A Sustainable Incentive Scheme for Federated Learning

Han Yu
Nanyang Technological University

Zelei Liu
Nanyang Technological University

Yang Liu
WeBank

Tianjian Chen
WeBank

Mingshu Cong
University of Hong Kong

Xi Weng
Peking University

Dusit Niyato
Nanyang Technological University

Qiang Yang
WeBank, Hong Kong University of Science and Technology

***Abstract*—In federated learning (FL), a federation distributedly trains a collective machine learning model by leveraging privacy preserving technologies. However, FL participants need to incur some cost for contributing to the FL models. The training and commercialization of the models will take time. Thus, there will be delays before the federation could pay back the participants. This temporary mismatch between contributions and rewards has not been accounted for by existing payoff-sharing schemes. To address this limitation, we propose the FL incentivizer (FLI). It dynamically divides a given budget in a context-aware manner among data owners in a federation by jointly maximizing the collective utility while minimizing the inequality among the data owners, in terms of the payoff received and the waiting time for receiving payoffs. Comparisons with five state-of-the-art payoff-sharing schemes show that FLI attracts high-quality data owners and achieves the highest expected revenue for a federation.**

Digital Object Identifier 10.1109/MIS.2020.2987774
Date of publication 16 April 2020; date of current version 3 September 2020.

IN TRADITIONAL MACHINE learning, the training dataset is usually stored in a central entity. Data need to be first collected from the data sources in order to facilitate learning. The rapid development of artificial intelligence (AI) has benefited from large-scale training data in real-world

1541-1672 © 2020 IEEE
Published by the IEEE Computer Society

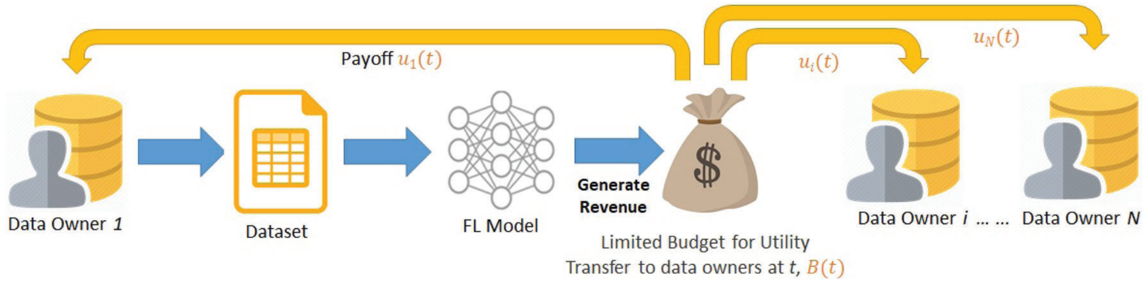

Figure 1. Transfer of utility under the FL settings.[3]

application. As AI becomes increasingly ubiquitous, nations are increasingly concerned about AI governance and privacy protection. They instituted new legislations such as the General Data Protection Regulation for these purposes. These new laws can potentially limit the development of AI in the long run. Federated learning (FL)[1] was proposed to enable AI to continue developing in this new regulatory landscape.

FL focuses on data integration methods, which comply with privacy and security laws. Under FL, *data owners* employ privacy protection techniques such as homomorphic encryption, secret sharing, and differential privacy to contribute model parameters trained on their own datasets to a *federation*.[1] The federation then combines these local model parameters in order to train a more effective collectively machine learning model. This allows the learning process to leverage the computational power of the data sources to train the model in a process similar to crowdsourcing.[2] Recently, Google has released the Tensorflow Federated toolkit to support the development FL applications. The technology has been applied to improve Google's keyboard query suggestions by crowdsourcing data from millions of mobile phone users in a privacy-preserving manner.

For a federation, data owners' continued participation in the FL process (through sharing of encrypted model parameters) is key to its long-term success. In essence, federations are competing for data owners in order to build high-quality models. Existing FL platforms such as Tensorflow Federated and the Federated AI Technology Enabler (FATE) assume that the federation already has a readily available group of participating data owners and do not provide any incentive mechanism to motivate participation. Such an assumption may not hold in practice, especially when data owners are companies rather than individuals.

The contributions by data owners to a federation are used to build a machine learning model which, in turn, can be used to generate revenue. Thus, the federation can allocate part of the revenue to data owners as incentives (see Figure 1). The research question here is how to quantify the payoff for each data owner in order to achieve long-term systemic wellbeing. In order to address this problem, a payoff-sharing scheme developed specially for FL is needed.

In game theoretic research, a number of payoff-sharing schemes exist. Standard coalition games with transferable utility.[4] Auctions, lotteries, and reputation can also be used to incentivize participation.[5] Payoff-sharing games such as the Labour union game and the Shapley game share payoff among players according to their marginal contribution in a coalition, while the Fair-value game does so according to the marginal loss as a result of a player leaving the coalition.[6] However, in FL, participants need to incur some cost for contributing to the FL models with their local datasets. The training and commercialization of the models will take time. Thus, there will be some delays before the federation has enough budget to pay back the participants. This temporary mismatch between contributions and rewards has not been accounted for by existing payoff-sharing schemes.

In order to sustain long-term stability in a data federation and attract more high-quality data owners over time, a fair incentive mechanism suitable for the FL context is needed.[1] For this purpose, we propose a dynamic payoff-sharing

[*] https://www.tensorflow.org/federated
[†] https://github.com/WeBankFinTech/FATE

scheme—FL Incentivizer (FLI). It is a polynomial time algorithm that can compute solutions for payoff-sharing by instalment in order to achieve fair treatment among data owners. It dynamically divides a given budget among data owners in a federation by jointly maximizing the collective utility generated while minimizing the inequality among the data owners in terms of the payoff and the waiting time for receiving the full payoff. Once the cost incurred by a data owner is fully compensated, FLI continues to pay the data owner following the baseline payoff-sharing scheme adopted by the federation (which are explained in more detailed in the "Related Work" section).

We make the following contributions in this article:

1) We model and describe the problem of motivating participation by high-quality data owners with incentives in the context of FL.
2) We provide a real-time algorithm to jointly achieve three fairness criteria 1) contribution fairness, 2) regret distribution fairness, and 3) expectation fairness), which are important to FL; and account for the interest of both the federation and the participating data owners.
3) We show the performance bounds of the proposed incentive scheme through theoretical analysis to show that FLI can produce near-optimal collective utility while limiting data owners' regret.
4) Extensive experimental comparisons with five existing payoff-sharing schemes show that FLI is the most attractive to high-quality data owners and least attractive to low-quality data owners, and achieves the highest expected revenue, thereby sustaining the long-term well being of a data federation.

To the best of our knowledge, this article is the first to study the issue of motivating continued participation by data owners in FL through dynamic payoff sharing. It provides a framework for incentive mechanism designers to sustain participation by data owners in FL to empower privacy respecting AI of the future.

The remaining parts of this article are organized as follows. The "Related Work" section reviews related work in profit-sharing and provides a detailed comparison of the relative advantages of FLI over existing approaches. The "FLI Payoff-Sharing Scheme" section introduces the system model of FL and explains the proposed FLI scheme in detail. The "Analytical Evaluation" section establishes the performance bounds of FLI through theoretical analysis. The "Experimental Evaluation" section introduces the experiment settings and interprets the experimental results. Finally, in the "Conclusion and Future Work" section, we conclude the article and discuss potential future research directions.

RELATED WORK

The problem studied in this article is related to the field of distributed welfare games. In a typical distributed welfare game, each player can select a subset of resources to generate welfare. The resulting welfare may depend on the subset of players who chose this resource, and the welfare generated at each resource is distributed among players who select it at the same time. Research in this field mainly focused on designing efficient schemes to fairly form coalitions by players in a distributed manner to reach (approximate) Nash equilibria. Similar research problems have also been studied under the topic of cost-sharing games. Most existing works in this field investigate how to design cost-sharing mechanisms in the context of congestion games in order to achieve efficient resource utilization.[7]

The research most closely related to our problem comes from the topic of profit-sharing games. In general, there are three categories of widely used profit-sharing schemes:

1) *Egalitarian:* any unit of utility produced by a data federation is divided equally among the data owners who help produce it.
2) *Marginal gain:* the payoff of a data owner in a data federation is the utility that the team gained when the data owner joined.
3) *Marginal loss:* the payoff of a data owner in a data federation is the utility that the team would lose if the data owner were to leave.

In general, a participant i's share of payoff from a total budget $B(t)$ in a given round of profit-sharing t, denoted as $\hat{u}_i(t)$, is computed as

$$\hat{u}_i(t) = \frac{u_i(t)}{\sum_{i=1}^{N} u_i(t)} B(t) \tag{1}$$

where $u_i(t)$ is the i's share of $B(t)$ among the peers computed following a given scheme.

Equal division is an example of egalitarian profit-sharing.[8] Under this scheme, the available profit-sharing budget $B(t)$ at a given round t is equally divided among all N participants. Thus, a participant i's payoff is

$$u_i(t) = \frac{1}{N}. \tag{2}$$

Under the individual profit-sharing scheme,[8] each participant i's own contribution to the collective (assuming the collective only contains i) is used to determine his share of the profit $u_i(t)$:

$$u_i(t) = v(\{i\}) \tag{3}$$

where $v(X)$ denotes the utility of a collective X.

The Labour Union game[6] profit-sharing scheme determines i's share of $B(t)$ based on his marginal contribution to the utility of the collective formed by his predecessors F (i.e., each participant's marginal contribution is computed based on the same sequence as they joined the collective)

$$u_i(t) = v(F \cup \{i\}) - v(F). \tag{4}$$

The Shapley game profit-sharing scheme[6] is also a marginal contribution-based scheme. Unlike the Labour Union game, Shapley game aims to eliminate the effect of the participants joining the collective in different sequences in order to more fairly estimate their marginal contributions to the collective. Thus, it averages the marginal contribution for each i under all different permutations of the i joining the collective relative to other participants:

$$u_i(t) = \sum_{P \subseteq P_j \setminus \{i\}} \frac{|P|!(|P_j| - |P| - 1)!}{|P_j|} [v(P \cup \{i\}) - v(P)] \tag{5}$$

where a collective is divided into m parties $(P_1, P_2, \ldots, P_m)$. Jia *et al.*[9] computed a Shapley value to split rewards among data owners. Such computations tend to be expensive.

For gradient-based FL approaches, the gradient information can be regarded as a type of data. However, in these cases, output agreement based rewards are hard to apply as mutual information requires a multitask setting, which is usually not present in such cases. Thus, among these three categories of schemes, model improvement is the most relevant way of designing rewards for FL. There are two emerging FL incentive schemes focused on model improvement.

A scheme which pays for marginal improvements brought about by model updates was proposed by Richardson *et al.*[10] The sum of improvements might result in overestimation of contribution. Thus, the proposed approach also includes a model for correcting the overestimation issue. This scheme ensures that payment is proportional to model quality improvement, which means the budget for achieving a target model quality level is predictable. It also ensures that data owners who submit model updates early receive a higher reward. This motivates them to participate even in early stages of the federated model training process. Similar to Richardson *et al.*,[10,9] Jia *et al.* computes a Shapley value to split reward among data owners.

These schemes are useful as baseline approaches to help a federation evaluate the contribution from a data owner. However, none of them accounts for the fairness of distributing profit over time with multiple contributions to a federation.

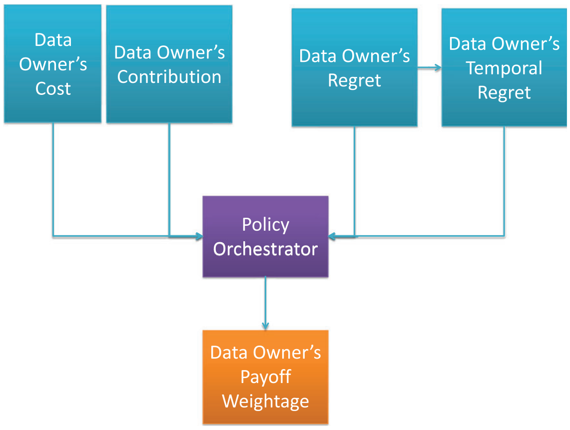

Figure 2. Overview of FLI.

FLI PAYOFF-SHARING SCHEME

In this section, we introduce the FL system model and derive the FLI payoff-sharing scheme. We explain each of the modules in the structure of FLI as shown in Figure 2.

Modeling Contribution

We assume that the data federation follows synchronous mode of model training commonly adopted by FL in which data owners share their model parameters in rounds. In round t, a data owner i can contribute his local model trained on a dataset to a federation. The federation is able to assess the contribution of i's data contribution to the federation following one of the profit-sharing schemes discussed in the previous section as the FLI baseline scheme.

To do so, a federation can run a sandbox simulation to estimate the effect of a data owner's contribution on model performance. The outcome is recorded by a variable $q_i(t) \geqslant 0$, which denotes the expected marginal revenue the federated model can gain with i's latest contribution. FLI is fully decoupled from how such a contribution score is produced. Thus, we do not focus on the exact mechanism by which $q_i(t)$ is produced, and treat it as an input for FLI.

Modeling Cost

Let $c_i(t)$ be the cost for i to contribute the local dataset, with size $d_i(t)$ and quality measure $\rho_i \in [0, 1]$, to the federation. There can be multiple ways to compute $c_i(t)$. Although it is possible to build computational models based on market research, a more practical solution is still auction-based self-report. A procurement auction can be used to estimate the cost when $c_i(t)$ is privately known. Specifically, the federation can ask each data owner to request a payment for the data contribution, and then select which data owner shall join the federation.

In this case, the delayed payment scheme can be separated from the procurement auction, where $c_i(t)$ can be interpreted as the payment to data owner i determined by the auction. This way, a clear separation of concern between the auction stage and the proposed incentive scheme can be achieved. Since this article focuses on developing the framework of incentive design for FL, we leave the topic of computing $c_i(t)$ to be treated in another work, and assume that this value is available here.

Modeling Regret

For each data owner i, the federation keeps track of the payoff gained from contributing data to the federation over time. As this value represents the difference between what the data owner has received so far and what he is supposed to receive, we refer to this term as *regret* $Y_i(t)$. The dynamics of $Y_i(t)$ can be regarded as a queueing system

$$Y_i(t+1) \triangleq \max[Y_i(t) + c_i(t) - u_i(t), 0] \tag{6}$$

where $u_i(t)$ is the payoff to be transferred to i by the federation. A large value of $Y_i(t)$ indicates that i has not been adequately compensated.

Modeling Temporal Regret

In some cases, the cost $c_i(t)$ may be too large to be fully covered by a single payment of $u_i(t)$ due to budget limitation in the federation. In such cases, the federation needs to compute instalments to be paid out to the data owners in multiple rounds. Their share of the current payout budget $B(t)$ depends on their regret as well as how long they have been waiting to receive the full payoff. We need to take into account how long a data contributor i has been waiting to be fully compensated through transfers of utility by instalment, when apportioning how big a share of the current $B(t)$ should be allocated to i.

For this purpose, we complement (6) with a *temporal queue* $Q_i(t)$ with queueing dynamics defined as

$$Q_i(t+1) \triangleq \max[Q_i(t) + \lambda_i(t) - u_i(t), 0] \tag{7}$$

where $\lambda_i(t)$ is an indicator function

$$\lambda_i(t) = \begin{cases} \hat{c}_i, & \text{if } Y_i(t) > 0 \\ 0, & \text{otherwise.} \end{cases} \tag{8}$$

This formulation means that as long as $Y_i(t)$ is not empty, the temporal queue $Q_i(t)$ will increase. The increment is based on i's average cost of data contribution to the federation $\hat{c}_i$ through past experience. Both queues decrease by the same amount when the federation pays i.

Equation (7) can be reexpressed as

$$Q_i(t+1) \geqslant Q_i(t) + \lambda_i(t) - u_i(t). \tag{9}$$

By rearranging the above inequality, we have

$$Q_i(t+1) - Q_i(t) \geqslant \lambda_i(t) - u_i(t). \tag{10}$$

Summing both sides of the above inequality over all $t \in \{0, ..., T-1\}$ yields

$$\sum_{t=0}^{T-1}[Q_i(t+1) - Q_i(t)] \geqslant \sum_{t=0}^{T-1}[\lambda_i(t) - u_i(t)]. \quad (11)$$

Thus, we have

$$Q_i(T) - Q_i(0) \geqslant \sum_{t=0}^{T-1}[\lambda_i(t) - u_i(t)]. \quad (12)$$

Since $Q_i(0) = 0$, the above inequality is simplified as

$$\frac{Q_i(T)}{T} \geqslant \frac{1}{T}\sum_{t=0}^{T-1}\lambda_i(t) - \frac{1}{T}\sum_{t=0}^{T-1}u_i(t). \quad (13)$$

Based on (13), by ensuring that the computed $u_i(t)$ values satisfy the queueing stability requirement of $\frac{1}{T}\sum_{t=0}^{T-1}u_i(t) \geqslant \frac{1}{T}\sum_{t=0}^{T-1}\lambda_i(t)$ for the temporal queue, the profit-sharing approach can ensure that data owners are compensated not only for their data contributions, but also for waiting to receive the full payoff, thereby making it "worth their while" to attract them to the federation.

Policy Orchestrator

In order to encourage data owners to continue participating in the federation, the federation needs to ensure that the data owners are treated fairly based on their individual contribution. Here, we define three fairness criteria that are important to the long-term sustainable operation of a federation:

1) *Contribution Fairness:* a data owner i's payoff shall be positively related to his contribution $q_i(t)$.
2) *Regret Distribution Fairness:* the difference of the regret and the temporal regret among data owners shall be minimized.
3) *Expectation Fairness:* the fluctuation of data owners' regret and temporal regret values shall be minimized.

In order to satisfy all the three fairness criteria, the federation shall maximize a "value-minus-regret drift" objective function over time. The collective utility derived from data owners' contributions is related to two factors: 1) the contribution to the federation by a data owner i ($q_i(t)$) and 2) the payoff that i receives from the federation for the contribution ($u_i(t)$). It is fair that a data owner who makes significant contributions to the federation shall receive a high payoff

$$U = \frac{1}{T}\sum_{t=0}^{T-1}\sum_{i=1}^{N}\{q_i(t)u_i(t)\}. \quad (14)$$

Maximizing U satisfies *Fairness criterion (1)*.

Since $Y_i(0) = 0$ for all i, if we consistently strive to minimize the variation in $Y_i(t)$ over time, the regret must not grow unbounded to drive data owners away. A federation needs to jointly consider the magnitude and distribution of regret among data owners and over time in order to treat them fairly. l_2-norm can capture simultaneously the magnitudes of the regret values and the distribution of regret among data owners. A large l_2-norm value means there are many data owners with non-zero regrets, and/or there are a few data owners with very large regret. Both shall be minimized.

Based on the l_2-norm technique, we formulate the Lyapunov function[2] of FLI as

$$L(t) = \frac{1}{2}\sum_{i=1}^{N}[Y_i^2(t) + Q_i^2(t)]. \quad (15)$$

For the simplicity of derivation later, we omit the $\sqrt{\cdot}$ operator in the standard l_2-norm calculation and multiply the whole term with $\frac{1}{2}$. These changes do not alter the desirable properties of l_2-norm for our formulation.

The drift in data owners' regret over time is

$$\begin{aligned}
\triangle &= \frac{1}{T}\sum_{t=0}^{T-1}[L(t+1) - L(t)] \\
&= \frac{1}{T}\sum_{t=0}^{T-1}\sum_{i=1}^{N}\left[\frac{1}{2}Y_i^2(t+1) - \frac{1}{2}Y_i^2(t)\right. \\
&\qquad \left. + \frac{1}{2}Q_i^2(t+1) - \frac{1}{2}Q_i^2(t)\right] \\
&\leqslant \frac{1}{T}\sum_{t=0}^{T-1}\sum_{i=1}^{N}\left[\frac{1}{2}(Y_i(t) + c_i(t) - u_i(t))^2\right. \\
&\quad \left. - \frac{1}{2}Y_i^2(t) + \frac{1}{2}(Q_i(t) + \lambda_i(t) - u_i(t))^2 - \frac{1}{2}Q_i^2(t)\right] \\
&\leqslant \frac{1}{T}\sum_{t=0}^{T-1}\sum_{i=1}^{N}\left[Y_i(t)c_i(t) - Y_i(t)u_i(t) + \frac{1}{2}c_i^2(t)\right. \\
&\quad - c_i(t)u_i(t) + \frac{1}{2}u_i^2(t) + Q_i(t)\lambda_i(t) - Q_i(t)u_i(t) \\
&\quad \left. + \frac{1}{2}\lambda_i^2(t) - \lambda_i(t)u_i(t) + \frac{1}{2}u_i^2(t)\right].
\end{aligned} \quad (16)$$

Since $u_i(t)$ is the control variable here, we extract only terms containing it from (16)

$$\triangle \leqslant \frac{1}{T}\sum_{t=0}^{T-1}\sum_{i=1}^{N}\{u_i^2(t) - u_i(t)[Y_i(t) + c_i(t) + Q_i(t) + \lambda_i(t)]\}. \tag{17}$$

The regret drift variable $\triangle$ jointly captures the distribution of regret (both $Y_i(t)$ and $Q_i(t)$) among data owners, as well as the fluctuation of regret over time. Minimizing $\triangle$ satisfies *Fairness criteria (2) and (3)*.

By jointly considering collective utility and the distribution of regret, the overall objective function for a given federation can be defined as "maximize the collective utility, to address fairness criteria (1), while minimizing the magnitude and distribution of regret over time, to address fairness criteria (2) and (3):"

$$\omega U - \triangle \tag{18}$$

which shall be maximized. Here, ω is a regularization term for a federation to control the tradeoff between the two objectives. Thus, the objective function of a federation is

Maximize:

$$\frac{1}{T}\sum_{t=0}^{T-1}\sum_{i=1}^{N}\{u_i(t)[\omega q_i(t) + Y_i(t) + c_i(t) + Q_i(t) + \lambda_i(t)] - u_i^2(t)\} \tag{19}$$

Subject to:

$$\sum_{i=1}^{N}\hat{u}_i(t) \leqslant B(t) \qquad \forall t \tag{20}$$

$$\hat{u}_i(t) \geqslant 0 \qquad \forall i, t \tag{21}$$

where $\hat{u}_i(t) \leqslant u_i(t)$ denotes the actual instalment payout from the federation to a data owner i in round t, which will be derived in the following section.

Computing Payoff Weightage

In order to optimize (19), we set its first derivative to 0 and solve for $u_i(t)$

$$\frac{d}{du_i(t)}[\omega U - \triangle] = 0. \tag{22}$$

Solving the above equation yields

$$u_i(t) = \frac{1}{2}[\omega q_i(t) + Y_i(t) + c_i(t) + Q_i(t) + \lambda_i(t)]. \tag{23}$$

The second derivative of (19) is

$$\frac{d^2}{du_i^2(t)}[\omega U - \triangle] = -1 < 0. \tag{24}$$

Thus, the solution maximizes the objective function.

Algorithm 1. Federated Learning Incentivizer (FLI)

Input: ω and $B(t)$ set by the system administrator; $Y_i(t)$ from all data owners at round t (with $Y_i(t) = 0$ for any i who just joined the federation); and $Q_i(t)$ from all data owners at round t (with $Q_i(t) = 0$ for any i who just joined the federation).

```
Initialize S(t) ← 0; //to hold the sum of all u_i(t) values.
for i = 1 to N do
  if d_i(t) > 0 then
    Compute c_i(t);
    Compute q_i(t);
  else
    c_i(t) = 0;
  end if
  u_i(t) ← (1/2)[ωq_i(t) + Y_i(t) + c_i(t) + Q_i(t) + λ_i(t)];
  S(t) ← S(t) + u_i(t);
end for
for i = 1 to N do
  û_i(t) ← (u_i(t)/S(t)) B(t)
  Y_i(t+1) ← max[0, Y_i(t) + c_i(t) − û_i(t)];
  Q_i(t+1) ← max[0, Q_i(t) + λ_i(t) − û_i(t)];
end for
return{û_1(t), û_2(t), ..., û_N(t)};
```

For contributing $d_i(t)$ amount of data of quality $q_i(t)$ at round t, the data owner i shall receive a total compensation of $u_i(t) = \frac{1}{2}[\omega q_i(t) + Y_i(t) + c_i(t) + Q_i(t) + \lambda_i(t)]$. The federation may need to pay out this in instalments over a period of time if not enough budget $B(t)$ is available to pay all data owners fully at round t. To share $B(t)$ among the data owners, the computed $u_i(t)$ values are used as weights to divide the budget $B(t)$. The actual payout instalment to i at t $\hat{u}_i(t)$ is

$$\hat{u}_i(t) = \frac{u_i(t)}{\sum_{i=1}^{N} u_i(t)} B(t). \tag{25}$$

The FLI payoff-sharing scheme is summarized in Algorithm 1. It accounts for both the magnitude and the temporal aspects of participating in a federation. Data owners who has contributed a large set of high quality data, and who has not been fully compensated for a long time will enjoy a higher share of subsequent revenues generated by the federation. The federation uses such a mechanism to ensure that the data owners' interests are well taken care of and in a timely manner.

The computational time complexity of Algorithm 1 is $O(N)$. Once $Y_i(t)$ and $Q_i(t)$ both reach 0 with no new cost incurred by i, $u_i(t) = \omega q_i(t)$. From then on, i will share future payoffs based on his contribution to the federation assessed using one of the baseline methods (e.g., the Shapley game payoff-sharing scheme). FLI prioritizes compensating the data owners with nonzero regret while taking into account their contributions.

ANALYTICAL EVALUATION

In this section, we analyze the performance bounds of FLI.

Collective Utility

Let $U^*(t)$ be the theoretical optimal collective utility achieved by a federation at round t following an Oracle payoff-sharing scheme. It is possible find positive values ω, ϵ, and δ such that

$$\omega U(t) - \Delta(t) \geqslant \omega U^*(\tau) + \epsilon \sum_{i=1}^{N} u_i(t) - \delta. \tag{26}$$

Summing both sides of the above inequality over $t \in \{0, 1, \ldots, T-1\}$, we have

$$\begin{aligned} &\omega \sum_{t=0}^{T-1} U(t) - \sum_{t=0}^{T-1} [L(t+1) - L(t)] \\ &\quad = \omega \sum_{t=0}^{T-1} U(t) - L(T) + L(0) \geqslant \omega \sum_{t=0}^{T-1} U^*(t) \\ &\qquad + \epsilon \sum_{t=0}^{T-1} \sum_{i=1}^{N} u_i(t) - T\delta. \end{aligned} \tag{27}$$

Since $L(0) = 0$, $L(T) \geqslant 0$ and $\epsilon \sum_{t=0}^{T-1} \sum_{i=1}^{N} u_i (t) \geqslant 0$, the above inequality can be reexpressed as

$$\begin{aligned} \omega \sum_{t=0}^{T-1} U(t) &\geqslant \omega \sum_{t=0}^{T-1} U^*(t) + \epsilon \sum_{t=0}^{T-1} \sum_{i=1}^{N} u_i(t) \\ &\quad + L(T) - T\delta \geqslant \omega \sum_{t=0}^{T-1} U^*(t) - T\delta. \end{aligned} \tag{28}$$

Taking the time-average of this inequality yields the *lower bound* on the long-term time-averaged collective utility

$$\lim_{T\to\infty} \inf \frac{1}{T} \sum_{t=0}^{T-1} U(t) \geqslant \frac{1}{T} \sum_{t=0}^{T-1} U^*(t) - \frac{\delta}{\omega}. \tag{29}$$

The result shows that, by following FLI, a federation produces long-term time-averaged collective utility within $O(\frac{1}{\omega})$ of the theoretical optimal time-averaged collective utility. By increasing the value of ω, a federation can achieve the time-averaged collective utility closer to the optimal time-averaged collective utility. Note that this collective utility is different from the social welfare. It expresses the desire to reward data owners fairly based on their contributions.

Total Payout by the Federation

Similarly, rearranging the terms in (26) yields

$$\begin{aligned} &\epsilon \sum_{t=0}^{T-1} \sum_{i=1}^{N} u_i(t) \\ &\leqslant \omega \sum_{t=0}^{T-1} U(t) + T\delta - \omega \sum_{t=0}^{T-1} U^*(t) - L(T) + L(0) \\ &\leqslant \omega \sum_{t=0}^{T-1} U(t) + T\delta. \end{aligned} \tag{30}$$

Taking the time-average of this inequality yields the *upper bound* on the long-term time-averaged total payout

$$\lim_{T\to\infty} \sup \frac{1}{T} \sum_{t=0}^{T-1} \sum_{i=1}^{N} u_i(t) \leqslant \frac{\omega}{T\epsilon} \sum_{t=0}^{T-1} U(t) + \frac{\delta}{\epsilon}. \tag{31}$$

The result shows that, by following FLI, the long-term time-averaged total payout by the federation to data owners is bounded by $O(\omega)$. Increasing the value of ω will signal the federation to pay out more of its revenue.

Regret

Recall that a data owner's regret when contributing data to a federation is defined by (6). Rearranging the terms and removing the $\max[\cdot, 0]$ operator in (6) yields

$$Y_i(t+1) - Y_i(t) \geqslant c_i(t) - u_i(t). \tag{32}$$

The long-term time-averaged change in i's regret is

$$\lim_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}[Y_i(t+1)-Y_i(t)] \geqslant \lim_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}[c_i(t)-u_i(t)]. \tag{33}$$

Based on (23), we have

$$\begin{aligned}\lim_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}u_i(t) &= \lim_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}[\omega q_i(t)+Y_i(t)+c_i(t)]\\ &\geqslant \lim_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}c_i(t).\end{aligned} \tag{34}$$

Therefore, the *lower bound* of a data owner i's long-term time-averaged regret is

$$\lim_{T\to\infty}\inf\frac{1}{T}Y_i(T)\geqslant\frac{1}{T}\sum_{t=0}^{T-1}[c_i(t)-u_i(t)] \tag{35}$$

where $\frac{1}{T}\sum_{t=0}^{T-1}[c_i(t)-u_i(t)]\leqslant 0$. Since

$$\begin{aligned}\epsilon\sum_{t=0}^{T-1}\sum_{i=1}^{N}u_i(t) &= \epsilon\sum_{t=0}^{T-1}\sum_{i=1}^{N}[\omega q_i(t)+Y_i(t)+c_i(t)]\\ &\leqslant \omega\sum_{t=0}^{T-1}U(t)+T\delta-\omega\sum_{t=0}^{T-1}U^*(t)\\ &\quad -L(T)+L(0).\end{aligned} \tag{36}$$

We have

$$\begin{aligned}\epsilon\sum_{t=0}^{T-1}\sum_{i=1}^{N}Y_i(t) &\leqslant \omega\sum_{t=0}^{T-1}U(t)+T\delta-\omega\sum_{t=0}^{T-1}U^*(t)\\ &\quad -L(T)+L(0)-\epsilon\sum_{t=0}^{T-1}\sum_{i=1}^{N}[\omega q_i(t)+c_i(t)]\\ &= T\delta+\left[\omega\sum_{t=0}^{T-1}U(t)-\omega\sum_{t=0}^{T-1}U^*(t)\right]\\ &\quad -L(T)-\epsilon\sum_{t=0}^{T-1}\sum_{i=1}^{N}[\omega q_i(t)+c_i(t)].\end{aligned} \tag{37}$$

As $\left[\omega\sum_{t=0}^{T-1}U(t)-\omega\sum_{t=0}^{T-1}U^*(t)\right]\leqslant 0$, the above inequality can be rewritten as

$$\epsilon\sum_{t=0}^{T-1}\sum_{i=1}^{N}Y_i(t)=\epsilon\sum_{i=1}^{N}Y_i(T)\leqslant T\delta. \tag{38}$$

Thus, the *upper bound* of the time-averaged regret is

$$\lim_{T\to\infty}\sup\frac{1}{T}\sum_{t=0}^{T-1}\sum_{i=1}^{N}Y_i(t)\leqslant\frac{\delta}{\epsilon}. \tag{39}$$

The result shows that, as long as the federation ensures that the condition of $\frac{1}{T}\sum_{t=0}^{T-1}[u_i(t)-c_i(t)]\geqslant 0$ holds, the long-term time-averaged regret for data owners is upper-bounded by a constant value and will not grow indefinitely

$$0\leqslant\lim_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}\sum_{i=1}^{N}Y_i(t)\leqslant\frac{\delta}{\epsilon}. \tag{40}$$

Thus, we show that FLI is stable in terms of data owners' regret regardless of the choice of value for ω since $u_i(t)\geqslant c_i(t),\forall i,t$ according to (23).

EXPERIMENTAL EVALUATION

To complement the analytical results and evaluate the performance of FLI under FL settings, we build a multiagent simulator which creates data owner agents with diverse characteristics and supports multiple payoff-sharing schemes. The relative performance of the schemes in the simulation is more important than the exact values.

Experiment Settings

In the experiments, there are seven federations each adopting one of the following six payoff-sharing schemes for data owner agents to join:

1) *Linear:* a data owner i's share of $B(t)$ is proportional to the total quantity of data contribution weighted by its data quality (this is a baseline payoff-sharing scheme that we designed for experimental comparison purposes only).
2) *Equal:* $B(t)$ is equally divided among data owners in this federation.[8]
3) *Individual:* i's share of $B(t)$ is proportional to his marginal contribution to the revenue of the federation.[8]
4) *Union:* i's share of $B(t)$ follows the Labour Union game[6] scheme and is proportional to his marginal contribution to the revenue of the federation formed by his predecessors, $v(F\cup\{i\})-v(F)$.
5) *Shapley:* i's share of $B(t)$ follows the Shapley value-based scheme proposed in.[9]
6) *FLI:* data owners receive payoff according to FLI (with $\omega=1$).

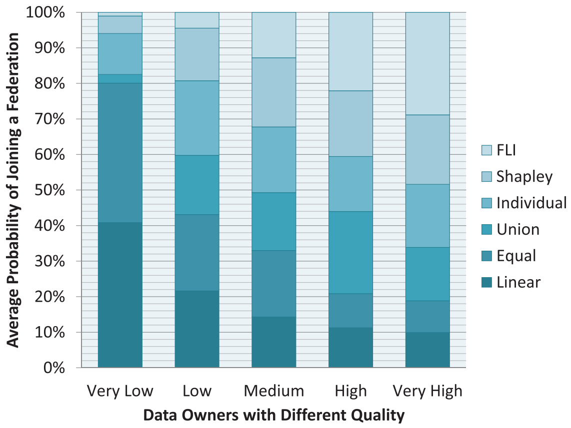

Figure 3. Data owner agents' final probability of following each scheme.

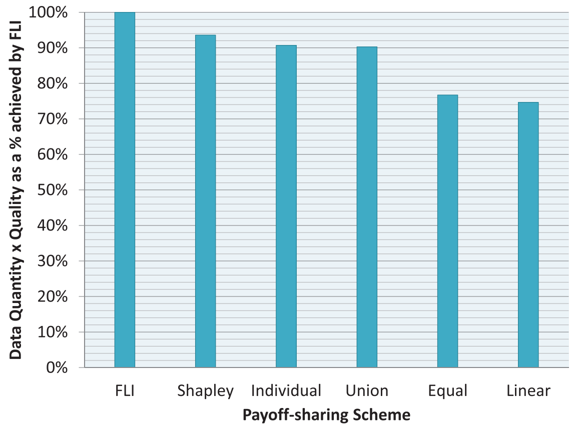

Figure 4. $\frac{1}{NT}\sum_i \sum_t d_i(t)\rho_i$ as a percentage of that achieved by FLI.

We follow the decreasing marginal utility assumption to map data quality and quantity to revenue generated by a federation. The revenue function used in the experiment is $\log$-linearly related to the product between the quality and quantity of data contributed to it. It is of the form $\log(1+\sum_i \sum_t d_i(t)\rho_i)$. In essence, federations are competing for data owner agents. The performance of the schemes is evaluated by the percentage and type of data owner agents who eventually choose to join each of them and the revenue generated. Each round of simulation consists of 1 000 epochs, and we repeat the simulation for 10 rounds with reinitialization to smoothen the effect of randomness.

We simulate 100 self-interested data owner agents in the experiment, each representing a company. Their ρ_i values are randomly initialized following a uniform distribution between 0 and 1 at the beginning of each round of experiment. In each epoch, each agent decides on which federation to join based on the cumulative payoffs received from each federation so far. The probability of an agent joining a federation at t equals to the cumulative payoff it received from this federation divided by total payoff received from all federations. Each agent joins federations following the ϵ-greedy approach, with equal starting probability for all federations.

Results and Discussion

Figure 3 shows different data owner agents' final probability of following each scheme. Data owner agents are divided into five types based on their individual ρ_i values. Agents with ρ_i values belonging to the range of [0,0.2), [0.2,0.4), [0.4,0.6), [0.6,0.8), and [0.8,1] are labeled as "Very Low," "Low," "Medium," "High," and "Very High" types, respectively. It can be observed that the medium, low, and very low type agents have the smallest probabilities of joining the federation adopting FLI. Shapley and Union follow are similar trend, but are less attractive to very high and high type agents compared to FLI. Individual, linear, and equal are more attractive to agents with lower data quality than those with higher data quality.

Figure 4 shows the total quantity of data received by each federation weighted by data quality ($\frac{1}{NT}\sum_i \sum_t d_i(t)\rho_i$). FLI achieves the best performance, outperforming the second best scheme, Shapley, by around 7%. Individual and

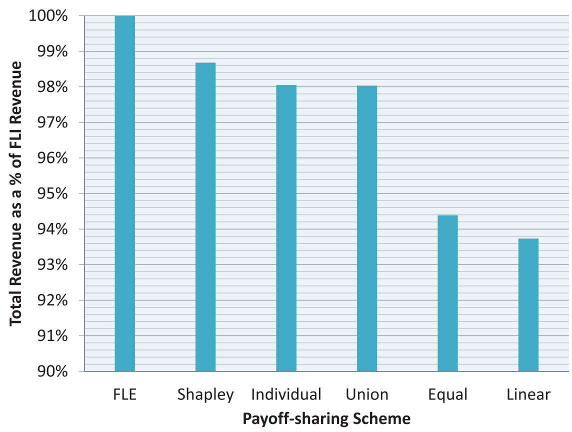

Figure 5. Revenue as a percentage of that achieved by FLI.

Union perform similarly, which is around 10% lower than that achieved by FLI. The performance of equal is more than 30% lower than that of FLI. Linear fares the worst, underperforming FLI by over 34%.

Figure 5 shows the total revenue as a result of the data received by each federation. Similar to the trend shown in Figure 4, FLI achieves the highest revenue, followed by Shapley, individual, union, equal, and linear. As long as the adopted revenue function monotonically increases with data quality and quantity, this performance trend holds.

CONCLUSION AND FUTURE WORK

In this article, we proposed the FLI payoff-sharing scheme incentivize FL data owners to contribute high-quality data to the data federation. Data owners who has contributed a large set of high-quality data, and who has not been fully compensated for a long time, will enjoy a higher share of subsequent revenues generated by the federation. A federation following FLI is able to dynamically adjust data owners' shares in order to fairly distribute benefits and sacrifice among them. Analytical evaluation established that FLI is able to produce near-optimal collective utility while limiting data owners' regret. With FLI accounting for the temporary mismatch between contributions and rewards due to the limitations of FL, thereby enabling a healthy FL ecosystem to emerge over time. To the best of our knowledge, FLI is the first incentive mechanism designed for FL. It jointly considers factors important to FL, with clear separation of concerns with respect to the delay for the federated model to start generating revenue.

In subsequent research, we will extend the fairness criteria being considered to include the likes of proportional fairness for instance. We will also study alternative ways to quantify the impact ρ_i of an encrypted dataset contributed by a data owner on the well being of a data federation. One of the most challenging task is to estimate data owners' cost incurred for joining the federation. A federation can run a sandbox simulation to estimate the effect of a data owner's contribution on model performance. A well-designed sandbox should be able to simulate the change in revenue as a result of a data owner's contribution. In this way, the mechanism is fully decoupled from how such a contribution score is produced.

ACKNOWLEDGMENTS

This work was supported, in part, by the Nanyang Assistant Professorship (NAP); NTU-WeBank JRI (NWJ-2019-007); AISG-GC-2019-003; NRF-NRFI05-2019-0002; and NTU-SDU-CFAIR (NSC-2019-011).

REFERENCES

1. Q. Yang, Y. Liu, Y. Cheng, Y. Kang, T. Chen, and H. Yu, *Federated Learning*. San Rafael, CA, USA: Morgan & Claypool, 2019.
2. H. Yu, C. Miao, Y. Chen, S. Fauvel, X. Li, and V. R. Lesser, "Algorithmic management for improving collective productivity in crowdsourcing," *Sci. Rep.*, vol. 7, 2017, Art. no. 12541.
3. H. Yu *et al.*, "A fairness-aware incentive scheme for federated learning," in *Proc. AAAI/ACM Conf. AI, Ethics, Soc.*, 2020, pp. 393–399.
4. B. Faltings, G. Radanovic, and R. Brachman, *Game Theory for Data Science: Eliciting Truthful Information*. San Rafael, CA, USA: Morgan & Claypool, 2017.
5. T. Luo, S. S. Kanhere, J. Huang, S. K. Das, and F. Wu, "Sustainable incentives for mobile crowdsensing: Auctions, lotteries, and trust and reputation systems," *IEEE Commun. Mag.*, vol. 55, no. 3, pp. 68–74, Mar. 2017.
6. S. Gollapudi, K. Kollias, D. Panigrahi, and V. Pliatsika, "Profit sharing and efficiency in utility games," in *Proc. 25th Ann. Eur. Symp. Algorithms (ESA'17)*, 2017, pp. 1–16.
7. K. Kollias and T. Roughgarden, "Restoring pure equilibria to weighted congestion games," *ACM Trans. Econ. Comput.*, vol. 3, no. 4, pp. 21:1–21:24, 2015.
8. S. Yang, F. Wu, S. Tang, X. Gao, B. Yang, and G. Chen, "On designing data quality-aware truth estimation and surplus sharing method for mobile crowdsensing," *IEEE J. Sel. Areas Commun.*, vol. 35, no. 4, pp. 832–847, Apr. 2017.
9. R. Jia *et al.*, "Towards efficient data valuation based on the Shapley value," in *Proc. Mach. Learn. Res.*, 2019, pp. 1167–1176.
10. A. Richardson, A. Filos-Ratsikas, and B. Faltings, "Rewarding high-quality data via influence functions," *arXiv:1908.11598*, 2019.

Han Yu is currently a Nanyang assistant professor with the School of Computer Science and Engineering (SCSE), Nanyang Technological University (NTU). From 2015 to 2018, he held the prestigious Lee Kuan Yew Postdoctoral Fellowship with the Joint NTU-UBC Research Centre of Excellence in Active Living for the Elderly. His research focuses on the ethics of artificial intelligence and federated learning. He received the B.Eng. (Hons.) and Ph.D. degrees from the SCSE, NTU, Singapore, in 2007 and 2014, respectively. Contact him at han.yu@ntu.edu.sg.

Zelei Liu is currently an RSB Postdoctoral Fellow with the School of Computer Science and Engineering, Nanyang Technological University, Singapore. His research focuses on federated learning and multi-agent systems. He received the B.Sc. and Ph.D. degrees from the School of Computer Science and Technology, Jilin University, China, in 2011 and 2019, respectively. Contact him at zelei.liu@ntu.edu.sg.

Yang Liu is a senior researcher with the AI Department of WeBank, China. Her research interests include machine learning, federated learning, transfer learning, multi-agent systems, statistical mechanics, and applications of these technologies in the financial industry. Her research has been published in leading scientific journals such as ACM TIST and Nature. She received the B.Sc. degree from Tsinghua University in 2007, and the Ph.D. degree from Princeton University in 2012. Contact her at yangliu@webank.com.

Tianjian Chen is the deputy general manager of AI Department of WeBank, China. He is now responsible for building the Banking Intelligence Ecosystem based on Federated Learning Technology. Before joining WeBank, he was the chief architect of Baidu Finance, Principal Architect of Baidu. He has over 12 years of experience in large-scale distributed system design and enabling technology innovations in various application fields. Contact him at tobychen@webank.com.

Mingshu Cong is currently working in the FinTech and Blockchain Laboratory, Department of Computer Science, the University of Hong Kong. She is also the CEO of LogiOcean Technologies, a financial technology company providing AI solutions for financial professionals. Her research focuses on financial intelligence analysis and economic aspects of federated learning. She received the B.Sc. degree in physics and Ph.D. degree in finance from Peking University. She received a M.Sc. degree in financial mathematics from the Hong Kong University of Science and Technology. Contact her at miranda.cong@logiocean.com.

Xi Weng is a tenured professor with Guanghua School of Management, Peking University. Prior to that, he received his bachelor's and master's degrees from Peking University, in 2004 and 2006, respectively. His research focuses on microeconomic theory, in particular, game theory, information economics, and organizational economics. He received the Ph.D. degree in economics from University of Pennsylvania, Philadelphia, PA, USA, in 2011. Contact him at wengxi125@gsm.pku.edu.cn.

Dusit Niyato is a professor with the School of Computer Engineering since July 2008. His research interests include wireless communications, and networking. He has done significant research work his research areas and published over 60 top quality international conference and journal papers. He received the bachelor's degree in computer engineering from King Mongkut's Institute of Technology Ladkrabang, Thailand, and the master's and Ph.D. degrees from the University of Manitoba, Winnipeg, Canada, in 2005 and 2008, respectively. Contact him at dniyato@ntu.edu.sg.

Qiang Yang is the head of AI at WeBank (Chief AI Officer) and the chair professor with the Computer Science and Engineering, Department of the Hong Kong University of Science and Technology (HKUST). His research interests include artificial intelligence, machine learning, and data mining, especially in transfer learning, automated planning, federated learning, and case-based reasoning. He received the Ph.D. degree from the Computer Science Department in 1989 and the M.Sc. degree in astrophysics in 1985, both from the University of Maryland, College Park. He received the B.Sc. degree in astrophysics from Peking University, Beijing, China, in 1982. He is a fellow of several international societies, including ACM, AAAI, IEEE, IAPR, and AAAS. Contact him at qyang@ust.hk.