$ from being allocated the model. Client i gets the best model if their corresponding value

$$b_i(\hat{s} - s_i)$$

surpasses this value. The value for $\pi(i)$ hence serves as minimum payment for i and incentivizes high bidding. If i 's value surpasses $\pi(i)$'s value, the system declares i a winner, i gets the best model \hat{w} which produces score \hat{s} .

Our federated evaluation function defines for each client a random set of receiver clients. Clients that bid higher are more likely to receive models to evaluate. This design choice takes into account that these are also most likely to receive the best model \hat{w} in the course of algorithm 1. Clients evaluate all single and the average of all the models they receive. By sending these models to other agents and collecting the median of reported scores, the server obtains accurate estimates of individual and best model quality. Payments for reports deviating from the median incentivize truthful evaluations.

3.3 DESIGN CONSIDERATIONS

In this subsection, we discuss design considerations for the methods used in algorithm 1, compare Table 1.

In this paper, we consider a semi-honest trust model in which clients follow the protocol specification, but may attempt to learn honest clients' private information from the models it receives, or may collude with other clients to learn private information. Private information of honest clients can be learned by performing model inversion and membership inference attacks (Nasr et al., 2018; Melis et al., 2019; Geiping et al., 2020). To make our system compatible with the semi-honest trust model, clients can implement $\text{Train}(i, w)$ to produce differentially private models (Abadi et al., 2016; Wang et al., 2019; 2020) to prevent these attacks.

The threshold function $\text{Thresh}(b)$ determines how many models clients see in peer evaluations. While a high value for $\text{Thresh}(b)$ yields more accurate estimates of performance, it also induces incentives for lowering bids for clients that hope to receive a high-quality model as an evaluator, and not as a winner of the auction.


The punishment function incentivizes truthful reporting. In cases with heterogenous clients, choosing small values of $\text{Punish}(d)$ might be preferable, as otherwise incentives to participate in the system for clients with small test sets might be lowered.

The aggregation function $\text{Aggr}(W)$ can use arbitrary aggregation routines, be they federated averaging McMahan et al. (2017), or involve search for the best model among the models to aggregate.


4 EVALUATION

We conducted experiments to evaluate the performance of our mechanism. All code can be found under <https://github.com/Indraos/MultiSidedFederation>. We present results for the MNIST dataset (LeCun et al., 1998). Our code also runs experiments for the CIFAR-10 (Krizhevsky et al., 2009) and Fashion-MNIST (Xiao et al., 2017) datasets. The MNIST and Fashion-MNIST datasets consists of 70000 28x28 images each. There are 60000 training images and 10000 test images. The images are normalized using the training set mean and variance. For MNIST and FashionMNIST, each client trained a simple network with two convolution layers followed by two fully connected layers. The CIFAR-10 dataset consists of 60000 32x32 colour images in 10 classes, with 6000 images per class. There are 50000 training images and 10000 test images. For MNIST and FashionMNIST, each client trained a simple network with two convolution layers followed by three fully connected layers.


Our experiments uses $k = 3$ clients, with valuations 0.1, 0.5 and 0.6. Other valuations qualitatively do not change our results. The clients split the data in an i.i.d. fashion. Training dataset sizes are 50%, 40% and 10% of the dataset. To highlight incentive and adaptivity properties of our system without peer evaluation, we set $\text{Punish}(d) = 0$ and $\text{Thresh}(b) = 0$.




(a) Accuracy trajectories for clients in the MNIST setup.




(b) Client objective value when bidding a report unequal to valuation 0.5 in the MNIST setup.




(c) Accuracy trajectories for clients in the FashionMNIST setup.



(d) Client objective value when bidding a report unequal to valuation 0.5 in the FashionMNIST setup.



(e) Accuracy trajectories for clients in the CIFAR-10 setup.



(f) Client objective value when bidding a report unequal to valuation 0.5 in the CIFAR-10 setup.

Figure 1: Experiments for MNIST (top row), FashionMNIST (middle row) and CIFAR-10 (bottom row). The left column shows the values in different clients in our first experiment. The figures in the right column shows for an agent with valuation 0.5 the objective function when submitting bids different from its valuation.

4.1 THE EFFECTS OF BIDS AND DATASET SIZES

Our first experiment shows the histories of model qualities for clients with different dataset sizes and bids. We assume that all clients bid their valuation truthfully. Figure 1 represents the accuracies of different clients for the first 6 FL rounds. We note that how the quality of models for different clients adapts in our model according to the bids placed. In particular, the FashionMNIST case shows that the high-bidding model ends up with a higher-quality model than the other clients despite a much smaller dataset, through a combination of being declared winner in the second-to-last round and last-round training.

4.2 PROFITABILITY OF OVERBIDDING

Next, we study the incentive properties of our system, compare the right column of Figure 1. We calculate the profit from bidding a value different from the valuation θ_i . For our exposition, we restrict to a setting with one client that has a fixed valuation $\theta_i = 0.5$ and the two other agents have valuations 0.1 and 0.6. We calculate the values (1) for client i when it bids $b_i \in \{0, 0.2, 0.4, 0.5, 0.6, 0.8, 1\}$. We observe that some deviations are profitable, but that bidding close to truthfully is a or close to a local maximum of the objectives for the deviating agent.

5 CONCLUSION

In this paper, we introduce a system to secure funds from clients participating in FL. Our system is adaptive as it gives models of different qualities to clients that have different valuations. We investigated incentive properties of the system and find some instances of profitability of deviations.

Our work complements existing incentive designs for reimbursing model contributors in FL. In future work, we plan to explore the integration of these systems into an incentive system for the demand *and* supply side of FL and to give theoretical revenue guarantees.

REFERENCES

- Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pp. 308–318, 2016.
- Donald Beaver. Foundations of secure interactive computing. In *Annual International Cryptology Conference*, pp. 377–391. Springer, 1991.
- Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloe Kiddon, Jakub Konečný, Stefano Mazzocchi, H Brendan McMahan, et al. Towards federated learning at scale: System design. *arXiv preprint arXiv:1902.01046*, 2019.
- Mengjing Chen, Yang Liu, Weiran Shen, Yiheng Shen, and Pingzhong Tang. Mechanism Design for Multi-Party Machine Learning. 2020.
- Edward H Clarke. Multipart pricing of public goods. *Public choice*, pp. 17–33, 1971.
- Michal Feldman and John Chuang. Overcoming Free-Riding Behavior in Peer-to-Peer Systems. *ACM SIGecom Exchanges*, 5(4):41–50, 2005.
- Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. Inverting gradients—how easy is it to break privacy in federated learning? *arXiv preprint arXiv:2003.14053*, 2020.
- Andrew V Goldberg and Jason D Hartline. Competitive auctions for multiple digital goods. In *European Symposium on Algorithms*, pp. 416–427. Springer, 2001.
- Jason D Hartline. Mechanism design and approximation. *Book draft. October*, 122, 2013.
- Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Measuring the effects of non-identical data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335*, 2019.

-
- Yutao Jiao, Ping Wang, Dusit Niyato, Bin Lin, and Dong In Kim. Toward an automated auction framework for wireless federated learning services market. *arXiv*, pp. 1–14, 2019. ISSN 23318422. doi: 10.1109/tmc.2020.2994639.
- Alex Krizhevsky et al. Learning multiple layers of features from tiny images. 2009.
- Yann LeCun, Corinna Cortes, and Christopher JC Burges. The mnist database of handwritten digits, 1998. URL <http://yann.lecun.com/exdb/mnist>, 10:34, 1998.
- Tian Li, Maziar Sanjabi, and Virginia Smith. Fair resource allocation in federated learning. *arXiv*, pp. 1–27, 2019. ISSN 23318422.
- Andreu Mas-Colell, Michael Dennis Whinston, Jerry R Green, et al. *Microeconomic theory*, volume 1. Oxford university press New York, 1995.
- Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pp. 1273–1282. PMLR, 2017.
- Luca Melis, Congzheng Song, Emiliano De Cristofaro, and Vitaly Shmatikov. Exploiting unintended feature leakage in collaborative learning. In *2019 IEEE Symposium on Security and Privacy (SP)*, pp. 691–706. IEEE, 2019.
- Vaikkunth Mugunthan, Ravi Rahman, and Lalana Kagal. BlockFLow: An Accountable and Privacy-Preserving Solution for Federated Learning. *arXiv*, 2020.
- Milad Nasr, Reza Shokri, and Amir Houmansadr. Comprehensive privacy analysis of deep learning: Stand-alone and federated learning under passive and active white-box inference attacks. *arXiv preprint arXiv:1812.00910*, 2018.
- L Wang, R Jia, and D Song. D2p-fed: Differentially private federated learning with efficient communication. *arxiv.org/pdf/2006.13039*, 2020.
- Yu-Xiang Wang, Borja Balle, and Shiva Prasad Kasiviswanathan. Subsampled rényi differential privacy and analytical moments accountant. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1226–1235. PMLR, 2019.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- Han Yu, Zelei Liu, Yang Liu, Tianjian Chen, Mingshu Cong, Xi Weng, Dusit Niyato, and Qiang Yang. A fairness-aware incentive scheme for federated learning. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pp. 393–399, 2020.