

LSHFed: Robust and Communication-Efficient Federated Learning with Locally-Sensitive Hashing Gradient Mapping

Guanjie Cheng¹, Mengzhen Yang¹, Xinkui Zhao¹, Shuyi Yu¹, Tianyu Du¹, Yangyang Wu¹,
Mengying Zhu¹, Shuiguang Deng²

¹School of Software Technology, Zhejiang University

²College of Computer Science and Technology, Zhejiang University

{chengguanjie,mzyang1610,zhaoxinkui,yushuyi}@zju.edu.cn,tydusky@gmail.com
{zjuwuyy,mengyingzhu,dengsg}@zju.edu.cn

Abstract

Federated learning (FL) enables collaborative model training across distributed nodes without exposing raw data, but its decentralized nature makes it vulnerable in trust-deficient environments. Inference attacks may recover sensitive information from gradient updates, while poisoning attacks can degrade model performance or induce malicious behaviors. Existing defenses often suffer from high communication and computation costs, or limited detection precision. To address these issues, we propose LSHFed, a robust and communication-efficient FL framework that simultaneously enhances aggregation robustness and privacy preservation. At its core, LSHFed incorporates LSHGM, a novel gradient verification mechanism that projects high-dimensional gradients into compact binary representations via multi-hyperplane locally-sensitive hashing. This enables accurate detection and filtering of malicious gradients using only their irreversible hash forms, thus mitigating privacy leakage risks and substantially reducing transmission overhead. Extensive experiments demonstrate that LSHFed maintains high model performance even when up to 50% of participants are collusive adversaries while achieving up to a **1000×** reduction in gradient verification communication compared to full-gradient methods.

Code — <https://github.com/HaftzK/LSHFed>

Introduction

Federated learning (FL) features collaboratively training models across distributed nodes. It allows different nodes to train on distinct datasets, enabling users who are reluctant to disclose private data to participate in training without the need to publicly share or transfer their datasets.

FL has gained significant attention in both academia and industry (Han et al. 2023; Jiang et al. 2023). This growing popularity has spurred the adoption of FL in innovative solutions such as LLMs and fine-tuning methods. For example, Kuang et al. (Kuang et al. 2024) leveraged FL to collaboratively enhance pre-trained LLMs for specialized applications, all while ensuring data privacy and reducing resource consumption.

However, the distributed architecture of FL renders it vulnerable to various challenges in trust-deficient environments. Several works indicate that the advancement of FL

is predominantly impeded by three interrelated challenges: *persistent privacy-and-security vulnerabilities, pronounced heterogeneity in client-side data and prohibitive communicational overhead* (Cui et al. 2021; Liu et al. 2022; Wen et al. 2023). Among the first challenge, two types of threats stand out as particularly critical for FL systems:

(1) Poisoning attacks: Malicious nodes can disrupt the aggregation process by uploading modified or tailored gradients (Liu et al. 2024), resulting in performance degradation in the aggregated model or inducing attacker-specified behaviors for specific tasks.

(2) Inference attacks: Attackers aim to extract private training data from other participants (Liu et al. 2024). Attackers may infer whether specific content exists in training sets—or even reconstruct partial training datasets—by analyzing gradient updates during the aggregation process.

Numerous methods have been proposed to defend against these attacks, including robust aggregation (e.g., Geometric Median (Pillutla, Kakade, and Harchaoui 2022), Marginal Trimmed Mean (Nishimoto et al. 2023)), selection-based filtering (Krum, Multi-Krum (Colosimo and De Rango 2023)), and clustering or similarity-based detection using Euclidean or cosine metrics (Stallmann and Wilbik 2022; Geiping et al. 2020; Nguyen et al. 2022; Cao et al. 2020; Awan, Luo, and Li 2021). While effective, many of these schemes incur substantial computational overhead. On the privacy side, techniques such as homomorphic encryption (Aziz et al. 2023), secure MPC (Liu et al. 2024), and differential privacy (Xiong et al. 2021; Badar et al. 2024) have been employed to safeguard gradient exchanges. Concurrently, communication-efficient strategies—including sparse matrix representations (Li et al. 2023) and gradient quantization (Mai, Yan, and Pang 2024; Xu et al. 2021; Peng et al. 2023; Oh et al. 2022)—reduce transmission costs but may degrade the accuracy of malicious update identification.

To address the aforementioned challenges, we propose LSHFed—a novel robust FL framework characterized by poisoning resistance and privacy preservation. Specifically, our Locally-Sensitive Hashing Gradient Mapping (LSHGM) algorithm precisely screens for malicious gradients while reducing communication overhead to 0.07% of that required to transmit full gradients. A distributed masking strategy is designed to achieve end-to-end privacy protection.

The main contributions of this work are:

(1) We develop a communication-efficient and maliciousness-sensitive gradient detection algorithm, namely LSHGM, which enables malicious gradient identification while only transmitting 0.07% of the original gradient at minimum, ensuring both computation and communication efficiency.

(2) Building upon the LSHGM algorithm, we integrate a mask-based privacy protection mechanism and a node reputation mechanism, namely the ScoreQ-Hash election algorithm, to establish LSHFed—a robust and privacy-preserving FL framework.

(3) Comprehensive experimental evaluations verify that LSHFed remains resilient to malicious attacks in FL tasks and consistently surpasses representative baseline methods across security and utility metrics.

Related Work

Poisoning Attacks

To defend against poisoning, robust aggregation rules and filtering strategies have been developed. Geometric Median (Wang et al. 2023; Pillutla, Kakade, and Harchaoui 2022) and Marginal Trimmed Mean (Nishimoto et al. 2023) mitigate outliers by computing central statistics or discarding extremes. Krum and its variants select updates closest to the majority (Bouhata et al. 2024; Colosimo and De Rango 2023), though their complexity and sensitivity to non-IID data can hinder scalability. Clustering and similarity-based methods further filter malicious updates via Euclidean distance (Herath, Rahulamathavan, and Liu 2023) or cosine similarity (Zhu et al. 2024; Stallmann and Wilbik 2022), as exemplified by ShieldFL (Ma et al. 2022), FLAME (Nguyen et al. 2022), FLTrust (Cao et al. 2020), and Contra (Awan, Luo, and Li 2021).

Privacy Leakages

Privacy-preserving mechanisms in FL include homomorphic encryption (Aziz et al. 2023), secure MPC (Liu et al. 2024), and differential privacy (Emelianov et al. 2022; Xiong et al. 2021). Local DP offers finer granularity but may degrade utility. RFLPA (Mai, Yan, and Pang 2024) integrates verifiable Shamir secret sharing with a dot-product aggregation to balance privacy and malicious detection, while TrustFed (Badar et al. 2024) employs Gaussian noise masking under LDP. Techniques such as noise-direction alignment (Wu et al. 2025) and noise-summation correction (VerifyNet) (Xu et al. 2019) seek to mitigate accuracy loss.

Communication & Computation Optimization

Efficient aggregation leverages sparsification (Wang et al. 2020), sign-based compression (Xu et al. 2021; Seide et al. 2014), and quantization (Oh et al. 2022). While these reduce communication, they can impair detection accuracy or impose computational costs. RFLPA (Mai, Yan, and Pang 2024) further enhances overhead reduction compared to BREA (So, Güler, and Avestimehr 2020) by combining secret sharing with compressed updates.

Locally-Sensitive Hashing

LSH enables sublinear nearest-neighbor search in high dimensions (Yuan et al. 2023; Jafari et al. 2021). Recent adaptations include cross-polytope hashing for MoE token clustering (Nie et al. 2024), reducing all-to-all communication with error compensation, and K-HashFed’s integration of K-means and LSH to compress and merge similar gradient centroids (Kapoor and Kumar 2025), achieving massive transmission savings under non-IID settings.

Method

Threat Model & Design Overview

LSHFed framework and its general working process are shown in Figure 1. LSHFed divides participants into three roles: Verifier (VR), Aggregators (AG), and Local Trainers (LT). LTs are selected from nodes that have datasets and are willing to train the global model (referred to as DataNodes, DN). AGs are chosen from nodes without datasets (non-DN). We assume that VRs are trustworthy participants such as rule makers, regulatory authorities and the task publishers. Since VR carries little computation and communication overhead, it can be deployed at a low cost. Furthermore, non-DNs will report data honestly but are curious about privacy data independently. DNs are not trustworthy, meaning that DNs may conduct poisoning attack or upload fake data.

Each AG is in charge of the aggregation of a group of LTs, and the VR is in charge of all the nodes, including role election. LTs train models on their local datasets and upload gradient updates. AGs aggregate the gradient updates uploaded by LTs within their group to form a group-level gradient update and submit its LSHGM result to the VR. The VR then verifies and selects the best one to aggregate into the global model, and notify the corresponding AG to perform the aggregation and distribute the aggregated model to other active nodes for the next round of training. When the entire training process is finished, VR collects the final global model from the selected AG of the previous round.

Throughout the training process, the LSHFed framework can be divided into three main components: LSHGM Malicious Gradient Detection, the Mask-Based Privacy Protection Scheme, and the ScoreQ-Hash Role Election. The following sections will introduce these three components in order, followed by a detailed illustration of the complete training workflow.

LSHGM Malicious Gradient Detection

LSHGM is a multi-hyperplane projection and bit string encoding algorithm that significantly reduces communication and storage overhead while preserving essential gradient features for effective malicious behavior recognition. It irreversibly projects gradients onto computable binary bit strings for malicious gradient detection. The detailed mechanism is as follows:

Generate Hyperplane Vectors. Suppose the gradient update of the current FL model is represented by N matrices $\{M_1, M_2, \dots, M_N\}$. The algorithm generates r certain hyperplane vectors $H_{i,k}, i \in [1, N], k \in [1, r]$ for each gradient matrix.

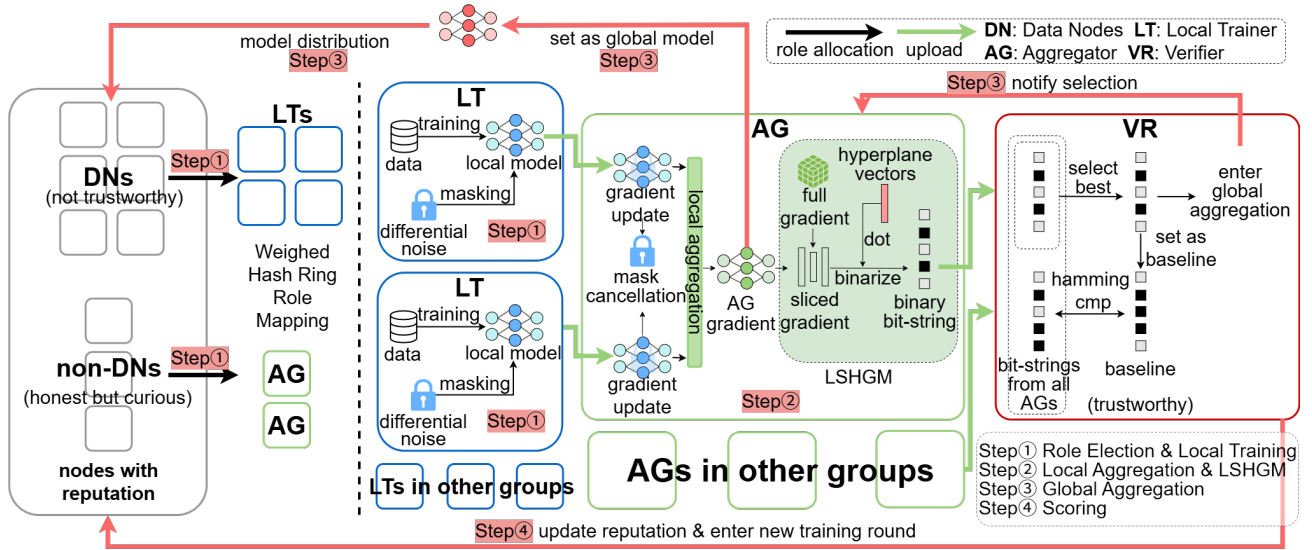


Figure 1: LSHFed Framework

Column Segmentation of the Gradient Matrix. The algorithm segments each gradient matrix $M_i, i \in [1, N]$ into multiple column vectors. This block-wise processing captures gradient features by allowing each column vector to independently reflect its characteristics in the low-dimensional representation.

Generate Bit String Results. The algorithm computes the dot product between each column vector and its corresponding hyperplane vector to generate a binary result. Assume that matrix M_i has m_i rows and n_i columns, for each segmented column vector $C_{i,j}, j \in [1, n_i]$ and hyperplane vector $H_{i,k}, k \in [1, r]$, the dot product is computed as:

$$s_{i,j,k} = H_{i,k} \cdot C_{i,j} \quad (1)$$

By applying the sign function, the dot product result $s_{i,j,k}$ is converted into a binary encoding value:

$$f(s_{i,j,k}) = \begin{cases} 1, & \text{if } s_{i,j,k} \geq 0 \\ 0, & \text{if } s_{i,j,k} < 0 \end{cases} \quad (2)$$

In this way, each column vector $C_{i,j}$ is mapped to a binary value under the hyperplane vector $H_{i,k}$, namely the generated bit strings. Bit string mapping calculation is irreversible, thus the exposure of hyperplane vectors $H_{i,k}, i \in [1, N], k \in [1, r]$ is privacy-secure.

Bit String Comparison and Anomaly Detection. By comparing the bit string generated by each AG with the benchmark bit string from the previous round, the algorithm computes their similarity to identify potential malicious behavior. Based on the gradient update result from the previous round, a benchmark bit string $B = [B_1, B_2, \dots, B_n]$ is generated using the same hyperplane vectors. In the first round of the entire training process, the LSHGM benchmark bit string needs to be obtained through local training by a trusted entity: in our experimental setup, this is

performed by the VR. Then, compute the Hamming distance between the current node's generated bit string $B' = [B'_1, B'_2, \dots, B'_n]$ and the benchmark bit string B , which can be efficiently calculated using a XOR operation. For each node, the Hamming distance between its generated bit string and the benchmark is computed and then sorted. LSHGM considers the gradient updates of nodes with a Hamming distance significantly higher than that of other nodes as malicious—in our setup, it only adopts gradients with the lowest Hamming distances.

Mask-Based Privacy Protection Scheme

Before each round of training, LTs with the same group generate their privacy masks based on a total mask sum R_{sum} . This total sum is a VR-defined constant that all nodes know. Each node's mask value is computed based on this total sum and the mask values of its preceding nodes. During aggregation, AGs subtract R_{sum} to remove all added noise.

In each training round, node n_j generates a random mask matrix $M_{j,i}$ and calculates a residual value r_i that needs to be passed on. Specifically, node n_i , based on the known sum of the masks of its preceding nodes, computes the residual value r_i to be passed to the next node using:

$$r_i = R_{\text{sum}} - \sum_{k=1}^{i-1} M_{j,k} \pmod{d} \quad (3)$$

where $\sum_{k=1}^{i-1} M_{j,k}$ represents the total mask generated by the first $i-1$ nodes, and d is the largest mask value. Node n_i uses this residual r_i^t to compute its own mask and then passes the residual value to the next node. Each node only computes the residual based on the known preceding masks without having to disclose its own generated mask.

Finally, after the last node receives the residual values from all preceding nodes, the correctness of the mask gener-

ation can be verified using:

$$\sum_{i=1}^n M_{j,i} = R_{\text{sum}} \quad \text{mod } d \quad (4)$$

ScoreQ-Hash Role Election

ScoreQ-Hash is a reputation-driven, weighted hash-ring election used to assign LT and AG roles dynamically. In the first round, P AGs are chosen uniformly at random from non-DNs, and all DNs act as LTs (suppose we have $R_{\text{non-DN}}$ non-DNs, then $P \leq R_{\text{non-DN}}$). From the second round onward, each node n_i receives a reputation score:

$$S(i) = \alpha_1 \phi_Q^{(1)}(i) + \alpha_2 \phi_Q^{(2)}(i) \quad (5)$$

where α_1 and α_2 are reputation weights, satisfying $\alpha_1 + \alpha_2 = 1$. $\phi_Q^{(1)}$ and $\phi_Q^{(2)}$ are ranked time-consumption and LSHGM-distance scores. $\phi_Q^{(1)}$ is ranked from shortest time to longest, and $\phi_Q^{(2)}$ is ranked from smallest LSHGM-distance to longest. Both scores are mapped into range $[0, 1]$ using quantile normalization. The quantile score is calculated using $\phi_Q^{(n)}(i) = \frac{R-r_i^{(n)}}{R-1}$, where $r_i^{(n)}$ is the ranking of node i in score $\phi^{(n)}$, R is the total node count. Since the reputation score is memoryless, malicious nodes that initially behave honestly will still have their selection probability promptly reduced once misbehavior is detected.

DNs and non-DNs are then mapped onto two separate hash rings: each node's arc length is proportional to its normalized weight:

$$P(i) = \frac{S(i)}{\sum_j S(j)} \quad (6)$$

A fresh random start $H_0 = \text{Hash}(\text{round_id} \parallel \text{seed})$ determines the traversal point. Moving clockwise from H_0 , the next Q DNs become LTs and the next P non-DNs become AGs (suppose we have R_{DN} DNs, then $Q \leq R_{\text{DN}}$). This mechanism ensures that higher-scoring nodes occupy larger ranges on the ring—and thus have a greater probability of selection—while low-score or malicious nodes, whose reputation falls with each detection event, are progressively excluded.

General Workflow

In this section, we illustrate the workflow of a training round in LSHFed:

Step 1: Role Election and Local Training. Before training begins, VR selects active DNs as LTs and non-DNs as AGs, which is decided using the ScoreQ-Hash role election algorithm. Meanwhile, LT groups are decided by order. LTs within the same group collaboratively generate encryption masks for gradients distributively. Subsequently, each LT computes gradient updates using its local dataset. After generating gradients, each LT applies a privacy-preserving mask $M_{j,i}$ on the original gradient $g_{j,i}$ as $\tilde{g}_{j,i} = g_{j,i} + M_{j,i}$. The masked gradient is then transmitted to the designated aggregation node AG_j^t within the group.

Step 2: Local Aggregation and LSHGM Bit String Generation. The AG node AG_j^t collects privacy-masked gradient updates from k LT nodes in the group and computes the locally aggregated gradient. Using the LSHGM algorithm, AG_j^t generates a compact bit string representation B_j^t .

Step 3: Global Aggregation. Upon receiving all bit strings from AG nodes, the VR computes the Hamming distance d_j^t between each bit string B_j^t and the previous round's chosen bit string B^{t-1} . VR ranks aggregation results by sorting Hamming distances in ascending order and select candidates within a predefined threshold as valid aggregation results. VR then notifies the selected AG to update the global model and distributes the updated model to other active nodes.

Step 4: Scoring. After completing a round of aggregation, VR scores nodes for the election in the next round.

Security Analysis

In this section, we analyze LSHFed's resilience against both privacy-and integrity-oriented threats under our assumed threat model. We show that (1) honest-majority mask generation and differential-privacy masking guarantee honest nodes' gradient confidentiality, (2) irreversible LSHGM ensures robust detection of malicious gradients, and (3) ScoreQ-Hash election prevents adversarial takeover.

(1) Gradient Privacy. Formally, let $\tilde{g}_{j,i} = g_{j,i} + M_{j,i}$ be the masked gradient; then for any coalition of up to $k < n$ LTs (with n the group size), the joint distribution of their masks is uniform over all vectors summing to at most R_{sum} , so the masked outputs $\{\tilde{g}_{j,i}\}$ remain statistically independent of the true gradients $\{g_{j,i}\}$. This achieves information-theoretic privacy of individual updates under the honest-majority assumption.

(2) Irreversibility of LSHGM. In LSHGM—where each gradient matrix is projected onto given hyperplanes and quantized to a bit string—no adversary (even a curious AG) can invert the bit string to recover information about the underlying gradient. Thus, our scheme provably bounds per-round privacy loss while providing provable non-invertibility of any transmitted summary.

(3) Robustness. Under the assumption that malicious LTs can control at most a fraction $\alpha < \frac{1}{2}$ of participants, classical results on Hamming-distance-based outlier detection guarantee that the honest updates cluster around the previous-round benchmark bit string, while malicious outliers incur significantly larger distances. Sorting distances and selecting the lowest- $(1 - \alpha)$ fraction provably excludes all adversarial updates whenever the distance gap between honest and malicious clusters exceeds a threshold τ ; in practice, we observe this gap to be large under common poisoning norms. Hence, LSHFed's global aggregation resists up to α Byzantine faults with high probability.

(4) Secure Role Election. ScoreQ-Hash assigns selection weights proportional to a combination of past reliability and LSHGM ranks. Since the hash-ring injects fresh randomness each round ($H_0 = \text{Hash}(\text{round_id} \parallel \text{seed})$), no coalition of compromised non-DNs can bias the choice of VR beyond

their proportional reputation share. As long as a majority of high-reputation nodes remain honest, the probability that an attacker controls all AG groups in any round is negligible.

(5) Mask-Integrity Robustness. A malicious LT that uploads incorrect privacy mask introduces an irrecoverable residual noise $\delta = \sum_{i \in \mathcal{G}} m_{j,i}^t - R_{\text{sum}} \neq 0$ into its group’s aggregated gradient. Since AGs subtract only R_{sum} , the noisy aggregate becomes $\tilde{G} + \delta$. Under LSHGM, this mask-induced perturbation is the same as noise injected by an attacker, yielding a bit string with a Hamming distance greater than that of honest gradients and thus filtered out as an outlier.

Consequently, the corresponding node’s gradient score $\phi_Q^{(2)}(i) = \frac{N-r_i^{(2)}}{N-1}$ collapses to its minimum quantile (where $r_i^{(2)}$ is its rank among N nodes), and its overall reputation drops. This in turn reduces its selection weight. Because ScoreQ-Hash re-elects roles with fresh randomness each round, the malicious node cannot hide in a stable honest group. Repeated mask tampering thus drives $P(i) \rightarrow 0$, making the attacker increasingly unlikely to be chosen as an LT and effectively excluding it from future rounds.

Combining the above guarantees, LSHFed achieves:

- **Confidentiality:** No individual gradient is revealed to any other party beyond what is exposed by the final, noise-free global model.
- **Robustness:** Malicious gradient injections or model-poisoning attempts are detected and excluded in every round, ensuring convergence to a clean model.
- **Liveness:** Honest nodes are guaranteed eventual participation in both LT and AG roles, so progress is not stalled by adversarial denial-of-service.

Experiments

Experimental Setup

We assess the performance and robustness of LSHFed in image classification tasks, evaluating LSHFed with five baseline FL frameworks on three public datasets. The detailed experiment configuration is shown in 1. The MNIST dataset is trained using a CNN model with two convolution layers, one pooling layer and two full connect layers, and CIFAR-10 is trained using a ResNet model, identical to the setup of RFLPA (Mai, Yan, and Pang 2024). We also evaluate LSHFed on the F-MNIST dataset using a vision transformer (ViT) and non-iid setups to fully assess its performance and scalability. To maximize the malicious gradient detection capability of LSHFed, we configure the ScoreQ-Hash parameters as $\alpha_1 = 0$ and $\alpha_2 = 1$ in experiments.

The experiments are conducted on a 24-core Ubuntu 20.04 server with an RTX4090, 128GB RAM, and a 20-core Windows laptop with an RTX4070, 32GB RAM. All code is implemented in Python 3.9.

Malicious Node Configuration: The proportion of malicious nodes is fixed at 0%, 10%, 30%, 40% and 50% on both the MNIST and CIFAR-10 datasets. The model accuracy is recorded over 50 training rounds to monitor its evolution. Malicious nodes are assumed to upload modified gradients dishonestly.

Param	MNIST	CIFAR	FMNIST
Number of Classes	10	10	10
Model Type	CNN	ResNet	ViT
Parameter Amount	206,922	6,575,370	805,130
Model Structure	c-c-p-f-f	c-c-r-c-c-r	e-n-h
DN Count (Total/LT)	10 / 5	10 / 5	10 / 5
non-DN Count (Total/AG/VR)	10 / 2 / 1	10 / 2 / 1	10 / 2 / 1
LSHGM Filter Rank	1	1	1

Table 1: Experiment Configuration.

*Model Structure: *c*=Convolution, *f*=Fully Connected, *r*=ResNet, *p*=Pooling, *e*=Transformer Encoder, *n*=LayerNorm, *h*=Linear. CIFAR dataset is CIFAR-10.

Attack Types: Attacks are categorized into *label flipping* (targeted) and *gradient manipulation* (untargeted) attacks. Targeted attacks flip the label from l to $N_l - l - 1$, where N_l is the number of classes. Conversely, untargeted attacks replaces gradient updates with generated gaussian noise.

Preliminary Experiment




Figure 2: LSHGM-Euclidean Distance Correlation

In the preliminary experiment, we explore the linear relationship between the number of differences produced by LSHGM and the Euclidean distance when calculating the same group of gradients. The experimental results demonstrates approximate linear correlation between the LSHGM distance and the Euclidean distance ($linearcorrelation > 0.98$), as shown in Figure 2. The result means LSHGM provides detection performance comparable to that of Euclidean distance.

Aggregation Performance

To evaluate the overall framework’s resistance to malicious attacks, robustness, and aggregation model accuracy, experiments are conducted upon LSHFed. In targeted attacks, all malicious nodes act in collusion; in untargeted attacks, they operate independently.

Table 2 exhibits the aggregation performance of LSHFed compared to a series of baseline methods under varying proportions of malicious clients. The final accuracy is measured as the average accuracy over the last 10 training epochs. The baselines include RFLPA (Mai, Yan, and Pang 2024), FedAvg, FedATM (Nishimoto et al. 2023), CKADA (Zhu et al.

Dataset	Method	Untargeted-attack					Targeted-attack					Control group
		10%	20%	30%	40%	50%	10%	20%	30%	40%	50%	
MNIST (CNN)	FedAvg	0.89	0.80	0.71	0.63	0.53	0.88	0.78	0.68	0.59	0.49	0.98
	FedATM (Nishimoto et al. 2023)	0.63	0.65	0.76	0.82	0.82	0.58	0.60	0.53	0.45	0.39	0.98
	CKADA (Zhu et al. 2024)	0.88	0.78	0.68	0.67	0.55	0.88	0.80	0.70	0.59	0.55	0.98
	Krum (Colosimo and De Rango 2023)	0.98	0.97	0.98	0.98	0.83	0.90	0.75	0.85	0.57	0.47	0.98
	Median (Xie et al. 2022)	0.87	0.84	0.85	0.78	0.12	0.80	0.73	0.65	0.55	0.44	0.96
	RFLPA (Mai, Yan, and Pang 2024)	0.96	0.95	0.95	-	-	0.96	0.95	0.95	-	-	0.96
	LSHFed (proposed)	0.98	0.98	0.97	0.97	0.97	0.97	0.97	0.97	0.97	0.98	
CIFAR-10 (ResNet)	FedAvg	0.91	0.82	0.73	0.64	0.50	0.90	0.80	0.70	0.60	0.50	0.99
	FedATM (Nishimoto et al. 2023)	0.12	0.12	0.17	0.14	0.18	0.10	0.11	0.10	0.18	0.17	0.99
	CKADA (Zhu et al. 2024)	0.89	0.83	0.77	0.64	0.47	0.90	0.81	0.70	0.58	0.53	0.99
	Krum (Colosimo and De Rango 2023)	0.99	0.99	0.99	0.99	0.58	0.85	0.85	0.73	0.63	0.48	0.99
	Median (Xie et al. 2022)	0.58	0.54	0.51	0.46	0.22	0.64	0.59	0.51	0.44	0.38	0.71
	RFLPA (Mai, Yan, and Pang 2024)	0.70	0.70	0.69	-	-	0.71	0.70	0.69	-	-	0.74
	LSHFed (proposed)	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	
F-MNIST (ViT)	FedAvg	0.77	0.69	0.62	0.55	0.47	0.88	0.80	0.70	0.59	0.50	0.88
	FedATM (Nishimoto et al. 2023)	0.73	0.67	0.60	0.51	0.41	0.72	0.64	0.56	0.48	0.40	0.88
	CKADA (Zhu et al. 2024)	0.76	0.67	0.62	0.56	0.41	0.78	0.70	0.54	0.55	0.46	0.88
	Krum (Colosimo and De Rango 2023)	0.85	0.85	0.85	0.85	0.65	0.72	0.76	0.64	0.57	0.41	0.87
	Median (Xie et al. 2022)	0.62	0.57	0.52	0.47	0.41	0.61	0.54	0.47	0.40	0.34	0.75
		LSHFed (proposed)	0.87	0.88	0.87	0.87	0.88	0.88	0.87	0.88	0.87	0.88

Table 2: Results of Aggregation Performance against Malicious Attacks

* The control group refers to the scenario where all participating nodes are honest, i.e., no adversarial behavior is introduced during training. This serves as an upper-bound reference for evaluating the robustness of each method under attack.




Figure 3: Per-iteration Model Performance

2024), Krum (Colosimo and De Rango 2023), and Median (Xie et al. 2022). Notably, RFLPA only supports up to 30% malicious clients and was not evaluated on the ViT-based F-MNIST setting; therefore, its results are unavailable for the 40% and 50% and F-MNIST settings.

As shown in Table 2, LSHFed consistently achieves the highest overall accuracy across all datasets and attack scenarios, and even when 50% of the clients are adversarial it maintains performance on par with the benign (0%) condition. This resilience stems from LSHFed’s ability to selectively aggregate only non-malicious updates, thereby preserving near-no-attack accuracy until the system’s theoretical failure point. In contrast, while Krum also demonstrates relatively strong accuracy at lower adversary fractions, it tends to collapse once the proportion of malicious participants exceeds 40%.

Non-IID Robustness

We also evaluate the performance of LSHFed in non-iid settings, as presented in Table 4. Two types of data heterogeneity are considered: *label skew* and *Dirichlet* distributions. The experiments are conducted on the F-MNIST dataset using a ViT model. The model configuration is consistent with that listed in Table 1. In the label skew scenario, each client receives 6000 training samples, where 70% of the labels are drawn from two common primary classes shared across clients, and the remaining 30% are randomly sampled from the other eight classes. In the Dirichlet setting, the training data is allocated among clients based on a Dirichlet distribution, which models varying degrees of label distribution skewness. Specifically, for each class, samples are distributed across three clients using Dirichlet concentration parameter $\alpha = 0.5$.

Method	Client computation (LT & AG)	Communication per client	Server computation (VR)	Server Comm. (first & last round)	Server Comm. (other rounds)
RFLPA	$\mathcal{O}((R^2 + N) \log^2 R)$	$\mathcal{O}(R + N)$	$\mathcal{O}((R + N) \log^2 R \log \log R)$	$\mathcal{O}((R + N) R)$	$\mathcal{O}((R + N) R)$
LSHFed	$\mathcal{O}(R_{AG} \sum_{i=1}^N n_i r) \approx \mathcal{O}(R)$	$\mathcal{O}(M + \sum_{i=1}^N n_i r) \approx \mathcal{O}(R + N)$	$\mathcal{O}(R_{AG} \sum_{i=1}^N n_i r) \approx \mathcal{O}(R)$	$\mathcal{O}(R \times M) \approx \mathcal{O}(RN)$	$\mathcal{O}(R_{AG} \sum_{i=1}^N n_i r) \approx \mathcal{O}(R)$

Table 3: Communication and Computation Comparison between RFLPA and LSHFed

Data distribution	Targeted-attack					Control group
	10%	20%	30%	40%	50%	
Dirichlet	0.91	0.91	0.91	0.91	0.91	0.91
Label skew partition	0.88	0.88	0.87	0.87	0.87	0.88

Table 4: LSHFed in Non-IID Scenario

Scalability

To assess the scalability of LSHFed, we further evaluate its performance under varying numbers of participating clients. The results are summarized in Table 5. The experiments are conducted on the F-MNIST dataset using a ViT model. The model configuration is consistent with that listed in Table 1.

Node Count	Accuracy in 50% Targeted-attack Scenario	Control Group
10 Nodes	0.88	0.88
20 Nodes	0.88	0.88
100 Nodes	0.85	0.88

Table 5: LSHFed Scalability

Communication & Computation Overhead

In this section, we provide both theoretical evaluations of the communication and computation overhead of LSHFed. Table 3 presents a theoretical comparison between LSHFed and RFLPA in terms of per-round communication and computation costs. Assume each client’s gradient update comprises N matrices and M parameters, and its LSH bit-string has length proportional to $\sum_{i=1}^N n_i \times r$ (with r hyperplanes). Let R be the total number of nodes, decomposed as R_{LT} local trainers and R_{AG} aggregators, and $R_{VR} = 1$ verifier.

Conclusion

We introduce a novel FL framework, LSHFed, which aims at improving the robustness of the aggregation process at a low cost. LSHFed enhances the aggregation process by incorporating a mask-based privacy protection scheme and the Score-Q Hash role election algorithm, effectively detecting and excluding malicious nodes while protecting gradient privacy. The core module in LSHFed, namely LSHGM, employs a multi-hyperplane projection technique to irreversibly map high-dimensional gradients into compact binary bit strings. By measuring the Hamming distance of the bit strings, LSHGM offers a communication- and computation-efficient method for identifying malicious gradients. Experimental results confirm that LSHFed outperforms baseline robust FL frameworks, maintaining high model accuracy even when up to 50% of participating nodes are malicious and reducing the communication overhead of gradient veri-

fication by at most three orders of magnitude than using full gradients.

Acknowledgments

This work was supported in part by NSFC grant 62502446, in part by the Ningbo Yongjiang Talent Programme Grant 2024A-402-G, in part by the the National Key R&D Program of China under Grant 2024YFC3308304, and in part by the fund of Oak Grove Ventures-School of Software Technology, Zhejiang University Blockchain Joint Lab.

References

- Awan, S.; Luo, B.; and Li, F. 2021. Contra: Defending against poisoning attacks in federated learning. In *Computer Security—ESORICS 2021: 26th European Symposium on Research in Computer Security, Darmstadt, Germany, October 4–8, 2021, Proceedings, Part I* 26, 455–475. Springer.
- Aziz, R.; Banerjee, S.; Bouzefrane, S.; and Le Vinh, T. 2023. Exploring homomorphic encryption and differential privacy techniques towards secure federated learning paradigm. *Future internet*, 15(9): 310.
- Badar, M.; Sikdar, S.; Nejdil, W.; and Fisichella, M. 2024. TrustFed: Navigating Trade-offs Between Performance, Fairness, and Privacy in Federated Learning. In *ECAI 2024*, 2370–2377. IOS Press.
- Bouhata, D.; Moumen, H.; Mazari, J. A.; and Bounceur, A. 2024. Byzantine fault tolerance in distributed machine learning: a survey. *Journal of Experimental & Theoretical Artificial Intelligence*, 1–59.
- Cao, X.; Fang, M.; Liu, J.; and Gong, N. Z. 2020. Fltrust: Byzantine-robust federated learning via trust bootstrapping. *arXiv preprint arXiv:2012.13995*.
- Colosimo, F.; and De Rango, F. 2023. Median-krum: A joint distance-statistical based byzantine-robust algorithm in federated learning. In *Proceedings of the Int’l ACM Symposium on Mobility Management and Wireless Access*, 61–68.
- Cui, Z.; Zhao, Y.; Cao, Y.; Cai, X.; Zhang, W.; and Chen, J. 2021. Malicious code detection under 5G HetNets based on a multi-objective RBM model. *IEEE network*, 35(2): 82–87.
- Emelianov, V.; Gast, N.; Gummadi, K. P.; and Loiseau, P. 2022. On fair selection in the presence of implicit and differential variance. *Artificial Intelligence*, 302: 103609.
- Geiping, J.; Bauermeister, H.; Dröge, H.; and Moeller, M. 2020. Inverting gradients-how easy is it to break privacy in federated learning? *Advances in neural information processing systems*, 33: 16937–16947.
- Han, S.; Ding, H.; Zhao, S.; Ren, S.; Wang, Z.; Lin, J.; and Zhou, S. 2023. Practical and robust federated learning with highly scalable regression training. *IEEE Transactions on Neural Networks and Learning Systems*.

- Herath, C.; Rahulamathavan, Y.; and Liu, X. 2023. Recursive Euclidean distance-based robust aggregation technique for federated learning. In *2023 IEEE IAS Global Conference on Emerging Technologies (GlobConET)*, 1–6. IEEE.
- Jafari, O.; Maurya, P.; Nagarkar, P.; Islam, K. M.; and Crushev, C. 2021. A survey on locality sensitive hashing algorithms and their applications. *arXiv preprint arXiv:2102.08942*.
- Jiang, M.; Roth, H. R.; Li, W.; Yang, D.; Zhao, C.; Nath, V.; Xu, D.; Dou, Q.; and Xu, Z. 2023. Fair federated medical image segmentation via client contribution estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16302–16311.
- Kapoor, A.; and Kumar, D. 2025. K-HashFed: Communication Efficient Federated Learning through Gradient Clustering and Hashing. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Kuang, W.; Qian, B.; Li, Z.; Chen, D.; Gao, D.; Pan, X.; Xie, Y.; Li, Y.; Ding, B.; and Zhou, J. 2024. Federatedscope-llm: A comprehensive package for fine-tuning large language models in federated learning. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 5260–5271.
- Li, Y.; Yu, Y.; Zhang, Q.; Liang, C.; He, P.; Chen, W.; and Zhao, T. 2023. Lospars: Structured compression of large language models based on low-rank and sparse approximation. In *International Conference on Machine Learning*, 20336–20350. PMLR.
- Liu, F.; Zheng, Z.; Shi, Y.; Tong, Y.; and Zhang, Y. 2024. A survey on federated learning: a perspective from multi-party computation. *Frontiers of Computer Science*, 18(1): 181336.
- Liu, J.; Huang, J.; Zhou, Y.; Li, X.; Ji, S.; Xiong, H.; and Dou, D. 2022. From distributed machine learning to federated learning: A survey. *Knowledge and Information Systems*, 64(4): 885–917.
- Ma, Z.; Ma, J.; Miao, Y.; Li, Y.; and Deng, R. H. 2022. ShieldFL: Mitigating model poisoning attacks in privacy-preserving federated learning. *IEEE Transactions on Information Forensics and Security*, 17: 1639–1654.
- Mai, P.; Yan, R.; and Pang, Y. 2024. Rflpa: A robust federated learning framework against poisoning attacks with secure aggregation. *Advances in Neural Information Processing Systems*, 37: 104329–104356.
- Nguyen, T. D.; Rieger, P.; De Viti, R.; Chen, H.; Brandenburg, B. B.; Yalame, H.; Möllering, H.; Fereidooni, H.; Marchal, S.; Miettinen, M.; et al. 2022. {FLAME}: Taming backdoors in federated learning. In *31st USENIX Security Symposium (USENIX Security 22)*, 1415–1432.
- Nie, X.; Qibin, L.; Fu, F.; Zhu, S.; Miao, X.; Li, X.; Zhang, Y.; Liu, S.; and Cui, B. 2024. LSH-MoE: Communication-efficient MoE Training via Locality-Sensitive Hashing. *Advances in Neural Information Processing Systems*, 37: 54161–54182.
- Nishimoto, K.; Chiang, Y.-H.; Lin, H.; and Ji, Y. 2023. FedATM: Adaptive Trimmed Mean based Federated Learning against Model Poisoning Attacks. In *2023 IEEE 97th Vehicular Technology Conference (VTC2023-Spring)*, 1–5. IEEE.
- Oh, Y.; Lee, N.; Jeon, Y.-S.; and Poor, H. V. 2022. Communication-efficient federated learning via quantized compressed sensing. *IEEE Transactions on Wireless Communications*, 22(2): 1087–1100.
- Peng, H.; Qin, S.; Yu, Y.; Wang, J.; Wang, H.; and Li, G. 2023. Birder: communication-efficient 1-bit adaptive optimizer for practical distributed DNN training. *Advances in Neural Information Processing Systems*, 36: 39529–39552.
- Pillutla, K.; Kakade, S. M.; and Harchaoui, Z. 2022. Robust aggregation for federated learning. *IEEE Transactions on Signal Processing*, 70: 1142–1154.
- Seide, F.; Fu, H.; Droppo, J.; Li, G.; and Yu, D. 2014. 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech DNNs. In *Interspeech*, volume 2014, 1058–1062. Singapore.
- So, J.; Güler, B.; and Avestimehr, A. S. 2020. Byzantine-resilient secure federated learning. *IEEE Journal on Selected Areas in Communications*, 39(7): 2168–2181.
- Stallmann, M.; and Wilbik, A. 2022. Towards Federated Clustering: A Federated Fuzzy *c*-Means Algorithm (FFCM). *arXiv preprint arXiv:2201.07316*.
- Wang, H.; Sievert, S.; Charles, Z.; and Papailiopoulos, D. 2020. Adaptive gradient compression for efficient distributed learning. *Advances in Neural Information Processing Systems*, 33: 4177–4189.
- Wang, X.; Zhang, H.; Bilal, A.; Long, H.; and Liu, X. 2023. WGM-dSAGA: Federated learning strategies with Byzantine robustness based on weighted geometric median. *Electronics*, 12(5): 1190.
- Wen, J.; Zhang, Z.; Lan, Y.; Cui, Z.; Cai, J.; and Zhang, W. 2023. A survey on federated learning: challenges and applications. *International Journal of Machine Learning and Cybernetics*, 14(2): 513–535.
- Wu, X.; Chen, Y.; Yu, H.; and Yang, Z. 2025. Privacy-preserving federated learning based on noise addition. *Expert Systems with Applications*, 267: 126228.
- Xie, M.; Ma, J.; Long, G.; and Zhang, C. 2022. Robust clustered federated learning with bootstrap median-of-means. In *Asia-Pacific Web (APWeb) and Web-Age Information Management (WAIM) Joint International Conference on Web and Big Data*, 237–250. Springer.
- Xiong, Z.; Cai, Z.; Takabi, D.; and Li, W. 2021. Privacy threat and defense for federated learning with non-iid data in AIoT. *IEEE Transactions on Industrial Informatics*, 18(2): 1310–1321.
- Xu, G.; Li, H.; Liu, S.; Yang, K.; and Lin, X. 2019. VerifyNet: Secure and verifiable federated learning. *IEEE Transactions on Information Forensics and Security*, 15: 911–926.
- Xu, J.; Huang, S.-L.; Song, L.; and Lan, T. 2021. Signguard: Byzantine-robust federated learning through collaborative malicious gradient filtering. *arXiv preprint arXiv:2109.05872*, 5.
- Yuan, C.; Liu, M.; Luo, Y.; and Chen, C. 2023. Recent advances in locality-sensitive hashing and its performance in

different applications. In *3rd International Conference on Internet of Things and Smart City (IoTSC 2023)*, volume 12708, 268–277. SPIE.

Zhu, T.; Guo, Z.; Yao, C.; Tan, J.; Dou, S.; Wang, W.; and Han, Z. 2024. Byzantine-robust federated learning via cosine similarity aggregation. *Computer Networks*, 254: 110730.