Decentralized and Incentivized Federated Learning: A Blockchain-Enabled Framework Utilising Compressed Soft-Labels and Peer Consistency

Leon Witt, Usama Zafar, KuoYeh Shen, Felix Sattler, Dan Li, Songtao Wang, and Wojciech Samek, *Member, IEEE*

***Abstract*—Federated Learning (FL) has emerged as a powerful paradigm in Artificial Intelligence, facilitating the parallel training of Artificial Neural Networks on edge devices while safeguarding data privacy. Nonetheless, to encourage widespread adoption, Federated Learning Frameworks (FLFs) must tackle (i) the power imbalance between a central authority and its participants, and (ii) the challenge of equitably measuring and incentivizing contributions. Existing approaches to decentralize and incentivize FL processes are hindered by (i) computational overhead and (ii) uncertainty in contribution assessment (Witt et al. 2023), limiting FL's scalability beyond use cases where trust between participants and the server is established. This work introduces a cutting-edge, blockchain-enabled federated learning framework that incorporates Federated Knowledge Distillation (FD) with compressed 1-bit soft-labels, aggregated through a smart contract. Furthermore, we present the Peer Truth Serum for Federated Distillation (PTSFD), which cultivates an incentive-compatible ecosystem by rewarding honest participation based on an implicit yet effective comparison of worker contributions. The primary innovation stems from its lightweight architecture that simultaneously promotes decentralization and incentivization, addressing critical challenges in contemporary FL approaches.**

Manuscript received 25 April 2023; revised 13 November 2023; accepted 13 November 2023. Date of publication 28 November 2023; date of current version 8 August 2024. This work was supported in part by the German Federal Ministry of Education and Research (BMBF) through the BIFOLD - Berlin Institute for the Foundations of Learning and Data under Grants 01IS18025A and 01IS18037I, and in part by the EU's Horizon 2020 project COPA EUROPE under Grant 957059, as well as in part by the National Nature Science Foundation of China under Grant U21B2022. *(Corresponding author: Wojciech Samek.)*

Leon Witt is with Tsinghua University, Beijing 100190, China, and also with the Department of Artificial Intelligence, Fraunhofer Heinrich Hertz Institute, 10587 Berlin, Germany (e-mail: leonmaximilianwitt@gmail.com).

Usama Zafar is with the Department of Information Technology, Uppsala University, 752 36 Uppsala, Sweden (e-mail: usama-zafar@outlook.com).

KuoYeh Shen and Dan Li are with the Tsinghua University, Beijing 100190, China (e-mail: kuoyehs@gmail.com; tolidan@tsinghua.edu.cn).

Felix Sattler is with the Department of Artificial Intelligence, Fraunhofer Heinrich Hertz Institute, 10587 Berlin, Germany (e-mail: felix.sattler@hhi.fraunhofer.de).

Songtao Wang is with Zhongguancun Laboratory, Beijing 100190, China (e-mail: kingplantwater@sina.com).

Wojciech Samek is with the Department of Electrical Engineering and Computer Science, Technical University of Berlin, 10587 Berlin, Germany, also with the Department of Artificial Intelligence, Fraunhofer Heinrich Hertz Institute, 10587 Berlin, Germany, and also with BIFOLD – Berlin Institute for the Foundations of Learning and Data, 10587 Berlin, Germany (e-mail: wojciech.samek@hhi.fraunhofer.de).

This article has supplementary downloadable material available at https://doi.org/10.1109/TSC.2023.3336980, provided by the authors.

Digital Object Identifier 10.1109/TSC.2023.3336980

***Index Terms*—Federated learning, blockchain, reward mechanism, federated distillation, decentralized machine learning.**

I. Introduction

THE ascent of *Machine Learning* (ML) has been marked by a growing emphasis on decentralized and privacy-preserving solutions. One of the leading solutions, *Federated Learning* (FL), allows training of *Deep Neural Networks* (NNs) across distributed devices, ensuring data remains localized, hence addressing privacy concerns. *Federated Averaging* (`FedAvg`) [2], a cornerstone algorithm in FL, achieves this by aggregating locally trained models to produce a global model. However, FL's transformative potential is curtailed by (i) A trust deficit emanating from the imbalance of power between workers and the central server and (ii) an absent practical reward mechanism to incentivize worker contributions. Notably, while blockchain's inherent transparency and immutability characteristics hold promise in addressing the trust issue, its effective integration with FL for scalable deployments has remained elusive [3], [4], [5]. Furthermore, designing mechanisms that effectively reward worker contributions without compromising data privacy remains open research in FL [1], [6]. Comparing and evaluating worker contributions in `FedAvg` is non-trivial since data always stays private [1], [7]. Traditional solutions like the Leave-one-out [8] or Shapley value [9], [10] introduce computational overhead and hinge on a centralized authority, constraining their adoption in decentralized, blockchain-based solutions. Lastly, the prohibitive cost of storing vast amounts of data on blockchain systems, compounded by the intensive computational demands, means popular methods like `FedAvg` struggle to fit within *General Purpose Blockchain Systems* (GPBS). This has prompted researchers towards *Application Specific Blockchain Systems* (ASBS) or off-chain aggregation [11], [12], [13], [14], [15] – both bringing their set of challenges.

A. Contributions

In response to these challenges, this work introduces the Peer Truth Serum for Federated Distillation (PTSFD), a blockchain-enabled and incentivized FL framework. It utilizes *Federated Knowledge Distillation* (FD) on 1-bit compressed soft-labels, combined with the Peer Truth Serum for Crowdsourcing [16], which is adjusted for the FD case.

© 2023 The Authors. This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

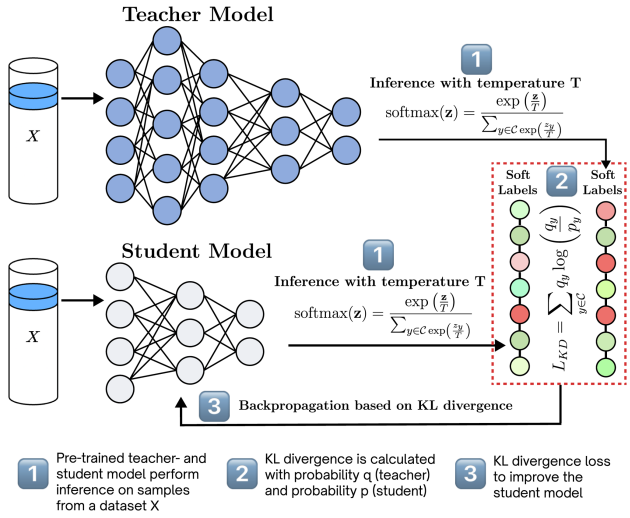

Fig. 1. Knowledge distillation.

1) *Incentivization:* We present the *Peer Truth Serum for Federated Distillation* (PTSFD), an informed-truthful multi-task peer prediction mechanism tailored for the FD case. It discerns contributions based on the correlation of reported 1-bit compressed soft-labels. This approach drastically reduces storage requirements, making it particularly suited for blockchain.
2) *Decentralization:* The reduced storage requirement and the simplicity of our method promote decentralization. The framework can be deployed on simple smart contracts hosted on prevalent blockchains, such as the Ethereum Virtual Machine [3], eliminating the need for specialized ASBS and simplifying the entire FL process.
3) *Scalability and Efficiency:* Our method, which builds upon Federated Knowledge Distillation, significantly reduces both communication overheads and blockchain storage needs. This sets the stage for large-scale, practical FL deployments without sacrificing efficiency or scalability.

We substantiate our contributions through theoretical validations and exhaustive experimental analyses. Our findings reveal a system that maintains a strong incentive-compatible equilibrium, demonstrating resilience against adversarial actions. Moreover, it showcases efficiency gains in storage and communication costs compared to `FedAvg` in various FL scenarios. The core of this work lies in its pioneering architecture, laying the groundwork for a lightweight, fully decentralized, incentivized, and efficient Federated Learning paradigm.

II. Background and Related Work

A. Federated Averaging

The most common algorithmic approach to FL problems is `FedAvg`, where the training process consists of an iteration of the following steps:

1) The central server selects a subset of clients $\mathcal{W}$, which participate in this training round.
2) The central server sends the current model θ to the selected clients.
3) The selected clients perform local training on their private data, leading to updated client models θ_i.
4) The updated models $\theta_i \ \forall i \in \mathcal{W}'$ are send back to the central server.
5) The central server aggregates the updated models to a new global model.

This training paradigm requires a two-way communication of the model θ (resp. θ_i) at every iteration, which can result in significant communication overhead for state-of-the-art NN models with hundreds of millions of parameters. To address this challenge, various approaches have been proposed, including pruning methods [17] and advanced compression techniques [18], [19], [20], [21], [22], [23], [24], [25]. However, despite these advances, the fundamental issue of scaling `FedAvg` to larger models remains, impeding the utilization of blockchain for storing or aggregating models [1].

B. Knowledge Distillation and Federated Settings

1) Knowledge Distillation: *Knowledge Distillation* (KD), depicted in Fig. 1, is a technique in deep learning where a smaller NN model (often called the "student") is trained to mimic the behavior of a larger, pre-trained model (referred to as the "teacher") [27]. This is accomplished not by transferring the model parameters directly, but rather by aligning the output distributions of both models. Traditional training methods involve training a model directly on ground-truth labels, using a cross-entropy loss that measures the discrepancy between the model's predictions and these true labels. In contrast, KD employs a divergence-based loss, such as the Kullback-Leibler (KL) divergence, to measure the difference between the student's predicted probabilities and those of the teacher model. This divergence provides insights into how closely the student is able to mimic the behavior of its teacher. A distinct feature of KD is the use of "softened" labels. In traditional classification tasks, hard labels are used, which unequivocally classify a data point into one category. However, the teacher model in KD provides "soft" labels in the form of probabilities, indicating the confidence levels across various categories. These probabilities can be further softened using a temperature parameter T to yield a smoother distribution, capturing the nuances of decision boundaries and offering richer guidance to the student model. This process allows the student to inherit not just the overt knowledge from the ground-truth labels but also the implicit, or "dark", knowledge embedded in the teacher model's predictions. Since only soft-labels are necessary to perform backpropagation, models with varying architectures can learn from the teacher. The appeal of KD lies in its ability to produce compact models with performance that closely mirrors that of much larger networks. These compact models are advantageous for deployment in resource-constrained environments, such as mobile devices or edge devices, without sacrificing much in terms of accuracy. With the above foundation in KD, we can delve deeper into its application in the FL setup, specifically focusing on Federated Distillation.

2) Federated Distillation: Drawing from the KD paradigm, Federated Distillation (FD) [28], [29], [30] extends the soft-label philosophy to a federated landscape, where the aggregated

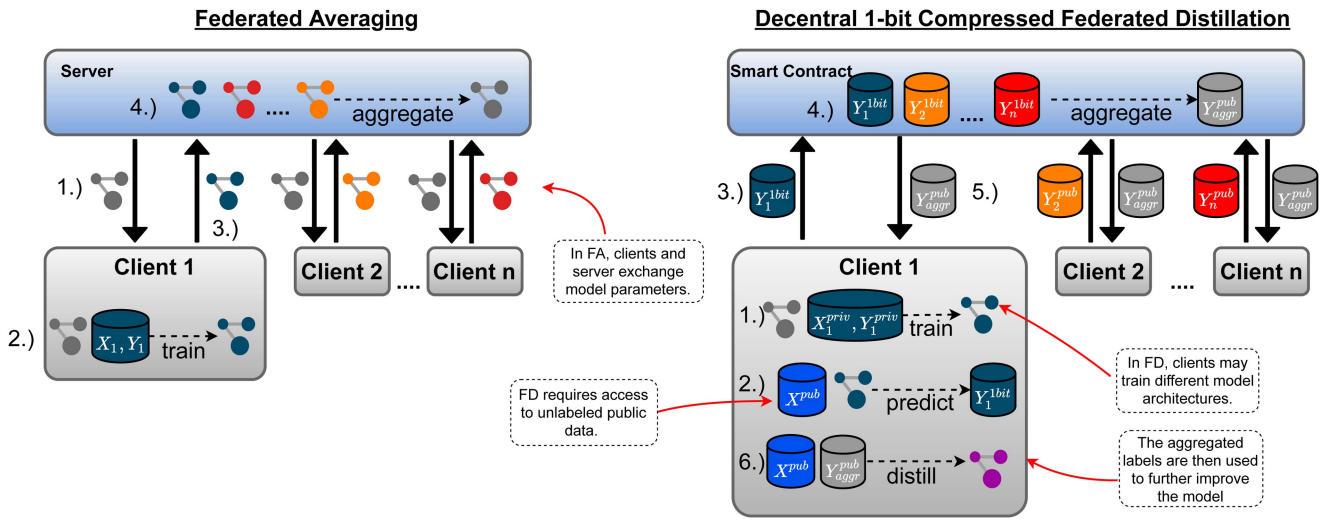

Fig. 2. Federated learning (from [26]) versus Decentral 1-bit compressed federated distillation on blockchain.

soft-logits from workers act akin to the soft-label output from an overarching teacher model. While FedAvg directly communicates the model parameters to transfer information between the central server and the clients, FD uses soft-label predictions Y_i^{pub} obtained on a separate public distillation dataset X^{pub} for this purpose. More precisely, the locally updated model $f_{\theta_i+\Delta\theta_i}$ is indirectly communicated to the central server by sending its predictions on the distillation dataset, i.e.,

$$Y_i^{pub} = \{f_{\theta_i+\Delta\theta_i}(x) \mid x \in X^{pub}\}. \tag{1}$$

Therefore, unlike traditional FL techniques such as FedAvg that mandate a consistent model structure across clients due to the aggregation of model parameter updates, FD does not require a single NN architecture but allows each worker to adopt a distinct architecture that might be best suited to its local data- or computational restrictions. Additionally, instead of with NN parameters, it scales with the size of the distillation dataset. This characteristic of FD can lead to communication savings [26], especially for large models. In this work, we modify a recently proposed, highly communication-efficient FD method [26], called *Compressed Federated Distillation (CFD)*, which is based on the multi-round protocol developed in [28], [30]. In our modified version of CFD, every client performs the following steps in each communication round, as depicted in Fig. 2:

1) *Train* on local datasets and improve model $\theta_i = \theta + \Delta\theta_i$ using X_i^{priv}, Y_i^{priv}.
2) *Predict labels* $Y_i^{1\,bit}$ using the improved model θ_i on X^{pub} to compute soft-labels Y_i^{pub} and perform 1-bit quantization $Y_i^{1\,bit} = Q^{1\,bit}(Y_i^{pub})$.
3) *Upload* the integer-encoded compressed soft-labels to the smart contract (in a two-step commit-reveal fashion outlined in Algorithm 3).
4) *(Blockchain) Aggregate* predictions Y_{aggr}^{pub} by majority vote over all Y_i^{pub}.
5) *Download* the aggregated predictions Y_{aggr}^{pub} from the blockchain.
6) *Distill* the current model θ using X^{pub} and Y_{aggr}^{pub}.

The authors of [26] showed that CFD largely reduces the information necessary for exchange by quantization Q and the use of a small public distillation dataset (e.g., random subset selection). The savings are in the order of two orders of magnitude when compared to Federated Distillation, and more than four orders of magnitude when compared to FedAvg. The possibility to apply binary soft-label quantization, i.e., Q_b with $b = 1$, ensure three important properties for a decentralized CFD on Blockchain, namely

- It reduces the amount of information processed in the aggregation process heavily.
- It makes contributions by workers explicit and comparable.
- It supersedes the need for additional encryption like noise inducing Differential Privacy or computational heavy secure multiparty computation.

C. Blockchain Technology in FL Context

Blockchain was initially introduced with Bitcoin by Satoshi Nakamoto in 2008 [31]. It is referred to as a distributed ledger managed by nodes in a peer-to-peer network, where cryptographic links of information ensure resistance to modification and immutability. The network is governed by a consensus mechanism [32] among peers, which supersedes the need for central coordination. The advent of general-purpose blockchains [3], with smart contract functionality supporting Turing-completeness, allows for a decentralized, immutable, and transparent business logic atop of blockchain. This technology is able to mitigate open problems of FL environments due to its inherent properties, namely:

Decentralization. In server-worker architectures, workers are exposed to a power imbalance and a single point of failure. A malicious server could (i) exclude workers arbitrarily or (ii) withhold reward payments. Furthermore, a server-worker design is not suitable for an environment where multiple entities share a common and equal interest in advancing their respective models. The decentral property of blockchain systems ensures a federal

TABLE I
SUMMARY OF DECENTRAL AND INCENTIVIZED FEDERATED LEARNING FRAMEWORKS

Ref.	FL	Type	NN-flex.	A	C	P	S	S on BC	CM	TA	Scal.
[45]	FedAvg	n.s.	✗	✓	✗	✓	✓	NN-p	n.s.	✓	limited
[44]	n.s.	n.s.	✗	✓	✗	✓	✓	NN-p, MD	Accuracy, data size	✓(SG)	limited
[52]	FedAvg	CS	✗	✓	✗	✗	✓	NN-p	Generic (Shapley values)	✗	limited
[46]	n.s.	CD	✗	✓	✗	✓	✓	NN-p	n.s.	✗	limited
[41]	FedAvg	CS	✗	✓	✓	✗	✓	NN-p	ED of model updates	✗	limited
[15], [43]	n.s.	n.s.	✗	✓	✓	✓	✓	NN-p, MD	Accuracy (generic)	✓(CT)	limited
[42]	n.s.	CD	✗	✗	✓	✗	✓	NN-p	Accuracy (loss)	✓	limited
[36]	FedAvg	n.s.	✗	✗	✓	✓	✓	NN-p	ED of model updates	✗	limited
[54]	n.s.	n.s.	✗	✗	✗	✗	✓	NN-p	Accuracy (loss)	✗	limited
[51]	n.s.	n.s.	✗	✗	✗	✗	✓	NN-p	n.s.	✓	limited
[50]	n.s.	n.s.	✗	✗	✗	✗	✓	NN-p,MD	Data size, sensing capacity	✓(RL)	limited
[49]	n.s.	n.s.	✗	✗	✗	✓	✓	NN-p	Accuracy	✓	limited
[48]	n.s.	n.s.	✗	✗	✗	✓	✓	NN-p	n.s.	✗	limited
[47]	FedAvg	n.s.	✗	✗	✓	✓	✓	NN-p,MD	Accuracy (generic)	✗	limited
[53]	FD	n.s.	✓	✓	✗	✓	✓	n.s.	Similarities to others	✗	(n.s.)
this work	FD	CS,(CD)	✓	✓	✓	✓	✓	logits, MD	PTSFD	✓	good

Ref. = Reference, NN-flex. = Neural Network flexibility, A = Aggregation, C = Coordination, P = Payment, S = Storage, S on BC = Storage on Blockchain, CM = Contribution measurement, TA = Theoretical analysis, Scal. = Scalability, MD = Metadata, n.s. = Not specified, CS = Cross-silo, CD = Cross-device, NN-p = Neural Network parameters (e.g. gradients, models, model updates), SG = Stackelberg game, CT = Contract theory, ED = Euclidean distance, RL = Reinforcement Learning.

system for entities with equal power without the need for a central server.

Transparency and Immutability. Since every peer in the system shares the same data, data on blockchain can only be updated and never deleted. A transparent and immutable reward logic in an FL context ensures trust on the worker side. On the other hand, each worker is audited and can therefore be held accountable for malicious behavior.

Cryptocurrency. Many general-purpose blockchain systems come with cryptocurrency functionality, e.g., the option to implement payment schemes within the business logic of the smart contract. Based on a reward mechanism of the FL system, workers can be rewarded immediately, automatically, and deterministically without the need for a trusted third party.

To analyze Blockchain systems, we categorize them into two main types:

1) *Application Specific Blockchain Systems (ASBS):* Blockchains which have to be adapted to a specific FL use-case require a novel infrastructure. This causes overhead in terms of complexity at the development, deployment, and hence more likely to introduce vulnerabilities.
2) *General Purpose Blockchain Systems (GPBS):* These are limited due to restricted virtual machines and predefined consensus layers, but allow for easy development, deployment, and operation utilizing already existing frameworks [4], [5], [33].

These types can either be public or permissioned/private. Public blockchains, like Ethereum, are open networks where anyone can participate, hence making it expansive to use as every transaction has to be duplicated by every node in the network. Permissioned blockchains, such as Hyperledger Fabric [5], restrict participation to authorized entities, offering a controlled, efficient, and private environment that may be preferable for FL scenarios with known and trusted participants.

D. Related Work

We focus on related Federated Learning Frameworks (FLF) that (i) are both decentral *and* reward participation as well as (ii) use blockchain at its core to decentralize FL. That is, parameters are aggregated *or* stored in a decentralized way.[1] We extended the systematic analysis established by [1] to compare FL, the application of Blockchain and the Contribution measurement in Table I. Note that the inherent complexity of FLF leads to heterogeneity in terms of application, overall design, special focus and details. [34] designs an FL system for home appliances using blockchain and a new normalization technique for differential privacy. Similarly, [40] introduces a regional FL framework for vehicles, integrating a reputation mechanism and a blockchain-secured trading platform. Focusing on robust mechanism deisigns, [15], [41] propose an FL protocol on blockchain using contest theory for worker engagement. [42] employs a Stackelberg game-based FL system considering contributions, deadlines, and upload times. [43] introduces DeepChain, a blockchain-secured FL framework with a special focus on privacy. [44] presents a two-layered blockchain for mobile edge networks. [45] minimizes communication costs in IoT FL through a double-layer aggregation model. [46] offers a specialized Democratic Learning (DemL) solution for on-device learning, including a unique consensus mechanism. [47] introduces Proof of FL (PoFL), an energy-efficient blockchain mechanism. [48] provides a secure FL framework for UAV-assisted sensing, incorporating differential privacy and reinforcement learning-based incentives. [49] designs a Mobile Crowdsensing framework that uses blockchain and edge intelligence for resource-constrained environments. [50] evaluates participant contributions transparently in its FL framework. [39] combines FL and blockchain for secure data sharing in neural

[1] Often, blockchain is only applied as an additional feature but not as part of the core infrastructure e.g., to store reputation [34], [35], [36], [37] or as a way to randomly choose an aggregator out of the client pool [38], [39].

training, using Shapley values for fair rewards. Lastly, [51] integrates blockchain and model distillation to accommodate model heterogeneity and enhance communication efficiency, yet it falls short in detailing blockchain operations and providing a theoretical analysis.

1) Incentivization of FLF: Measuring contributions in FL to fairly reward clients remains an open research challenge [1], [6]. Various metrics and methods are currently used for this purpose, each with its own set of challenges and limitations. [42], [46], [48], rely on self-reported information such as data size to determine rewards. However, this approach is susceptible to malicious behavior as false reporting leads to a maximal return. Alternatives include using similarity measures like the Euclidean distance of model updates [34], or employing voting systems for contribution assessment [41], [45]. Despite their utility, these methods lack rigorous theoretical and experimental validation and are vulnerable to attacks. Explicit methods like Shapley value [39], [50] or simple test-set accuracy [40], [47] have been utilized for explicit reward measurement. However, when applied in a decentralized context, these explicit methods (i) either require complex adjustments to the blockchain consensus mechanism, (ii) cause infeasible overhead (especially Shapley value), or require a central authority that measures the contribution against the test set, introducing a single point of failure. While [44] acknowledges multiple factors like data quality and task satisfaction as affecting rewards, it remains vague about its contribution measurement methodology. Similarly, [49], [52], and [46] lack specificity in this regard.

2) Decentralization of FLF: The trade-off in using blockchain lies between scalability and decentralization. Although it is theoretically favorable to decentralize all FL operations - namely Aggregation (A), Coordination (C), Payment (P), and Storage (S) - on-chain, doing so may introduce prohibitive computational and storage costs. This is because all blockchain nodes must replicate both computation and storage at all times. Specifically, the need to store and manipulate data-heavy objects, such as millions of NN parameters, on-chain restricts the framework to a limited number of participants. In summary, Table I compares decentralized and incentivized FLF to the approach presented in this work. Our approach is unique in allowing for NN flexibility (see Section II-B) while simultaneously maintaining full decentralization and scalability.

III. Problem Statement and Reward Mechanism

A. Problem Statement

We assume a federation $\mathcal{F}$ of workers $\mathcal{W}$ who have a common interest in advancing their private Neural Networks based on (i) additional data from other participants and (ii) the unlabeled public dataset X^{pub} through Federated Distillation (FD). We consider an environment where all participants of $\mathcal{F}$ have equal power. For example, no central entity such as a central server should have the power to either censor or manipulate the reward distribution. Each worker participating in the training is responsible for submitting predictions on the public dataset X^{pub} based on their locally trained model and label distribution $labelCount_i$ of the predictions. To enable decentralization, a smart contract atop a blockchain will replace the central server. This contract will (i) aggregate the workers' predictions and (ii) calculate the rewards considering other contributions. To ensure accountability and to prevent free-riding, each worker must stake a deposit D_i. $\mathcal{D} = \sum_{i \in \mathcal{F}} D_i$ will be used to pay $\bar{\tau}_i$ for each contribution at the end of the training process. Note that $\bar{\tau}_i \geq D_i$ if worker i's contributions are above average to $\mathcal{F}$ and $\bar{\tau}_i \leq D_i$ otherwise. Malicious behaviors, such as (i) withholding after committing and (ii) committing an incorrect label distribution $labelCount_i$, will result in the slashing of the deposit and exclusion from $\mathcal{F}$. The worker selection process is beyond the scope of this work. Reputation systems [36], [53] or required registrations might be feasible solutions. Our proposed framework is designed to be lightweight and blockchain agnostic. By employing 1-bit compressed logits on a public test set, instead of aggregating millions of parameters of modern NN (`FedAvg`), and incorporating a computationally simple, correlation-based reward mechanism, our framework uniquely enables (i) on-chain aggregation and (ii) on-chain reward calculation, while maintaining compatibility with both ASBS and GPBS. While theoretically possible, many promising public blockchain projects are still in their technological infancy, either lacking smart contract functionality or facing scalability restrictions. These constraints currently make deploying our system on public blockchains economically infeasible, due to high transaction fees and limited transactions per second, resulting in scalability issues. Consequently, our framework is specifically designed for the *cross-silo case* on *permissioned blockchains*. We assume the following properties:

1) *Honest Majority Assumption:* We assume that the majority of the nodes in the blockchain network are honest and follow the protocol. This is critical for the blockchain's consensus mechanism to function correctly.
2) *Sybil Attack Resistance:* We assume that our blockchain network is resistant to Sybil attacks, where an adversary controls multiple nodes. This is especially important for the GPBS deployment, where entry to the network is more open.
3) *Confidentiality and Integrity:* We assume that the blockchain ensures the confidentiality and integrity of the data and code.

B. Reward Mechanism Motivation

As no entity is in possession of the true labels of X^{pub} in the decentralized Federated Learning setting, workers' evaluations cannot be verified. This might encourage workers to report random data without actually classifying X^{pub}. This can be mitigated by rewarding peer consistency, e.g., the reward depends on its consistency with the label given by other workers. However, the best strategy in such schemes is for all workers to report the same answer without investing effort in finding the real label. The solution to these issues is to set up a mechanism, where the expected profit for each individual worker is maximized, if they put high effort into solving the task while acting truthful. In contrast to a server-worker relationship, our framework assumes multiple stakeholders with common interest in improving their

Algorithm 1: The Peer Truth Serum for Crowdsourcing [16].

Step 1
`CalculateFrequency` (*all tasks except worker i's report*)
Let $R_i(x) = \frac{\text{num}(x)}{\sum_y \text{num}(y)}$ be the frequency of reports (excluding worker i), where num counts the occurrences of reported values x

Step 2 `SelectPeerWorker` (*task j*)
Select the peer worker p for task j

Step 3 `RewardForReporting` (x_w)
Worker i is rewarded for reporting x_{ij} on task j with the score:

$$\tau(x_{ij}, x_{pj}) = \lambda \cdot (\tau_0(x_{ij}, x_{pj}) - 1)$$

where x_{pj} is the report from worker p on task j, and $\lambda > 0$. The function τ_0 is given by:

$$\tau_0\left(x_{ij}, x_{pj}\right) = \begin{cases} \frac{1}{R_i(x_{ij})} & \text{if } x_{ij} = x_{pj} \\ 1 & \text{if } x_{ij} \neq x_{pj} \end{cases}$$

respective model. The initially staked deposit D which will be used to pay τ manifests this mutual interest. Yet, contributions may be of different quality to the overall federation. Low quality workers may even have a negative effect on the overall federation even if their intention is truthful. At the same time, some classes in X^{pub} may be less common and therefore are more important to classify correctly. Hence, a mechanism is required to:

1) incentivize only workers with the best abilities for the task
2) incentivize these workers to invest their utmost effort in obtaining the most accurate answer
3) incentivize workers who are able to classify uncommon samples in X^{pub} with higher rewards

C. Peer Truth Serum for Federated Distillation

The Peer Truth Serum for Crowdsourcing (PTSC) is a promising Multi-task Peer Prediction mechanism. Through a scoring rule τ, it rewards workers for surprisingly common reports, encouraging honest and high-effort behavior without the need for ground-truth knowledge [16]. PTSC merges the reward mechanism of [54] with the Peer Truth Serum concept [55], [56], ensuring incentive compatibility across a non-binary solution space suitable for heterogeneous workers. Introducing the Peer Truth Serum for Federated Distillation (PTSFD), we adopt the PTSC framework, as described in Algorithm 1, for the Federated Distillation setting. In this scenario, a group of workers perform statistically independent tasks, where a task refers to classifying a sample j, with $j \in X^{pub}$. The discrete density function is represented as

$$R_i(x) : \mathcal{C} \mapsto [0, 1], \sum_{x \in \mathcal{C}} R_i(x) = 1. \tag{2}$$

Here, $R_i(x)$ excludes the contribution from worker i and denotes the fraction of reported labels, given by

$$R_i(x) = \frac{\text{labelCount}(x)}{\sum_{y \in \mathcal{C}} \text{labelCount}(y)}. \tag{3}$$

Furthermore, PTSFD incorporates an adjustable penalty term β. This modification acknowledges that the primary motivation might be the utility of an improved model, making the payment secondary (as seen in (6)). Consequently, the reward for each sample is

$$\tau_{ij}\left(x_{ij}\right) = \lambda \cdot \left(\frac{1}{n_{\text{peers}}} \sum_p \tau_0\left(x_{ij}, x_{pj}\right) - \beta \right). \tag{4}$$

where λ adjusts the payment magnitude and β modulates the reward-accuracy ratio. The cumulative reward for worker i is computed over all tasks as

$$\bar{\tau}_i = \sum_{j \in X^{pub}} \tau_{ij}. \tag{5}$$

D. Game-Theoretic Analysis

The setting can be considered a two-stage game. In stage 1, workers choose the amount of effort e they want to invest in classifying X^{pub}. To simplify the analysis, we assume two levels of effort, high e_1 and low e_0. Here, e_1 represents the best work possible exerted by the worker, and e_0 represents no effort (i.e., no local NN training). Unlike in `FedAvg`-based systems, the proposed framework applies FD, hence it does not require a uniform NN among the clients but allows for flexible architectures appropriate for the hardware constraints of the respective workers (see Section II-B). In stage 2, workers decide on what to report. The baseline model assumes each worker solves every task. Yet, without loss of generality, workers could be randomly allocated to solve tasks such that each sample of X^{pub} is classified by at least two different workers.

Workers. We assume workers to be individually rational, aiming to maximize their expected profit $\Pi_i = \text{Rewards}_i - \text{Costs}_i$

$$\max \mathbb{E}\left(\Pi_i\right) = U_i(\theta_i^{improved}) + U_i\left(\bar{\tau}_i - [c_i(e_i) + c_{\mathcal{S}}^{fix}]\right). \tag{6}$$

U_i represents the expected utility function of worker i, which can vary among workers. The expected rewards of contributing to the federation $\mathcal{F}$ are twofold: (i) the expected utility of the improved model $\theta_i^{improved}$ and (ii) the utility of the expected monetary reward from $\mathcal{S}$ for contributing to classify X^{pub}. We assume that the training process incurs variable costs $c_i(e_i)$, where c_i is an increasing function of effort e_i. Specifically, $c_i(e_1) > c_i(e_0)$, where e_0 denotes no effort and e_1 denotes high effort of worker i. Effort represents the quality and quantity of private data, model quality, number of training iterations, etc. A detailed Pareto-optimal cost analysis [57], [58] under real-world assumptions will be explored in future work. Additionally, to offset free-riding of inactive but registered workers of $\mathcal{S}$ who benefit from an improved model $U_i(\theta_i^{improved})$ without contributing, fixed participation costs $c_{\mathcal{S}}^{fix}$ are necessary. The

initially staked deposit is used to pay contributing workers. Thus, $c_{\mathcal{S}}^{fix} = D_i^{before} - D_i^{after}$ describes the implicit costs for accessing Y_{aggr}^{pub}.

Incentive Compatibility. In order to evaluate PTSFD in game theoretic terms, we analyze each workers expected profit $\Pi_i = Rewards_i - Costs_i$. We assume *Individual Rationality* (IR), e.g., workers try to maximize their expected profit and do not participate if $\Pi \leq 0$. For the sake of simplicity, we further assume that the gain in model improvement $U_i(\theta_i^{improved})$ is offset by $U_i(c_{\mathcal{C}}^{fix})$. When a worker classifies a sample, it obtains an evaluation Y_j^{eval} which can be different from the reported value Y_j^{report}. In stage two, workers face three different strategies $\forall j \in X^{pub}$ [16]

1) *Honest* Invest high effort e_1 to obtain Y_j^{eval} and report honestly, s.t. $Y_j^{report} = Y_j^{eval}$
2) *Strategic* Invest high effort e_1 to obtain Y_j^{eval} but reports $Y_j^{report} \neq Y_j^{eval}$
3) *Heuristic* Do not invest any effort e_0 and randomly report Y_j^{report} based on the a-priori known distribution of labels in X^{pub}

We define the mechanism to be incentive compatible, if the honest strategy is the dominant strategy for every worker. We use an equilibrium analysis to determine the resulting behavior of each worker. In particularly, $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n)$ represents a strategy profile of each worker. This profile is an equilibrium $\bar{\sigma}$ if for any worker $i \in \mathcal{W}$, the workers expected profit is maximized with the honest strategy profile $\bar{\sigma}$. Suppose that worker i believes that the peer workers are honest and their answer on a given sample j is positively correlated with the worker i's answer x, when obtained with high effort e_1. Specifically, worker i believes that answer x is not less likely for sample j than in the distribution over all tasks.

Honest Strategy. For every sample j in X^{pub}, the worker calculates the probability scores over all possible classes in $\mathcal{C}$ (output of the softmax layer of a NN). Let us further assume worker i is in possession of a trained model θ_i, with an overall accuracy $Accuracy_{\theta_i}$. We define the relative certainty $\mathbb{A}_{ij}$ of any prediction of client i on an element j of X^{pub} as the product of the local classifier accuracy and the sample-specific maxprobabilityscore.

$$\mathbb{A}_{ij} = Accuracy_{\theta_i} \cdot MaxProbabilityScore_{ij}. \tag{7}$$

Under the assumption that the local client data X_i^{priv} is representative of the entire data distribution D, this metric will give a heuristic measure for the data specific certainty in the model prediction. Based on this metric, each worker will make the decision whether to report predicted labels, discarding those for which reward is expected to be negative. This leads to the expected profit

$$\mathbb{E}(\Pi_{ij}) = \mathbb{A}_{ij} \cdot \lambda \left(\frac{1}{R(x_{ij})} - \beta \right) + (1 - \mathbb{A}_{ij}) \cdot \lambda(-\beta) - c_i(e_i). \tag{8}$$

Assuming individual rationality, $\mathbb{E}(\Pi_{i,j}) \geq 0$ in order to incentive worker i to submit a vote on sample j. Following 8, we can derive minimum prediction quality

$$\mathbb{A}_{ij} \geq R(x_{ij}) \cdot \left(\frac{c_i(e_i)}{\lambda} + \beta \right), \tag{9}$$

required to incentivize worker i to participate, e.g., $\Pi_i \geq 0$. Notice that the federation can set the overall quality threshold by adjusting hyperparameter λ and β appropriately, assuming similar variable costs $c(e)$ on the workers side.

Heuristic Strategy. The heuristic strategy assumes that worker i does not exert any effort to obtain $Y_j^{eval} = x \in \mathcal{C}$. The expected reward is based on the probability of matching a peer's answer. Given that answer x is independent of the task, the probability of coincidentally matching a peer is equivalent to the frequency of answer $x \in \mathcal{C}$.

$$\mathbb{E}(\Pi_{ij}) = R(x_{ij}) \cdot \left(\frac{1}{R(x_{ij})} - \beta \right) + (1 - R(x_{ij})) \cdot (-\beta) = 1 - \beta. \tag{10}$$

It's important to note that the expected profit for $\beta = 1$ is 0, and it is strictly negative for $\beta > 1$. This holds irrespective of what answer x is, or what the worker knows about the distribution $R(x)$ across labels in X^{pub}. Given that the noise introduced to classify X^{pub} diminishes the overall model quality, as per (8), a rational worker is unlikely to participate if $\beta \geq 1$.

Strategic Strategy. Assuming the honest participation of other workers, exerting Y_j^{eval} while reporting $Y_j^{report} \neq Y_j^{eval}$ consistently results in a negative expected profit for all $j \in X^{pub}$.

$$\mathbb{E}(\Pi_{ij}) = \lambda \cdot (-\beta) - c_i(e_i). \tag{11}$$

This is true provided the self-predicting condition [16] is met, i.e.,

$$\frac{A_{ij}(x|x)}{R(x)} > \frac{A_{ij}(\bar{x}|x)}{R(\bar{x})}, \forall \bar{x} \neq x. \tag{12}$$

Considering a scenario where workers collude (they report x for both $Y_j^{eval} = x$ and $Y_j^{eval} = y$), R will adjust such that $R^{col}(x) = R(x) + R(y)$.

$$\mathbb{E}(\Pi_{ij}) = \begin{cases} \mathbb{A}_{ij} \cdot \lambda \left(\frac{1}{R_i(x_{ij}) + R_i(y_{ij})} - \beta \right) \\ +(1 - \mathbb{A}_{ij}) \cdot \lambda(-\beta) - c_i(e_i) & \text{if} \quad Y_{ij}^{eval} = x_{ij} \\ \mathbb{A}_{ij} \cdot \lambda \left(\frac{1}{R_i(x_{ij}) + R_i(y_{ij})} - \beta \right) \\ +(1 - \mathbb{A}_{ij}) \cdot \lambda(-\beta) - c_i(e_i) & \text{if} \quad Y_{ij}^{eval} = y_{ij} \end{cases}. \tag{13}$$

This change in R entirely offsets the increase in the probability of a match. Therefore, only an honest strategy paired with a high-quality model will yield a positive expected reward for a given worker. This results in an equilibrium $\bar{\sigma}^{honest}$ for the PTSFD mechanism, thereby demonstrating its incentive compatibility.

IV. 1-Bit Compressed Federated Distillation Framework With Smart Contract Logic

The protocol consists of the following steps: (i) Task Specification & Contract Deployment, (ii) Worker Registration & Deposit, (iii) Local Model Training, (iv) Hash Commit Submission,

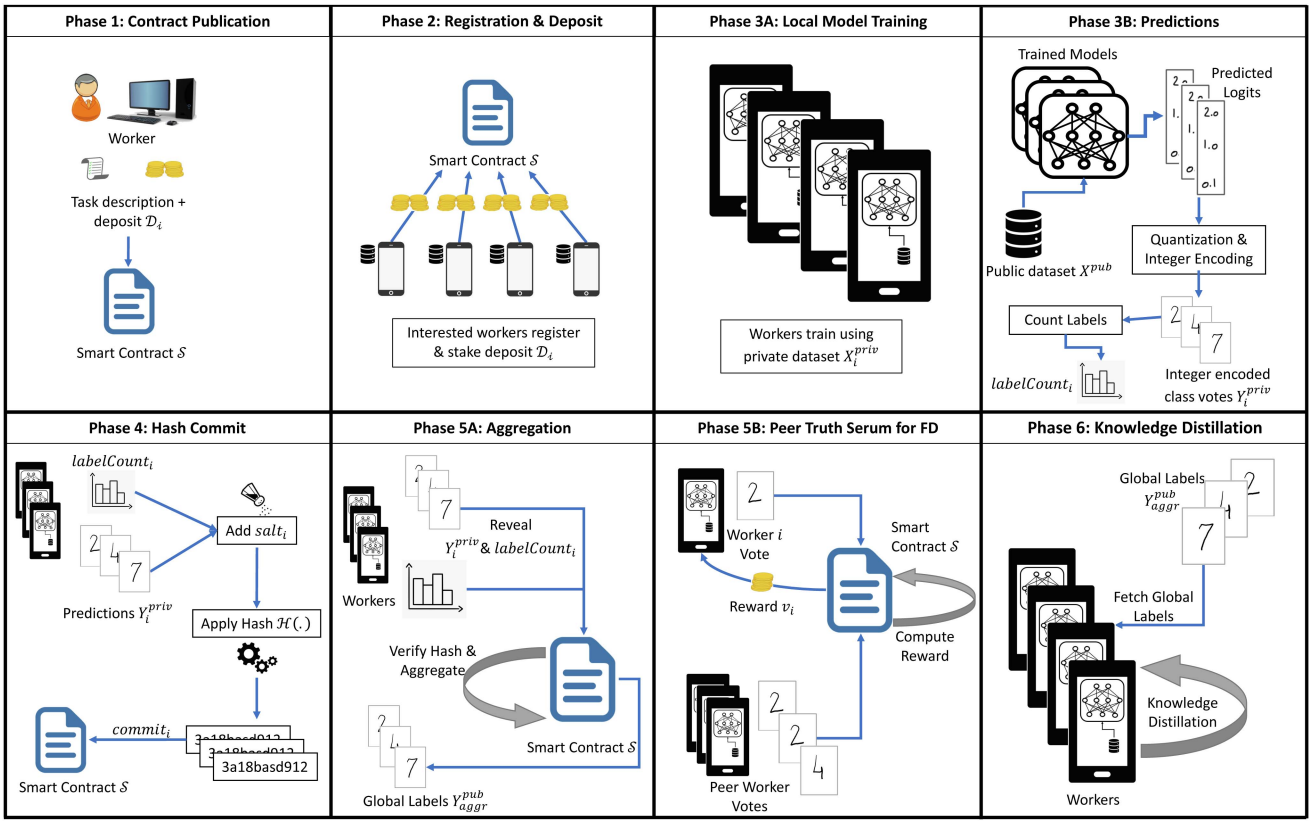

Fig. 3. Iterative process of on-blockchain federated distillation.

(v) Reveal Predictions, (vi) Aggregation & Reward Distribution, and (vii) Knowledge Distillation from X^{pub} as depicted in Fig. 3.

A. Task Specification & Smart Contract Deployment

To form Federation $\mathcal{F}$, participants with similar interests must agree on the requirements and specifics of an FD task, specifically:

1) Task description and data distribution (e.g., images of a certain type).
2) Reference to a public data set X^{pub} and potential classes $\mathcal{C}$ for the Federated Distillation pipeline. This will subsequently be utilized by workers to predict the labels on each sample of the dataset.
3) Reference to the address of $\mathcal{S}$.
4) Deposit amount D_i that each worker must stake.
5) PTSFD and reward mechanism details (λ and β values).

After forming a federation $\mathcal{F}$, either an external third party or one of the workers from $\mathcal{F}$ deploys the governing smart contract $\mathcal{S}$, stakes the necessary deposit D_i, and lists the addresses of all eligible workers in $\mathcal{F}$, as well as the aggregation and PTSFD logic of the FD task.

B. Worker Registration & Deposit Submission

Based on the task specifications, interested workers register on the smart contract $\mathcal{S}$ using their respective blockchain address (public key) and submit the required deposit D_i. $\mathcal{S}$ verifies if the applying worker belongs to the federation. Given that $|\mathcal{F}| >> |\mathcal{W}|$, PTSFD motivates valuable workers for $\mathcal{F}$ in terms of data and computational capacity to engage while dissuading low-quality workers, as demonstrated in Section IV-E. To preclude free-riding, workers in $\mathcal{F}$ who aren't registered shouldn't access $\mathcal{S}$. This restriction can be implemented by deploying $\mathcal{S}$ on a suitable blockchain system or by shuffling X^{pub}, ensuring only registered clients can access the correct indices.

C. Local Model Training and Prediction

The entire training procedure encompasses two phases: the local model training phase on local data X_i^{priv}, Y_i^{priv} and the KD phase from X^{pub}, Y_{aggr}^{pub}, which occurs as the protocol's final step, as detailed in Section II-B.

Training on Local Data. Each worker either has a pre-trained model or begins training a NN on their specific private data until convergence (optionally, until a predetermined minimum accuracy agreed upon within $\mathcal{F}$ is achieved). Notably, unlike `FedAvg`, FD doesn't mandate a common shared NN architecture across all workers, thus enabling them to select an optimal architecture tailored to their computational resources.

Label Prediction. Upon completing the training, workers compute the soft labels $Y_i^{pub} = \{f_{\theta_i+\Delta\theta_i}(x) \mid x \in X^{pub}\}$ and subsequently quantize these to 1-bit as $Y_i^{1\,bit} = Q_{1\,bit}(Y_i^{pub})$.

Algorithm 2: Local Label Count for Worker i.

input : Integer encoded class votes x_{ij}, where $i \in \mathcal{W}' \subseteq \mathcal{W}$
output: $labelCount_i$

init $labelCount_i$
  var $labelCount_i \in \mathbb{N}^{|\mathcal{C}|} = (0, 0, \ldots, 0)$
foreach $j \in X^{pub}$ **do** `// iterate over data samples`
  $labelCount_i(x_{ij})$ += 1
return $labelCount_i$

Label Count. The PTSFD mechanism necessitates data on the label distribution $R(x)$ over X^{pub} for reward calculations. Hence, each worker i must compute the label count $labelCount_i \in \mathbb{N}^{|\mathcal{C}|}$ for each label present in X^{pub}, to minimize computational overhead on the blockchain (as described in Algorithm 2). The supplemental validation function to ensure the accurate computation of $labelCount_i$ depends on the specific blockchain system and is outside this work's purview.

D. Commit and Reveal

Information on the blockchain is transparent to every node. Even in a private blockchain setup, workers in $\mathcal{W}$ could wait for peers to publish $Y_p^{1\,bit}$ and replicate their results without expending any effort. To prevent this kind of copying and to ensure that workers apply effort to classify X^{pub}, a two-step commit and reveal scheme is employed.

Commit. Prior to publishing the results to $\mathcal{S}$, where all peer workers could view the submission, a cryptographic hash $hashCommit_i = \mathcal{H}(Y_i^{1\,bit}, salt_i, labelCount_i)$ is computed to obfuscate $Y_i^{1\,bit}$ and $labelCount_i$. The property of pre-image resistance of a cryptographic hash function (e.g., it should be challenging to find any message m such that $commit_i = \mathcal{H}(m)$) and the property of collision resistance (e.g., it should be challenging to find two distinct messages $m1$ and $m2$ with $\mathcal{H}(m_1) = \mathcal{H}(m_2)$) ensure that no worker can either retrieve $Y_i^{1\,bit}$ or alter their previously committed $Y_i^{1\,bit}$. Each worker i transmits $hashCommit_i$ to $\mathcal{S}$ upon completing their training. It's noteworthy that the commit phase on $\mathcal{S}$ concludes once $|\mathcal{W}'| \subseteq |\mathcal{W}|$ workers have registered with $\mathcal{S}$ or when the maximum time T_{commit}^{max} is reached.

Reveal. During the reveal phase on $\mathcal{S}$, each worker that successfully committed in the commit phase must disclose $Y_i^{1\,bit}$, $labelCount_i$, and $salt_i$ within the timeframe T_{reveal}^{max} via a transaction function call to $\mathcal{S}$. To counteract withholding attacks, a worker's deposit D_i is forfeited if worker i fails to reveal within the allotted time T_{reveal}^{max}. The smart contract then verifies the commitment's validity to confirm that $\mathcal{H}(Y_i, labelCount_i, salt_i) == hashCommit_i$.

Algorithm 3 provides the pseudocode for this scheme in Solidity on the Ethereum blockchain.

Algorithm 3: Commit and Reveal Protocol.

Data: `32byte` $hashCommit_i \leftarrow \mathcal{H}\left(Y_i^{1bit}, labelCount_i, salt_i\right)$
, **Init**
  var `commitments` $\leftarrow$ Mapping($address_i \rightarrow$ byte32) $\forall i \in \mathcal{W}$
  var `userIsCommitted` $\leftarrow$ Mapping($address_i \rightarrow$ bool) $\forall i \in \mathcal{W}$
Phase I `commit` (*hashCommit*)
  foreach $i \in \mathcal{W}'$ **do**
    `require` (*userIsRegistered(msg.sender)*) `// registered in` $\mathcal{W}$
    `require` (*!userIsCommited(msg.sender)*)
    $commitments.append(commit_i)$
    $isCommitted(msg.sender) = True$
Phase II `reveal` $(Y_i^{1bit}, salt)$
  foreach $i \in \mathcal{W}'$ **do**
    `require` (*userIsCommitted(msg.sender)*)
    `require` $(\mathcal{H}\left(Y_i^{1bit}, labelCount_i, salt_i\right) == commitments(msg.sender))$

E. Aggregation & Reward Distribution

We apply PTSFD to calculate the reward distribution for each worker. In order to calculate the rewards, $\mathcal{S}$ aggregates $labelCount_i$ across all workers $i \in \mathcal{W}'$ first to obtain the global label count $G = \sum_{\mathcal{W}'} labelCount_i \in \mathbb{N}^{|\mathcal{C}|}$. G is a helper variable to calculate R_i

$$R_i = \frac{1}{m \times n} \times (G - labelCount_i)\,. \tag{14}$$

The worker is rewarded for its prediction on sample j with respect to it's peers regarding (4). The final *rewardScore* for worker i is a sum of all individual rewards over X^{pub}, given by

$$\begin{aligned} \bar{\tau_i} &= rewardScore(i) \\ &= \lambda \cdot \left(\frac{1}{n_j^{peers}} \sum_j \sum_p \tau_0(x_{ij}, x_{pj}) \right) \forall i \in \mathcal{W}', \end{aligned} \tag{15}$$

where parameter λ describes a scaling parameter for the reward and n_j^{peers} describes the number of peer workers who also submitted a label prediction on j. The aggregated predictions Y_{aggr}^{pub} are calculated by majority vote of $Y_i^{1\,bit}$ $\forall i \in \mathcal{W}'$. We merge the reward computation and aggregation into a single algorithm outlined in Supplementary Materials B. Note that implementation details may differ fundamentally depending on the underlying blockchain architecture.

F. Knowledge Distillation on Public Dataset

Finally, workers download the aggregated predictions Y_{aggr}^{pub} from the blockchain and perform several epochs of KD using X^{pub} and Y_{aggr}^{pub} to improve their respective models ($\theta_i^{improved} \rightarrow \theta_i + \Delta\theta_i$). Optionally, the training process of each client, as shown in Section II-B2, can be repeated until a specific threshold is achieved as specified in the Smart Contract

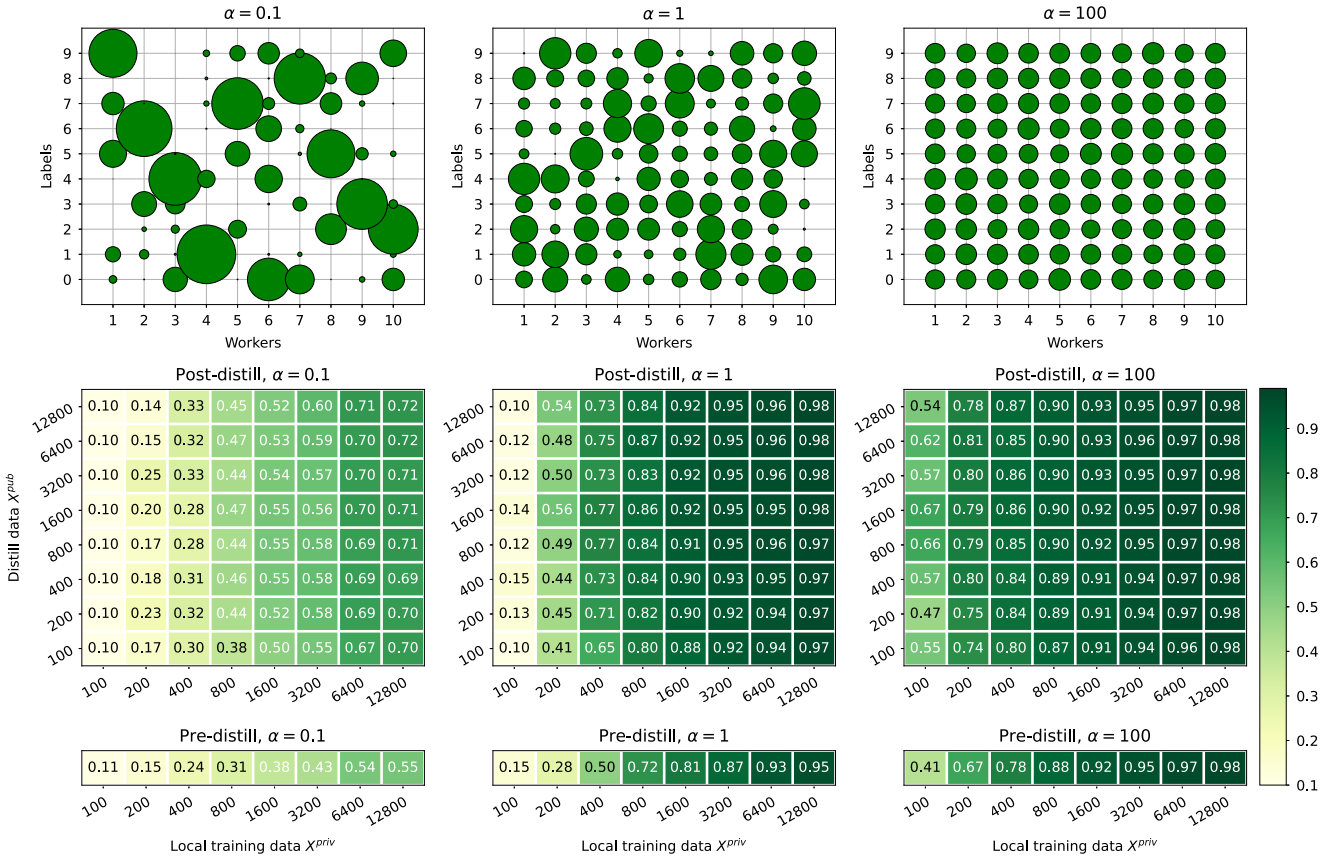

Fig. 4. Influence of X^{pub} and X^{priv} on model accuracy with corresponding dirichlet α setting. These experiments were run using LeNet on MNIST/EMNIST respectively as training/distillation datasets.

$\mathcal{S}$. Note that λ should decrease with every consecutive round, as most evaluated labels will not change.

G. Complexity Analysis

Because we are executing this protocol on the blockchain, it's essential to understand the computational and storage costs involved. In this section, we discuss the overheads related to computation and storage that our proposed algorithm introduces. While the actual implementation on a general-purpose blockchain system might differ based on the underlying virtual machine, our PTSFD implementation, as illustrated in Supplementary Materials B, provides a useful reference to estimate the complexity.

1) Computational Complexity: The algorithm presented in Supplementary Materials B first calculates the global label distribution and counts class votes across all workers (*lines 7–12*). This computational overhead is $\mathcal{O}(m \cdot n)$, where $m = |X^{pub}|$ and $n = |\mathcal{W}'|$. Subsequently, we examine each data sample in X^{pub}, rewarding or penalizing a worker based on its peers. We also determine the aggregated class label for each sample during this algorithm phase (*lines 13–29*). The process of calculating the reward for each worker based on its peers results in a computational overhead of $\mathcal{O}(m \cdot \sum_{i=1}^{n} n_{peers})$. The global label calculation adds a cost of $\mathcal{O}(m \cdot |\mathcal{C}|)$. In the baseline scenario, where each worker processes all data samples from the public dataset and is considered a peer of all other workers, the overall computational cost is given by (16).

$$\mathcal{O}(m \cdot (n^2 + |\mathcal{C}|)). \tag{16}$$

For more efficient solutions, we distribute public dataset samples among workers such that each sample is classified by a maximum of two workers. Implementing PTSFD in this manner would reduce the overhead as detailed in (17).

$$\mathcal{O}(m \cdot (2n + |\mathcal{C}|)). \tag{17}$$

2) Storage Complexity: Two types of storage costs are associated with the proposed algorithm: permanent storage and temporary memory variables. $Votes$, M, S, R_i, and τ_0 need memory storage during the computation, resulting in $\mathcal{O}(|\mathcal{C}| \cdot (m+2) + n)$ additional memory storage. Whether the reported frequencies $labelCount_i$ or each worker's final reward share $rewardScore$ need permanent blockchain storage depends on the requirements of the specific blockchain system. Ideally, only $globalLabels = Y^{pub}_{aggr}$ is stored permanently on the blockchain. The minimum amount of data required for each round is represented by (18), where η accounts for the overhead due to encoding necessities.

$$\mathrm{b}_{globalLabels} = n \times |\mathcal{C}| \times 1\mathrm{bit} + \eta. \tag{18}$$

H. PTSFD in Comparison to Shapley Value

Quantifying individual contributions in FL is essential for the equitable distribution of rewards and the growth of FL systems. The Shapley value is a concept in cooperative game theory that distributes total gains among players by measuring the marginal contributions each player makes to different possible coalitions. For a worker i in a set $\mathcal{W}$, with utility evaluation function $V : 2^{\mathcal{W}} \rightarrow \mathbb{R}$, the Shapley value $\phi_i(V)$ is given by

$$\phi_i(V) = \frac{1}{|\mathcal{W}|!} \sum_{\Pi} \left[V\left(S_i^{\Pi} \cup \{i\}\right) - V\left(S_i^{\Pi}\right) \right], \tag{19}$$

where Π is the set of all permutations of $\mathcal{W}$, and $S_i^{\Pi} \subset \mathcal{W}$ is the set of workers preceding i in permutation Π. The Shapley value adheres to several axioms for fairness:

1) *Efficiency:* Total utility is fully distributed among the workers: $\sum_{i \in \mathcal{W}} \phi_i(V) = V(\mathcal{W})$.
2) *Symmetry:* Workers contributing identically to every subset receive identical rewards.
3) *Null Player:* Workers who do not enhance utility for any coalition yield no reward.
4) *Additivity:* The Shapley value is linear over the utility functions of combined games.

Despite these favorable properties, the computation of the Shapley value scales with $O(2^n)$. In FL, $V(S) = V(\theta_S)$, where θ_S is the FL model trained on the subset of datasets $\{X, Y\}_S = \{X_i, Y_i\}, i \in S$ from scratch for every permutation:

$$V(S) = V(\theta_S) = V(A(\theta^{init}, \{X, Y\}_S)), \tag{20}$$

where $A(\cdot)$ is a learning algorithm and θ^{init} denotes the initial model. The computational demands involved in calculating the Shapley Value render it challenging to efficiently execute even a single step on a general-purpose blockchain. Recent studies have attempted to approximate the Shapley value to mitigate this computational overhead [59], [60], [61].In contrast, PTSFD offers a lightweight method to elicit truthful contributions implicitly, bypassing the exhaustive computation of the marginal utilities. PTSFD, as an alternative, is designed to be compatible with blockchain technology, enabling scalable and decentralized FL without excessive computation, while still aspiring to maintain fairness in incentives.

I. Limitations

Despite the benefits of our decentralized FD protocol, our framework encounters the following limitations.

Public Dataset. Although FD offers numerous advantages, such as reduced information exchange and the flexibility of independent NN architectures, the FD training process necessitates access to a public dataset X^{pub}. This may not be available for certain use cases. While [62] demonstrated that highly disparate data distributions might suffice for FD, relying on $\mathbb{A}_{ij}$ as a heuristic for the evaluation certainty of a sample limits the variance of distributions between X^{pub} and X^{priv}. The scoring method proposed by [63] appears promising in addressing this.

Public Blockchains. Despite significantly reducing computational and storage demands, our framework remains unsuitable for current public blockchain systems due to (i) the high costs associated with storing Y_{aggr}^{pub} and the computational overhead of PTSFD, and (ii) the transparency of Y_{aggr}^{pub} to nodes that are not members of $\mathcal{W} \in \mathcal{F}$ and hence did not make a deposit. Both issues may be addressed by upcoming advancements in the public blockchain sector.

Self Predicting Condition. PTSFD is incentive-compatible and yields an optimal outcome when workers are honest. However, the mechanism's incentive compatibility is contingent upon the satisfaction of (12). If classes are evenly distributed across X^{pub}, the conditions will always be met. Example: Let $Pr(x = a) = 0.8$ and $Pr(x = b) = 0.2$, but $R(a) = 0.9$ and $R(b) = 0.1$. Although the worker's $Y_i^{eval} = a$, their expected reward would be greater if $Y_i^{report} = b$, as $\frac{0.8}{0.9} - \beta < \frac{0.2}{0.1} - \beta$.

V. Experiments

In this section, we empirically evaluate the PTSFD framework and analyze the reward distribution under different levels of effort as well as its robustness in the event of malicious behavior. We do not consider explicit variable costs $c_i(e)$. Furthermore, we set the reward scaling parameter $\lambda = 1$ for all experiments. We do not account for lagging workers, so $\mathcal{W}' = \mathcal{W}$ for all experiments. All experiments are based on multiple rounds of the proposed protocol. The code for our experiments is made publicly available.[2] Specifically, we experimentally validate the following properties of PTSFD:

1) *Performance:* Choosing to participate in the federation should lead to a significant improvement in model accuracy for each worker.
2) *Fairness:* The greater the effort a worker exerts in terms of training accuracy and amount of training data, the higher the reward they should receive.
3) *Robustness:* Malicious workers should receive substantially less reward, even under high collusion rates.

A. Data Sets and Models

We analyze the decentralized 1-bit compressed FD with the PTSFD protocol on three different federated image classification problems, using EMNIST [64] / MNIST [65], CIFAR-10 [66] / STL-10 [67] and Fashion MNIST [68] datasets on ResNet-18 [69] and LeNet [70] and respectively as training/distillation data. Our Federation comprises 4, 10 and 25 workers for different experiments. The training data is distributed among workers according to a Dirichlet distribution with the Dirichlet parameter α. Fig. 4 top row illustrates the data distribution of 10 labels across 10 different workers for $\alpha = 100$, $\alpha = 1$, and $\alpha = 0.1$. The different alphas simulate various data distributions such as iid and non-iid. We first train models locally on X^{priv} and then perform KD using the public dataset X^{pub}. Even though in real-world PTSFD applications, workers might train different model architectures and vary the number of local training epochs based on their hardware constraints, we employ a single default NN architecture for simplicity. We simulate heterogeneity through varying local training accuracy (early stopping), non-iid data, and different sizes of X^{pub}. It's important to note that the

[2] https://github.com/Tsinghua-FL-Team/decentralized-FD

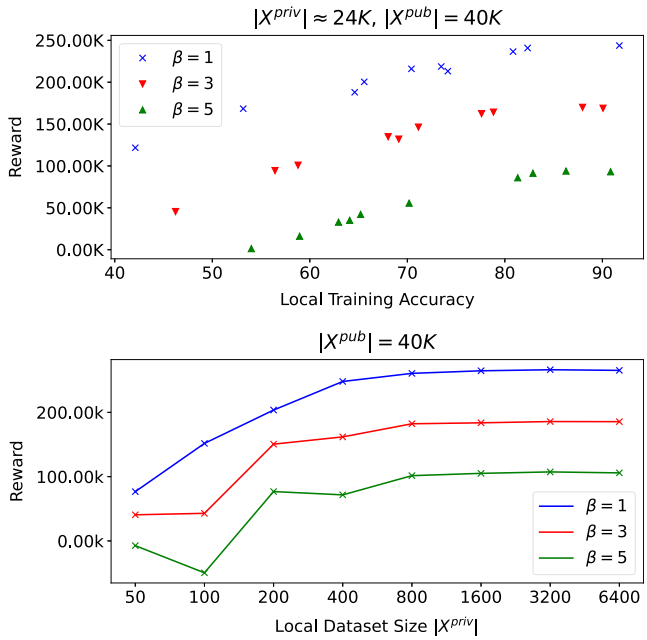

Fig. 5. Effect of local training accuracy & local data size on reward using LeNet on EMNIST.

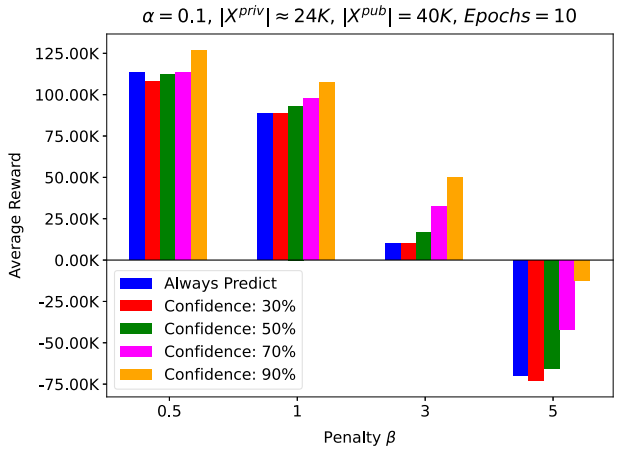

Fig. 6. Effect of confidence based predictions on reward with different β (penalty) using LeNet on MNIST.

distribution of the distillation data deviates from the worker's data, mirroring realistic FL scenarios (e.g., MNIST contains handwritten digits, while EMNIST features a different set of handwritten numbers; similarly, CIFAR-10 includes a distinct set of images compared to the STL-10 dataset). We use the Adam optimizer [71] with a fixed learning rate of 0.001 for both the distillation and training processes. We minimize cross-entropy loss for local model training on X^{priv}, Y^{priv} and minimize Kullback-Leibler Divergence on X^{pub}, Y^{pub}_{aggr}.

B. Storage and Communication Cost

Given that storage and computation costs on the blockchain are critical scalability constraints for decentralized FLFs [1], we assess our framework's storage costs in comparison to `FedAvg` for specific target accuracies. Table II shows the communication cost (upstream/downstream) as well as the storage cost required to achieve a specific accuracy target for ResNet-18 and LeNet on the CIFAR-10, MNIST, and Fashion-MNIST datasets, respectively. All experiments were run under different data distributions by varying the Dirichlet parameter α while simulating 4, 10, and 25 workers respectively. As observed from the summarized results, our framework achieves accuracy similar to that of a typical Federated Learning system using `FedAvg` as the aggregation mechanism but at a fraction of the communication and storage costs. For instance, when training ResNet on CIFAR-10 with $\alpha = 100$, we can achieve the target accuracy with a minimal communication cost of 0.84 MB and storage cost of 0.92 MB, compared to the substantial costs incurred by `FedAvg` (5.33 GB and 5.37 GB for communication and storage costs, respectively, an improvement of roughly 5,000x). The remaining experiments, showcasing model quality improvement, reward fairness, and simulation of collusion or heuristic behavior, were conducted using LeNet on EMNIST/MNIST as training and distillation datasets. We believe that these experiments sufficiently demonstrate the desired results and would yield comparable outcomes for other datasets or models.

C. Model Quality Improvement

Fig. 4 illustrates the impact of the sizes of the local dataset $|X^{priv}|$ and the public dataset $|X^{pub}|$ on the accuracy for EMNIST/MNIST under both non-iid distributions ($\alpha = 0.1$ & $\alpha = 1.0$) and the iid distribution ($\alpha = 100$) for 10 clients. A marked improvement in model quality is evident for every worker following their KD execution. While the size of the local training dataset plays a more substantial role than the distillation dataset, the significance of the latter should not be underestimated, particularly in the context of non-iid distributions which can be attributed to the inclusion of additional data from the public dataset.

D. Fair Effort-Reward Correlation

Fig. 5 illustrates the correlation between effort, in terms of both heterogeneous training efforts (top) and varying data quantities (bottom), and the subsequent reward distribution. PTSFD facilitates realistic FL scenarios where workers, due to hardware constraints, can employ various local model architectures and train them for different numbers of epochs. For instance, we emulated heterogeneity by training 10 workers using diverse early stopping criteria, where higher local training accuracy, indicative of greater effort, corresponded to better rewards. Concurrently, variations in private data quantity, while assuming consistent data quality, also influence contribution quality. Overall, superior training accuracy and more extensive local dataset result in greater rewards.

E. Robustness of PTSFD

In order to ensure the desired quality of label predictions, the federation can adjust the parameter λ to scale the reward based

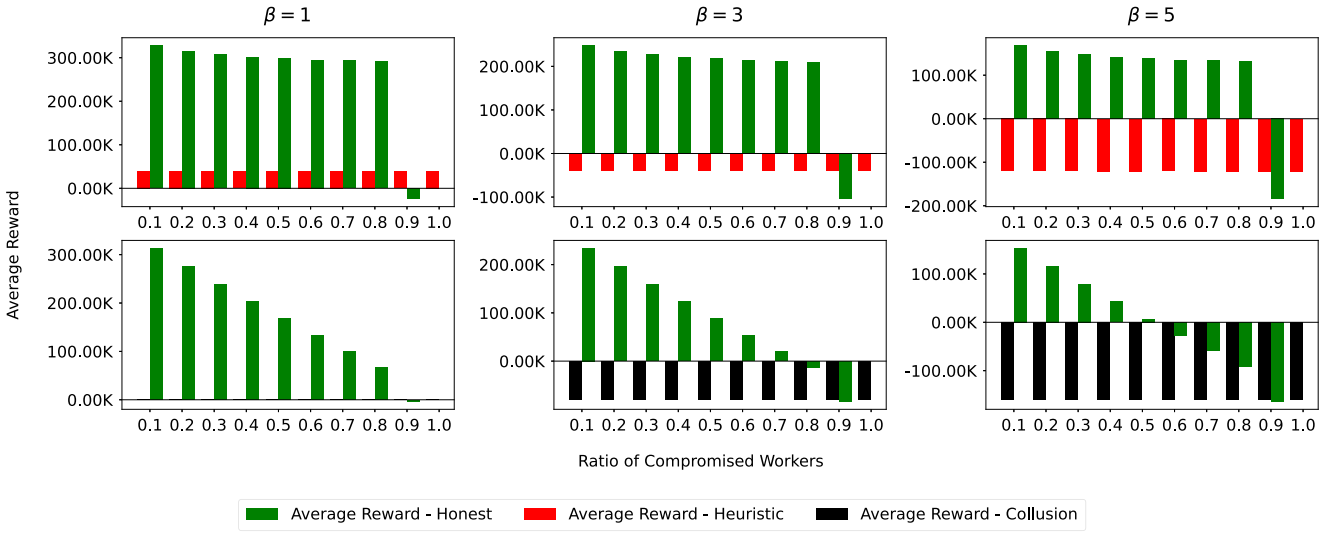

Fig. 7. Average reward with varying ratio of colluding and heuristic workers under different penalty β. Federated Learning setting with 10 workers running LeNet on EMNIST digits for 10 epochs. For all experiments, 40000 data points from the MNIST data set were used as public dataset X^{pub}.

TABLE II
COMMUNICATION AND STORAGE COSTS (IN MB) FOR ACHIEVING SPECIFIC ACCURACY IN FL ACROSS VARIOUS DATASETS, ARCHITECTURES, AND DATA HETEROGENEITY LEVELS

Model	Train / Distill Dataset	Target Accuracy	α	Up/Down	FedAvg	1-bit FedDistill
ResNet-18	CIFAR-10/Testset/[STL-10]	0.74 [0.68]	100.0	up/down	5332.50 (5)	0.61 (16)
				store	5375.16 (5)	0.77 (16)
		0.67	1.0	up/down	5332.50 (5)	1.83 (48)
				store	5375.16 (5)	2.29 (48)
		0.50	0.1	up/down	8532.00 (8)	1.22 (32)
				store	8574.66 (8)	1.53 (32)
LeNet	Fashion-MNIST/Testset	0.86	100.0	up/down	78.00 (13)	0.19 (14)
				store	78.24 (13)	0.22 (14)
		0.85	1.0	up/down	78.00 (13)	0.11 (9)
				store	78.24 (13)	0.14 (9)
		0.77	0.1	up/down	72.00 (12)	0.16 (12)
				store	72.24 (12)	0.19 (12)
LeNet	MNIST/EMNIST	0.97	100.0	up/down	36.00 (6)	0.23 (6)
				store	36.24 (6)	0.24 (6)
		0.97	1.0	up/down	36.00 (6)	0.27 (7)
				store	36.24 (6)	0.28 (7)
		0.94	0.1	up/down	54.00 (9)	0.15 (4)
				store	54.24 (9)	0.16 (4)

We use 10′000, 12′800, and 5′000 samples respectively for CIFAR-10, MNIST and fashion-mnist as distillation data. Communication rounds for target accuracy are in parentheses.

on its underlying collateral (with $\lambda = 1$ consistently used in our experiments). Additionally, β can be set to modify the penalty for incorrect answers, thus tuning the confidence threshold required for rationally individual workers (as described in (9)) to submit a prediction. The initially staked deposit acts as a safeguard against malicious behavior, as such actions can lead to losses. In our experiment, workers can opt to withhold their reports if they lack confidence in their predictions. Fig. 6 demonstrates the reward variations with different penalty factors β across diverse confidence levels. For this experiment, the local training data was divided based on a Dirichlet distribution with parameter $\alpha = 0.1$, mimicking scenarios where workers might not have access to uniform data. As a result, certain workers' local models could be ill-equipped to predict classes previously unavailable to them. These workers will only submit their predictions if their confidence in the most probable label surpasses a certain

benchmark. Our findings indicate that, by adjusting β, PTSFD can effectively deter subpar contributions from contaminating the federated training process.

F. Robustness in Case of Malicious Behavior

Building on the game-theoretic analysis detailed in Section II-I-D, we experimentally confirm our theoretical claims in a real FL setting. We have demonstrated that both heuristic behaviors (like bypassing local training to randomly report labels on a public dataset) and strategic tactics like collusion yield an expected reward of $1-\beta$. Fig. 7 (top) depicts the reward differences between colluding and honest workers. Colluding predictions are structured as

$$Y_i^{report} = \begin{cases} 0 & \text{if } Y_i^{eval} \in \{0,1,2,3,4\} \\ 9 & \text{if } Y_i^{eval} \in \{5,6,7,8,9\} \end{cases}.$$

Conversely, Fig. 7 (bottom) contrasts the rewards of heuristic workers, who predict randomly on the public dataset, with those of diligent participants. Overall, the data indicate that genuine engagement results in the most significant rewards, even amidst prevalent malicious actions. In-depth cost considerations are reserved for future studies.

VI. Conclusion

In this work, we introduced a novel decentralized and reward-based 1-bit compressed Federated Distillation scheme on the blockchain, incorporating the Peer Truth Serum [16] specifically for Federated Distillation. The 1-bit compression ensures explicit comparability between contributions, a critical feature for automatically computing rewards on a smart contract atop a general-purpose blockchain system, where each worker is regarded as an equitable member of the federation. We have demonstrated that, in terms of storage on the blockchain and communication overhead, our framework is significantly more efficient than Federated Averaging. Additionally, the system not only offers flexibility in neural network architecture but also allows adaptation to various thresholds of contribution quality by adjusting the penalty term β. Furthermore, both theoretical insights and experimental evidence suggest our proposed mechanism is resilient to random reporting and collusion. We are confident that our findings will further the scalability of Federated Learning tasks in fully decentralized environments, where all entities have an equal interest in enhancing their models.

References

[1] L. Witt, M. Heyer, K. Toyoda, W. Samek, and D. Li, "Decentral and incentivized federated learning frameworks: A systematic literature review," *IEEE Internet Things J.*, vol. 10, no. 4, pp. 3642–3663, Feb. 2023.

[2] H. Brendan McMahan, E. Moore, D. Ramage, S. Hampson, and B. Agüera y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. 20th Int. Conf. Artif. Intell. Statist.*, 2017, pp. 1273–1282.

[3] G. Wood, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum Project Yellow Paper*, vol. 151, no. 2014, pp. 1–32, 2014.

[4] D. G. Wood, "Polkadot: Vision for a heterogeneous multi-chain framework." Accessed: Dec. 2, 2023. [Online]. Available: https://polkadot.network/PolkaDotPaper.pdf

[5] E. Androulaki et al., "Hyperledger fabric: A distributed operating system for permissioned blockchains," in *Proc. 13th EuroSys Conf.*, 2018, Art. no. 30.

[6] P. Kairouz et al., "Advances and open problems in federated learning," 2019, *arXiv:1912.04977*.

[7] T. Nishio, R. Shinkuma, and N. B. Mandayam, "Estimation of individual device contributions for incentivizing federated learning," in *Proc. IEEE Globecom Workshops*, 2020, pp. 1–6.

[8] R. D. Cook, "Detection of influential observation in linear regression," *Technometrics*, vol. 19, no. 1, pp. 15–18, Apr. 1977. [Online]. Available: http://www.jstor.org/stable/1268249

[9] X. Tu, K. Zhu, N. C. Luong, D. Niyato, Y. Zhang, and J. Li, "Incentive mechanisms for federated learning: From economic and game theoretic perspective," 2021, *arXiv:2111.11850*.

[10] R. Jia et al., "Towards efficient data valuation based on the shapley value," 2020, *arXiv:1902.10275*.

[11] J. Weng, J. Weng, J. Zhang, M. Li, Y. Zhang, and W. Luo, "DeepChain: Auditable and privacy-preserving deep learning with blockchain-based incentive," *IEEE Trans. Dependable Secure Comput.*, vol. 18, no. 5, pp. 2438–2455, Sep./Oct. 2021.

[12] M. Shayan, C. Fung, I. Beschastnikh, and C. J. Yoon, "Biscotti: A ledger for private and secure peer-to-peer machine learning," 2018, *arXiv:1811.09904*.

[13] X. Bao, C. Su, Y. Xiong, W. Huang, and Y. Hu, "FLChain: A blockchain for auditable federated learning with trust and incentive," in *Proc. 5th Int. Conf. Big Data Comput. Commun.*, 2019, pp. 151–159.

[14] Y. Liu, Z. Ai, S. Sun, S. Zhang, Z. Liu, and H. Yu, "FedCoin: A peer-to-peer payment system for federated learning," in *Lecture Notes in Computer Science*. Spriger, Cham, vol. 12500, 2020, pp. 125–138.

[15] K. Toyoda, J. Zhao, A. N. S. Zhang, and P. T. Mathiopoulos, "Blockchain-enabled federated learning with mechanism design," *IEEE Access*, vol. 8, pp. 219744–219756, 2020.

[16] G. Radanovic, B. Faltings, and R. Jurca, "Incentives for effort in crowdsourcing using the peer truth serum," *ACM Trans. Intell. Syst. Technol.*, vol. 7, no. 4, Mar. 2016, Art. no. 48.

[17] Y. LeCun, J. S. Denker, and S. A. Solla, "Optimal brain damage," in *Proc. Adv. Neural Inf. Process. Syst.*, 1990, pp. 598–605.

[18] M. Courbariaux, Y. Bengio, and J. David, "BinaryConnect: Training deep neural networks with binary weights during propagations," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 3123–3131.

[19] F. Li, B. Zhang, and B. Liu, "Ternary weight networks," 2016, *arXiv:1605.04711*.

[20] J. Konečnỳ, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," 2016, *arXiv:1610.05492*.

[21] F. Sattler, S. Wiedemann, K.-R. Müller, and W. Samek, "Sparse binary compression: Towards distributed deep learning with minimal communication," in *Proc. Int. Joint Conf. Neural Netw.*, 2019, pp. 1–8.

[22] J. Xu, W. Du, R. Cheng, W. He, and Y. Jin, "Ternary compression for communication-efficient federated learning," 2020, *arXiv:2003.03564*.

[23] F. Sattler, S. Wiedemann, K.-R. Müller, and W. Samek, "Robust and communication-efficient federated learning from non-iid data," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 31, no. 9, pp. 772–785, Sep. 2020.

[24] D. Neumann et al., "DeepCABAC: Plug & play compression of neural network weights and weight updates," in *Proc. IEEE Int. Conf. Image Process.*, 2020, pp. 21–25.

[25] S. Wiedemann et al., "DeepCABAC: A universal compression algorithm for deep neural networks," *IEEE J. Sel. Topics Signal Process.*, vol. 14, no. 4, pp. 700–714, May 2020.

[26] F. Sattler, A. Marban, R. Rischke, and W. Samek, "CFD: Communication-efficient federated distillation via soft-label quantization and delta coding," *IEEE Trans. Netw. Sci. Eng.*, vol. 9, no. 4, pp. 2025–2038, Jul./Aug. 2022.

[27] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," 2015, *arXiv:1503.02531*.

[28] E. Jeong, S. Oh, H. Kim, J. Park, M. Bennis, and S. Kim, "Communication-efficient on-device machine learning: Federated distillation and augmentation under non-IID private data," 2018, *arXiv:1811.11479*.

[29] T. Lin, L. Kong, S. U. Stich, and M. Jaggi, "Ensemble distillation for robust model fusion in federated learning," in *Proc. 34th Int. Conf. Neural Inf. Process. Syst.*, 2020, Art. no. 198.

[30] S. Itahara, T. Nishio, Y. Koda, M. Morikura, and K. Yamamoto, "Distillation-based semi-supervised federated learning for communication-efficient collaborative training with non-IID private data," 2020, *arXiv:2008.06180*.

[31] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," *Consulted*, vol. 1, pp. 1–9, 2008.
[32] S. S. Shetty, C. A. Kamhoua, and L. L. Njilla, *Distributed Consensus Protocols and Algorithms*, Hoboken, NJ, USA: Wiley-IEEE Press, 2019, pp. 25–50.
[33] V. Buterin, "Ethereum: A next-generation smart contract and decentralized application platform," *Ethereum White Paper*, vol. 3, no. 37, pp. 1–36, 2014.
[34] Y. Zhao et al., "Privacy-preserving blockchain-based federated learning for IoT devices," *IEEE Internet Things J.*, vol. 8, no. 3, pp. 1817–1829, Feb. 2021.
[35] Q. Zhang, Q. Ding, J. Zhu, and D. Li, "Blockchain empowered reliable federated learning by worker selection: A trustworthy reputation evaluation method," in *Proc. IEEE Wireless Commun. Netw. Conf. Workshops*, 2021, pp. 1–6.
[36] J. Kang, Z. Xiong, D. Niyato, S. Xie, and J. Zhang, "Incentive mechanism for reliable federated learning: A joint optimization approach to combining reputation and contract theory," *IEEE Internet Things J.*, vol. 6, no. 6, pp. 10700–10714, Dec. 2019.
[37] Z. Li, J. Liu, J. Hao, H. Wang, and M. Xian, "CrowdSFL: A secure crowd computing framework based on blockchain and federated learning," *Electronics*, vol. 9, no. 5, 2020, Art. no. 773.
[38] A. Fadaeddini, B. Majidi, and M. Eshghi, "Privacy preserved decentralized deep learning: A blockchain based solution for secure AI-driven enterprise," in *Proc. High- Perform. Comput. Big Data Anal.*, 2019, pp. 32–40.
[39] C. He, B. Xiao, X. Chen, Q. Xu, and J. Lin, "Federated learning intellectual capital platform," *Pers. Ubiquitous Comput.*, vol. 27, pp. 1525–1536, 2023.
[40] Y. Zou, F. Shen, F. Yan, J. Lin, and Y. Qiu, "Reputation-based regional federated learning for knowledge trading in blockchain-enhanced IoV," in *Proc. IEEE Wireless Commun. Netw. Conf.*, 2021, pp. 1–6.
[41] K. Toyoda and A. N. Zhang, "Mechanism design for an incentive-aware blockchain-enabled federated learning platform," in *Proc. Int. Conf. Big Data*, 2019, pp. 395–403.
[42] S. Jiang and J. Wu, "A reward response game in the blockchain-powered federated learning system," *Int. J. Parallel Emergent Distrib. Syst.*, vol. 37, no. 1, pp. 68–90, 2022.
[43] J. Weng, J. Weng, J. Zhang, M. Li, Y. Zhang, and W. Luo, "DeepChain: Auditable and privacy-preserving deep learning with blockchain-based incentive," *IEEE Trans. Dependable Secure Comput.*, vol. 18, no. 5, pp. 2438–2455, Sep./Oct. 2021.
[44] L. Feng, Z. Yang, S. Guo, X. Qiu, W. Li, and P. Yu, "Two-layered blockchain architecture for federated learning over mobile edge network," *IEEE Netw.*, vol. 36, no. 1, pp. 45–51, Jan./Feb. 2022.
[45] S. Xuan, M. Jin, X. Li, Z. Yao, W. Yang, and D. Man, "DAM-SE: A blockchain-based optimized solution for the counterattacks in the internet of federated learning systems," *Secur. Commun. Netw.*, vol. 2021, Jul. 2021, Art. no. 9965157.
[46] R. Zhang et al., "Democratic learning: Hardware/software co-design for lightweight blockchain-secured on-device machine learning," *Int. J. High Perform. Syst. Archit.*, vol. 118, Sep. 2021, Art. no. 102205.
[47] X. Qu, S. Wang, Q. Hu, and X. Cheng, "Proof of federated learning: A novel energy-recycling consensus algorithm," *IEEE Trans. Parallel Distrib. Syst.*, vol. 32, no. 8, pp. 2074–2085, Aug. 2021.
[48] Y. Wang, Z. Su, N. Zhang, and A. Benslimane, "Learning in the air: Secure federated learning for UAV-assisted crowdsensing," *IEEE Trans. Netw. Sci. Eng.*, vol. 8, no. 2, pp. 1055–1069, Second Quarter 2021.
[49] Q. Hu, Z. Wang, M. Xu, and X. Cheng, "Blockchain and federated edge learning for privacy-preserving mobile crowdsensing," *IEEE Internet Things J.*, vol. 10, no. 14, pp. 12000–12011, Jul. 2023.
[50] S. Ma, Y. Cao, and L. Xiong, "Transparent contribution evaluation for secure federated learning on blockchain," in *Proc. IEEE 37th Int. Conf. Data Eng. Workshops*, 2021, pp. 88–91.
[51] Y. Li, J. Zhang, J. Zhu, and W. Li, "HBMD-FL: Heterogeneous federated learning algorithm based on blockchain and model distillation," in *Proc. Int. Symp. Emerg. Inf. Secur. Appl.*, J. Chen, D. He, and R. Lu, eds. Cham: Switzerland: Springer Nature, 2022, pp. 145–159.
[52] Y. Li, X. Tao, X. Zhang, J. Liu, and J. Xu, "Privacy-preserved federated learning for autonomous driving," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 7, pp. 8423–8434, Jul. 2022.
[53] M. H. ur Rehman, K. Salah, E. Damiani, and D. Svetinovic, "Towards blockchain-based reputation-aware federated learning," in *Proc. IEEE Conf. Comput. Commun. Workshops*, 2020, pp. 183–188.
[54] A. Dasgupta and A. Ghosh, "Crowdsourced judgement elicitation with endogenous proficiency," in *Proc. 22nd Int. Conf. World Wide Web*, 2013, pp. 319–330.
[55] R. Jurca and B. Faltings, "Incentives for answering hypothetical questions," in *Proc. Workshop Social Comput. User Generated Content*, Jan. 2011, pp. 1–9.
[56] B. Faltings, J. J. Li, and R. Jurca, "Incentive mechanisms for community sensing," *IEEE Trans. Comput.*, vol. 63, no. 1, pp. 115–128, Jan. 2014.
[57] J. Wang, C. Jiang, H. Zhang, Y. Ren, K.-C. Chen, and L. Hanzo, "Thirty years of machine learning: The road to pareto-optimal wireless networks," *IEEE Commun. Surv. Tuts.*, vol. 22, no. 3, pp. 1472–1514, Third Quarter 2020.
[58] S. Wang et al., "Adaptive federated learning in resource constrained edge computing systems," *IEEE J. Sel. Areas Commun.*, vol. 37, no. 6, pp. 1205–1221, Jun. 2019.
[59] Z. Liu, Y. Chen, H. Yu, Y. Liu, and L. Cui, "GTG-shapley: Efficient and accurate participant contribution evaluation in federated learning," *ACM Trans. Intell. Syst. Technol.*, vol. 13, no. 4, May 2022, Art. no. 60. [Online]. Available: https://doi.org/10.1145/3501811
[60] A. Ghorbani and J. Zou, "Data shapley: Equitable valuation of data for machine learning," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 2242–2251.
[61] R. Jia et al., "Towards efficient data valuation based on the shapley value," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2019, pp. 1167–1176.
[62] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE Signal Process. Mag.*, vol. 37, no. 3, pp. 50–60, May 2020.
[63] F. Sattler, T. Korjakow, R. Rischke, and W. Samek, "FedAUX: Leveraging unlabeled auxiliary data in federated learning," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 34, no. 9, pp. 5531–5543, Sep. 2023.
[64] G. Cohen, S. Afshar, J. Tapson, and A. van Schaik, "EMNIST: An extension of MNIST to handwritten letters," 2017, *arXiv:1702.05373*.
[65] L. Deng, "The MNIST database of handwritten digit images for machine learning research," *IEEE Signal Process. Mag.*, vol. 29, no. 6, pp. 141–142, Nov. 2012.
[66] A. Krizhevsky, "Learning multiple layers of features from tiny images," Univ. Toronto, Tech Rep., pp. 1–60, 2009. [Online]. Available: https://www.cs.toronto.edu/kriz/learning-features-2009-TR.pdf
[67] A. Coates, A. Ng, and H. Lee, "An analysis of single-layer networks in unsupervised feature learning," in *Proc. 14th Int. Conf. Artif. Intell. Statist.*, Fort Lauderdale, FL, USA: PMLR, 2011, pp. 215–223. [Online]. Available: https://proceedings.mlr.press/v15/coates11a.html
[68] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms," 2017, *arXiv:1708.07747*.
[69] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. 2016 IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 770–778.
[70] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
[71] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learn. Representations*, 2017, pp. 1–15.

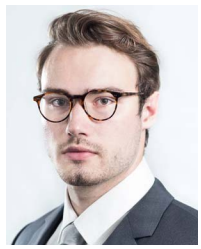

Leon Witt received the master's degree in mechanical engineering and business adminstration from RWTH Aachen, Germany, with exchange semesters in Zurich and Los Angeles, and the second master's degree in industrial engineering from Tsinghua University, in 2017. He is currently working toward the PhD degree with the Department of Computer Science and Technology, Tsinghua University in Beijing. His research interests lie with the intersection of federated artificial intelligence, blockchain and mechanism design.

Usama Zafar received the bachelors of engineering (BE) degree in software engineering from the National University of Sciences and Technology (NUST), in 2015, and the master's (MSc) degree in computer science from Tsinghua University, in 2019. He is currently working toward the PhD degree with the Department of Information Technology, Uppsala University, Sweden. His research interests include distributed machine learning, as well as security and privacy issues in federated machine learning.

KuoYeh Shen received the master's degree of mechanical engineering and MBA degree of finance from National Taiwan University. He is currently working toward the PhD degree with the Department of Computer Science and Technology, Tsinghua University in Beijing. His research interests lie with the intersection of federated artificial intelligence, blockchain and mechanism design.

Felix Sattler received the BSc degree in mathematics, in 2016, the MSc degree in computer science, in 2018, the MSc degree in applied mathematics, in 2018, all from Technische Universität Berlin, and the PhD degree with distinction from the Technical University of Berlin, in 2021. His research interests include efficient and robust machine learning, federated learning and multi-task learning.

Dan Li is currently a full professor with the Department of Computer Science and Technology, Tsinghua University. He leads the NASP (Network Architecture, System and Protocols) research group, which is part of the networking research lab. He joined the faculty of Tsinghua University, in March 2010, after two years working in the Wireless & Networking Group of Microsoft Research Asia as an associate researcher. His main research direction includes Trustworthy Internet, Data Center Network and Data-Driven Networking.

Songtao Wang received the master's degree of system on chip from Southampton University, in 2011, and the PhD degree of computer science and technology from Tsinghua University, in 2020. He is currently an assistant researcher with Zhongguancun Laboratory in Beijing, P.R.China. His research interests include distributed machine learning, data center network and data-driven networking.

Wojciech Samek (Member, IEEE) received the diploma degree in computer science from Humboldt University Berlin, in 2010, and the PhD degree from the Technical University of Berlin, in 2014. He is a professor with the Technical University of Berlin and head of the Department of Artificial Intelligence, Fraunhofer Heinrich Hertz Institute (HHI), Berlin, Germany. He is fellow with the Berlin Institute for the Foundation of Learning and Data (BIFOLD), senior editor of *IEEE Transactions on Neural Networks and Learning Systems*, and an elected member of the IEEE MLSP Technical Committee. He has been serving as an AC for NeurIPS'23, was a recipient of multiple best paper awards, including the 2020 Pattern Recognition Best Paper Award and the 2022 Digital Signal Processing Best Paper Prize. His research interests include deep learning, explainable and trustworthy AI, and federated learning.