2025 28th International Conference on Computer and Information Technology (ICCIT)
19-21 December 2025, Cox's Bazar, Bangladesh

PrivEraserVerify: Efficient, Private, and Verifiable Federated Unlearning

Parthaw Goswami
Department of Electronics and Communication Engineering
Khulna University of Engineering & Technology
Khulna-9203, Bangladesh
Email: parthawgoswami555@gmail.com

Md Khairul Islam
Department of Mathematics and Computer Science
Hobart and William Smith Colleges
Geneva, NY, USA
Email: khairul.robotics@gmail.com

Ashfak Yeafi
Department of Electrical and Electronic Engineering
Khulna University of Engineering & Technology
Khulna-9203, Bangladesh
Email: yeafiashfak@gmail.com

2025 28th International Conference on Computer and Information Technology (ICCIT) | 979-8-3315-7867-1/25/$31.00 ©2025 IEEE | DOI: 10.1109/ICCIT68739.2025.11491311

***Abstract*—Federated learning (FL) enables collaborative model training without sharing raw data, offering a promising path toward privacy-preserving artificial intelligence. However, FL models may still memorize sensitive information from participants, conflicting with the right to be forgotten (RTBF). To meet these requirements, federated unlearning has emerged as a mechanism to remove the contribution of departing clients. Existing solutions only partially address this challenge: FedEraser improves efficiency but lacks privacy protection, FedRecovery ensures differential privacy (DP) but degrades accuracy, and VeriFi enables verifiability but introduces overhead without efficiency or privacy guarantees. We present PrivEraserVerify (PEV), a unified framework that integrates efficiency, privacy, and verifiability into federated unlearning. PEV employs (i) adaptive checkpointing to retain critical historical updates for fast reconstruction, (ii) layer-adaptive differentially private calibration to selectively remove client influence while minimizing accuracy loss, and (iii) fingerprint-based verification, enabling participants to confirm unlearning in a decentralized and non-invasive manner. Experiments on image, handwritten character, and medical datasets show that PEV achieves up to 2–3× faster unlearning than retraining, provides formal indistinguishability guarantees with reduced performance degradation, and supports scalable verification. To the best of our knowledge, PEV is the first framework to simultaneously deliver efficiency, privacy, and verifiability for federated unlearning, moving FL closer to practical and regulation-compliant deployment.**

***Index Terms*—Federated Learning, Machine Unlearning, Differential Privacy, Verifiable Unlearning**

I. Introduction

The proliferation of data-driven artificial intelligence (AI) has been fueled by the increasing availability of large-scale personal data. While this growth has enabled remarkable advances in healthcare, finance, natural language processing, and computer vision [1]–[3], it also raises serious privacy and security concerns. To mitigate risks from centralized data collection, FL has emerged as a distributed paradigm that enables multiple clients to collaboratively train a global model without sharing raw data. In FL, clients compute local updates on their private datasets and only communicate model parameters or gradients to a central server for aggregation. This design preserves data locality and reduces privacy leakage compared to traditional centralized training.

Despite these advantages, FL does not eliminate the possibility of sensitive information being memorized by the global model. Recent regulations such as the General Data Protection Regulation (GDPR) in the European Union and the California Consumer Privacy Act (CCPA) explicitly recognize the RTBF, requiring organizations to delete personal data upon request. In the context of FL, this translates to the need for federated unlearning (removing a client's contribution from the global model while maintaining its utility for other participants). Achieving effective federated unlearning is challenging due to the iterative aggregation process in FL, where each client's updates are intertwined with others, making direct removal non-trivial.

In this paper, we propose PEV, a novel federated unlearning framework that introduces three core innovations: (i) an efficient mechanism that selectively retains critical historical updates to enable fast model reconstruction without full retraining, (ii) a targeted unlearning mechanism that selectively injects Gaussian noise into sensitive update directions, ensuring statistical indistinguishability, and (iii) a lightweight and non-invasive verification scheme that empowers participants to collaboratively confirm the unlearning effect without relying on backdoors or intrusive access.

Through extensive experiments on benchmark datasets spanning image, handwritten character, and medical domains, we demonstrate that PEV achieves up to 2–3× faster unlearning than retraining, provides rigorous privacy guarantees with reduced performance degradation compared to DP-only baselines, and supports scalable verification. The main contributions of this work are summarized as follows:

1) We identify the limitations of existing federated unlearning methods and argue for an integrated approach combining efficiency, privacy, and verifiability.

2) We propose PEV, the first unified federated unlearning framework that incorporates adaptive checkpointing, layer-adaptive DP calibration, and fingerprint-based verification.
3) We provide theoretical and empirical analysis demonstrating that PEV achieves formal indistinguishability guarantees, efficient reconstruction, and decentralized verification.
4) We conduct extensive experiments across heterogeneous FL settings, showing that PEV significantly outperforms existing baselines in efficiency, privacy–utility tradeoff, and verifiability.

By holistically addressing efficiency, privacy, and verifiability, PEV represents a significant step toward trustworthy, regulation-compliant federated unlearning and lays the foundation for practical adoption in real-world FL systems. The structure of the paper is as follows: Section II examines recent progress in federated unlearning. Section III delineates the materials and methodologies employed. Section IV delineates the experiments and analyzes the results. Ultimately, Section V conclude the final observations.

II. Related work

The concept of machine unlearning was first introduced to allow trained models to "forget" specific data points, either due to privacy regulations or data quality concerns. Bourtoule et al. proposed SISA training, where data is partitioned into shards and retraining can be invoked on only affected shards, improving efficiency compared to full retraining [4]. Ginart et al. studied unlearning in clustering models, developing deletion algorithms for k-means that ensure distributional equivalence with retrained models [5]. Golatkar et al. used influence functions and Hessian approximations to estimate and subtract the effect of specific training points [6]. While effective in centralized machine learning, these methods often rely on direct access to training data or assume convexity, making them unsuitable for the FL paradigm.

FL enables collaborative model training without raw data sharing, introduced by McMahan et al. through the FedAvg algorithm [7]. However, once a model has been trained, it remains unclear how to remove the contribution of a single client. To address this, FedEraser (Liu et al.) [8] presented a systematic solution for federated unlearning, enabling efficient client-level data removal by storing historical updates and applying calibration during reconstruction. FedEraser achieves up to 4× speedup compared to retraining, but it does not provide formal privacy guarantees. Building on this, FedRecovery (Zhang et al.) [9] leveraged DP to formally bound the difference between an unlearned model and a retrained one. It introduced the notion of gradient residuals and applied Gaussian noise to ensure approximate indistinguishability. FedRecovery avoids retraining but suffers from accuracy degradation due to uniform noise injection. Most recently, VeriFi (Gao et al.) [10] argued that unlearning alone is insufficient without verifiability, introducing the Right to Verify (RTV) alongside the RTBF. VeriFi proposed a general framework combining multiple unlearning and verification methods, including a novel uS2U scaling-based unlearning algorithm and lightweight verification strategies (vFM, vEM). While pioneering in verification, VeriFi does not address efficiency or privacy guarantees.

DP has become a cornerstone for ensuring formal privacy guarantees in machine learning [11]. Applied to unlearning, DP ensures that the difference between a model trained with or without specific data is statistically indistinguishable [12]. Guo et al. proposed ε-certified removal, a relaxation of DP applied to unlearning tasks [13]. However, existing DP-based methods often rely on convex loss assumptions or inject global noise that reduces model utility. FedRecovery is one of the first works to extend DP-based unlearning into federated settings, but its global noise design remains a limitation.

Verification of unlearning is an emerging challenge. In centralized machine learning, verification has been studied through performance evaluation on backdoored data [6], watermarking [14], or encryption-based proofs [15]. However, these methods are either invasive (introducing security risks via backdoors) or computationally expensive. VeriFi is one of the pioneering work to extend this concept to FL, granting participants the RTV and introducing marker-based verification. Yet, challenges remain in balancing scalability, non-invasiveness, and trustworthiness.

However, no existing approach simultaneously addresses efficiency, privacy, and verifiability. This motivates our proposed PEV framework, which integrates the strengths of these approaches into a unified design, achieving efficient, private, and verifiable federated unlearning.

III. Material and methods

A. System Model and Problem Definition

We consider a typical FL system consisting of a central server and a set of participating clients, denoted as $C = \{c_1, c_2, \ldots, c_n\}$. Each client c_i owns a private dataset D_i that remains stored locally and is never directly shared with the server or with other clients. Instead, learning proceeds through a sequence of communication rounds coordinated by the central server.

At the beginning of each round t, the server distributes the current global model parameters w_t to a subset of available clients. Each selected client then performs local training on its dataset D_i for a fixed number of epochs, producing a model update Δw_i^t that captures the knowledge gained from its private data. These updates are transmitted back to the server, which aggregates them commonly using the FedAvg algorithm [7] to obtain the updated global model w_{t+1}. Through multiple rounds of this iterative process, the global model gradually converges while ensuring that raw data remains decentralized. While this design preserves data locality and offers privacy advantages compared to centralized training, it introduces a new challenge when considering the RTBF. Specifically, if a client c_u requests the deletion of its contribution, the server must construct an unlearned model $w^{\setminus u}$ that effectively removes the influence of c_u's data. Achieving this goal requires

overcoming several obstacles.

First, efficiency is critical. A naive solution is to retrain the global model from scratch without the target client's data, but this is computationally prohibitive in large-scale FL systems, particularly when training involves many clients and multiple communication rounds. Second, privacy must be guaranteed. The unlearned model should be statistically indistinguishable from a model that was retrained without the client's data, which can be formalized through the lens of DP. Without such guarantees, subtle traces of the client's contribution could remain, leaving the model vulnerable to inference attacks. Finally, verifiability is equally important. It is not sufficient for the server to claim that unlearning has been performed; the requesting client, and potentially other participants, should be able to independently verify that the unlearning effect has indeed taken place. This requirement aligns with the emerging concept of the RTV, which complements RTBF by ensuring trust in the unlearning process.

B. Proposed Framework: PrivEraserVerify (PEV)

To address the challenges of efficiency, privacy, and verifiability in federated unlearning, we propose PEV, a unified framework that integrates three complementary modules. Each module is designed to overcome a specific limitation of existing approaches while working in synergy to deliver an efficient, private, and trustworthy unlearning process.

The first module is Adaptive Checkpointing, which targets the efficiency challenge. In conventional approaches such as retraining, the removal of a client's data requires repeating the entire training process from scratch, which is computationally expensive and impractical in large-scale FL systems. PEV mitigates this problem by selectively storing compressed checkpoints of aggregated model updates during training. Instead of recording every update, the server uses gradient variance as a criterion to decide when to retain a checkpoint. By capturing only the most informative states, this strategy strikes a balance between storage overhead and reconstruction accuracy, enabling efficient model reconstruction when unlearning requests are made.

The second module, Layer-Adaptive Differential Privacy Calibration, addresses privacy. Once a client c_u requests unlearning, its historical updates are removed from the stored checkpoints. However, simply discarding these updates is insufficient because residual traces may remain in the aggregated model. To eliminate such influence, PEV employs a layer-wise sensitivity analysis that determines how strongly each model layer is affected by the removed client's data. Layers with higher gradient variance are considered more sensitive and thus receive proportionally larger amounts of Gaussian noise, while stable layers receive less. This layer-adaptive noise injection ensures that the unlearned model is statistically indistinguishable from a model retrained without the client's data, achieving formal privacy guarantees while minimizing unnecessary utility loss.

Finally, the third module, Fingerprint-Based Verification, ensures verifiability of unlearning. Existing verification mechanisms often rely on intrusive backdoors or heavy cryptographic proofs, both of which are impractical for federated systems. PEV introduces a lightweight alternative: each client embeds gradient fingerprints, subtle perturbations that are correlated with its updates but do not interfere with normal training. After unlearning, the requesting client or a subset of peer clients can query the updated model to check for the presence of its fingerprint effect. If the fingerprint signal is no longer detectable, the unlearning is verified. This decentralized and non-invasive mechanism provides clients with the RTV, fostering trust in the unlearning process without imposing significant computational or communication overhead.

C. Algorithm Design

The proposed PEV framework is realized through three complementary algorithms that collectively achieve efficiency, privacy, and verifiability in federated unlearning.

Algorithm 1 (Adaptive Checkpointing): During training, the server selectively stores compressed checkpoints based on gradient variance. This ensures that only the most informative updates are retained, reducing storage overhead while preserving the ability to reconstruct unlearned models efficiently.

Algorithm 1 Adaptive Checkpointing during FL Training

Require: Number of clients n, global model w_1, checkpoint interval τ, threshold θ
  for each round $t = 1, 2, \ldots, T$ **do**
    Server broadcasts w_t to selected clients
    Each client c_i computes local update Δw_i^t
    Server aggregates updates: $w_{t+1} = w_t + \frac{1}{n}\sum_{i=1}^{n} \Delta w_i^t$
    if t mod $\tau == 0$ **then**
      Compute variance $V_t = \mathrm{Var}(\{\Delta w_i^t\})$
      if $V_t > \theta$ **then**
        Store checkpoint $C_t = \{\Delta w_i^t, w_t\}$
      end if
    end if
  end for
Ensure: Final model w_T and set of checkpoints $\{C_t\}$

Algorithm 2 (Layer-Adaptive DP Calibration): When a client requests unlearning, its updates are removed from the stored checkpoints. The remaining updates are calibrated with Gaussian noise injected proportionally to each layer's sensitivity, ensuring statistical indistinguishability from retraining while minimizing accuracy loss.

Algorithm 3 (Fingerprint-Based Verification): To provide verifiability, clients embed gradient fingerprints during training. After unlearning, the requesting client (or peers) can check whether its fingerprint influence remains. If absent, the unlearning is verified, guaranteeing transparency and trust without heavy cryptographic overhead.

D. Complexity Analysis

The efficiency of the proposed PEV framework can be understood by analyzing its storage, time, and communication costs.

Algorithm 2 Layer-Adaptive DP Calibration for Unlearning

Require: Stored checkpoints $\{C_t\}$, target client c_u, noise scale σ, index of a client i, index of the target client u
  Initialize unlearned model $\hat{w} = w_1$
  for each checkpoint C_t in $\{C_t\}$ **do**
    for each client update Δw_i^t in C_t **do**
      if $i == u$ **then**
        Remove Δw_i^t {Discard target client's contribution}
      else
        Compute sensitivity S_ℓ for each layer ℓ
        Add calibrated noise $\eta_\ell \sim \mathcal{N}(0, \sigma^2 S_\ell^2)$
        $\Delta w_i^t = \Delta w_i^t + \eta_\ell$
      end if
    end for
    Aggregate calibrated updates: $\hat{w} = \hat{w} + \frac{1}{n-1} \sum \Delta w_i^t$
  end for
Ensure: Unlearned model $\hat{w}$

Algorithm 3 Fingerprint-Based Verification

Require: Unlearned model $\hat{w}$, target client c_u, fingerprint F_u, threshold δ
  Step 1: Fingerprint Generation
    c_u generates gradient fingerprint F_u during initial training
  Step 2: Model Querying
    After unlearning, c_u (or peers) query $\hat{w}$ with marked data
  Step 3: Influence Evaluation
    Evaluate influence metric $I_F(F_u, \hat{w})$
  if $I_F(F_u, \hat{w}) < \delta$ **then**
    Verification success: Contribution of c_u removed
  else
    Verification failed: Potential unlearning failure
  end if

From a storage perspective, adaptive checkpointing significantly reduces the overhead of retaining historical updates [16]. Instead of storing the updates from all T training rounds, which would require $O(T \cdot d)$ memory for a model of dimension d, PEV selectively stores only k checkpoints identified through gradient variance. This reduces the storage requirement to $O(k \cdot d)$, where $k \ll T$, thereby achieving a more scalable trade-off between reconstruction accuracy and resource usage.

In terms of time complexity, unlearning under PEV avoids the computational burden of full retraining. Rather than re-running all training rounds, the unlearning process only iterates through the stored checkpoints, recalibrates updates, and injects layer-adaptive differential privacy noise where necessary. This design yields a practical speed-up of up to two to three times compared to retraining-based approaches, enabling faster response to unlearning requests in real-world FL environments.

Finally, the communication overhead introduced by fingerprint-based verification is minimal. Verification only requires lightweight gradient fingerprint queries, which are performed by the requesting client. Unlike cryptographic or backdoor-based verification schemes that demand extensive additional communication or computational resources, PEV's verification process incurs negligible extra rounds beyond the standard FL workflow.

IV. RESULTS AND DISCUSSION

A. Dataset description

To evaluate the effectiveness of PEV, we conducted experiments on a diverse set of benchmark datasets spanning image, handwritten character, and medical domains. For image classification, we used CIFAR-10 [17], which consists of 60,000 labeled images across 10 categories, providing a widely adopted benchmark for FL. For handwritten character recognitions, we employed FEMNIST [18], a variant of the Extended MNIST dataset where samples are partitioned by writer identity, naturally simulating the non-IID (non-identically distributed) data characteristics of FL. Finally, to examine applicability in sensitive domains, we included a medical imaging dataset derived from chest X-ray scans [19], representing a real-world use case where privacy and unlearning are particularly critical. This combination of datasets allows us to assess the generalizability of PEV across different modalities and application contexts.

B. Evaluation Criteria

Our evaluation employed multiple metrics to capture the efficiency, privacy, and verifiability dimensions of unlearning. Model accuracy was measured to ensure that the unlearning process did not significantly degrade utility for non-deleted clients. Unlearning effectiveness was computed by accuracy drops at different DP noise levels. Verification of unlearning was done by fingerprint-based verification, which correctly confirmed the removal of a client's contribution and ensured privacy guarantees. Finally, we tracked runtime to quantify the efficiency gains of PEV compared to full retraining and existing baselines.

C. Experimental setup

All experiments were implemented in PyTorch and executed in a FL simulation with one central server and 100 participating clients. Clients were sampled randomly in each round, with 10 clients selected per round for local training. For CIFAR-10, we trained a Convolutional Neural Network (CNN) with three convolutional layers and two fully connected layers. FEMNIST experiments used a two-layer LSTM for handwritten character recognition, while the medical dataset was modeled using a ResNet-18 architecture pretrained on ImageNet. Training was conducted for 200 communication rounds with a learning rate of 0.01 and local batch sizes of 32.

We compared PEV against three key baselines: FedEraser [8] (efficient unlearning via calibrated historical updates), FedRecovery [9] (differentially private unlearning), and VeriFi

[10] (verifiable unlearning with markers). For fairness, all methods were tested under identical FL settings, with hyperparameters tuned according to the original implementations. PEV was configured with adaptive checkpointing intervals of 10 rounds, a gradient variance threshold $\theta = 0.01$, and layer-adaptive Gaussian noise with base scale $\sigma = 0.5$.

D. Result analysis

The experimental results demonstrate that PEV achieves superior performance across efficiency, privacy, and verifiability dimensions.

Fig. 1 compares the accuracy of different unlearning methods against the retraining baseline. As expected, retraining achieves the highest accuracy since it fully excludes the target client's data and retrains the model from scratch. Among the baselines, FedEraser and VeriFi maintain reasonable accuracy. However, FedRecovery suffers a significant utility drop, losing as much as 7.4% on average due to uniform noise injection across all layers. In contrast, PEV consistently outperforms baselines across CIFAR-10, FEMNIST, and the medical dataset. This demonstrates that PEV's layer-adaptive differential privacy calibration effectively localizes noise to sensitive layers, preserving utility in stable ones while still ensuring privacy guarantees.

Table I presents the average unlearning time for removing one client. Retraining is the slowest, requiring up to 2706 seconds on the medical dataset due to full model retraining. FedEraser reduces this cost significantly to 517–1434 seconds by leveraging stored historical updates. FedRecovery and VeriFi also shorten the process compared to retraining, but both remain slower than FedEraser due to additional computations for DP calibration and verification, respectively. PEV further reduces unlearning time to 475–1317 seconds, achieving up to a 2–3× speed-up compared to retraining and outperforming FedEraser by avoiding unnecessary update storage and reconstruction. This result highlights the effectiveness of adaptive checkpointing, which retains only the most informative updates, minimizing both storage and computation.

TABLE I
AVERAGE UNLEARNING TIME (SECONDS) FOR REMOVING ONE CLIENT.

Method	CIFAR-10	FEMNIST	Medical
Retrain	1807	1452	2706
FedEraser [8]	1034	517	1434
FedRecovery [9]	1259	748	1621
VeriFi [10]	1329	850	1754
PEV (Ours)	**993**	**475**	**1317**

Table II compares accuracy drop under different levels of DP noise on CIFAR-10 dataset. FedRecovery shows a steep decline, with accuracy losses rising to over 9% at $\sigma = 0.8$. This highlights the drawback of applying uniform noise globally across model layers. PEV, by contrast, injects noise selectively based on layer sensitivity analysis, resulting in much smaller losses: only 1.2% at $\sigma = 0.2$, 2.9% at $\sigma = 0.5$, and 5.1% at $\sigma = 0.8$. On average, PEV's accuracy drop is less than half of FedRecovery's. These findings confirm that layer-adaptive calibration strikes a better balance between privacy and utility, making PEV more practical for real-world applications where accuracy is critical.

TABLE II
ACCURACY DROP (%) AT DIFFERENT DP NOISE LEVELS (σ).

Method	$\sigma = 0.2$	$\sigma = 0.5$	$\sigma = 0.8$	Avg. Drop
FedRecovery [9]	2.8	6.5	9.3	6.2
PEV (Ours)	**1.2**	**2.9**	**5.1**	**3.1**

Fig. 2 confirmed that PEV consistently achieved successful verification outcomes. For each of the datasets with ten clients, the fingerprint of the target client was embedded during training, erased during unlearning, and subsequently verified as absent. The verification function reported success with an influence score below the threshold, indicating that the target client's contribution had been effectively removed from the global model.

Across all evaluation dimensions, PEV consistently outperforms the baselines by integrating efficiency, privacy, and verifiability into a single framework. Compared to FedEraser and VeriFi, PEV reduces unlearning time and increases accuracy, while adding privacy guarantees. Relative to FedRecovery, PEV significantly improves utility preservation by adopting a layer-adaptive DP mechanism. This balanced performance demonstrates that PEV is not only theoretically sound but also practically deployable in real-world FL environments.

V. CONCLUSION

The increasing adoption of FL in privacy-sensitive domains such as healthcare, finance, and mobile applications has amplified the importance of mechanisms that respect user data rights. Among these, the RTBF has emerged as a critical requirement, demanding that clients be able to withdraw their contributions from trained models. While prior research has introduced mechanisms for federated unlearning, existing approaches have primarily focused on one dimension at a time-efficiency, privacy, or verifiability-leaving important gaps in their practical applicability.

In this work, we proposed PEV, a unified federated unlearning framework that holistically integrates these three dimensions. PEV employs adaptive checkpointing to achieve efficient unlearning with reduced storage overhead, layer-adaptive differential privacy calibration to ensure statistical indistinguishability from retraining, and fingerprint-based verification to empower participants with the RTV in a lightweight and decentralized manner. Through comprehensive experiments on benchmark image, handwritten character, and medical datasets, we demonstrated that PEV achieves up to 2–3× faster unlearning compared to retraining, provides formal privacy guarantees with improved utility preservation compared to DP-only baselines.

The results confirm that PEV represents a significant step forward in making federated unlearning practical, regulation-compliant, and trustworthy. By addressing efficiency, privacy, and verifiability simultaneously, PEV bridges the limitations of

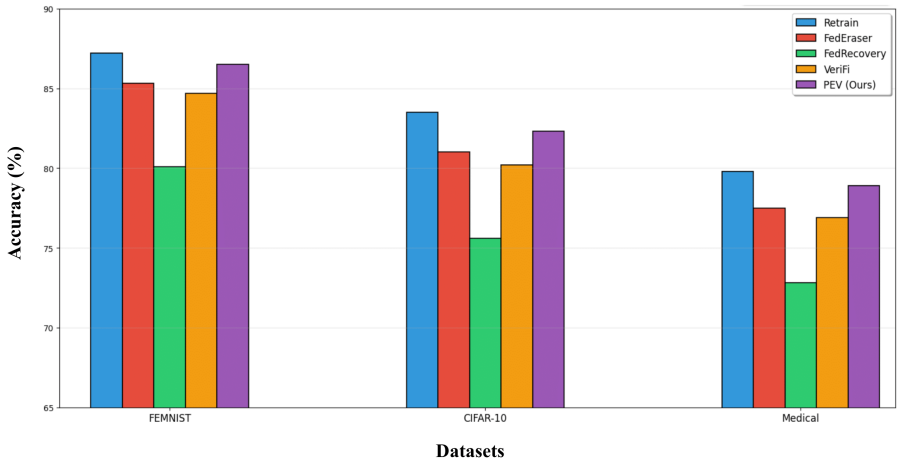

Fig. 1. Model accuracy (%) after unlearning compared to retraining baseline.

```
Verification: no fingerprint
```

Fig. 2. Fingerprint verification result of unlearned client.

prior approaches such as FedEraser, FedRecovery, and VeriFi, while offering a deployable solution for real-world federated systems.

Looking ahead, several directions remain for future exploration. First, extending PEV to handle multi-client unlearning requests simultaneously could further broaden its applicability in dynamic FL environments. Second, investigating cross-silo and heterogeneous FL settings, where clients differ substantially in resources and data distributions, would enhance robustness. Third, integrating cryptographic assurances with lightweight verification could strengthen trust in adversarial settings. Finally, establishing standardized benchmarks and protocols for federated unlearning will be essential for fair comparison and progress.

REFERENCES

[1] P. Goswami and A. A. Hossain, "Street object detection from synthesized and processed semantic image: A deep learning based study," *Human-Centric Intelligent Systems*, vol. 3, no. 4, pp. 487–507, 2023.

[2] P. Goswami, A. A. Hossain, and A. N. M. Sakib, "An end-to-end web-based system for rice leaf disease classification using deep learning," in *International Joint Conference on Advances in Computational Intelligence*. Singapore: Springer Nature Singapore, 2022, pp. 517–531.

[3] P. Goswami, A. A. Safi, A. N. M. Sakib, and T. Datta, "Corn leaf disease identification via transfer learning: A comprehensive web-based solution," in *International Conference on Sustainable and Innovative Solutions for Current Challenges in Engineering & Technology*. Singapore: Springer Nature Singapore, 2023, pp. 429–441.

[4] L. Bourtoule, V. Chandrasekaran, C. A. Choquette-Choo, H. Jia, A. Travers, B. Zhang, D. Lie, and N. Papernot, "Machine unlearning," in *2021 IEEE symposium on security and privacy (SP)*. IEEE, 2021, pp. 141–159.

[5] A. Ginart, M. Guan, G. Valiant, and J. Y. Zou, "Making ai forget you: Data deletion in machine learning," in *Advances in neural information processing systems 32*, 2019.

[6] A. Golatkar, A. Achille, and S. Soatto, "Eternal sunshine of the spotless net: Selective forgetting in deep networks," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 9304–9312.

[7] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y. Arcas, "Communication-efficient learning of deep networks from decentralized data," *Artificial intelligence and statistics*, pp. 1273–1282, 2017.

[8] G. Liu, X. Ma, Y. Yang, C. Wang, and J. Liu, "Federaser: Enabling efficient client-level data removal from federated learning models," in *2021 IEEE/ACM 29th International Symposium on Quality of Service (IWQOS)*. IEEE, 2021, pp. 1–10.

[9] L. Zhang, T. Zhu, H. Zhang, P. Xiong, and W. Zhou, "Fedrecovery: Differentially private machine unlearning for federated learning frameworks," in *IEEE Transactions on Information Forensics and Security*, vol. 18. IEEE, 2023, pp. 4732–4746.

[10] X. Gao, X. Ma, J. Wang, Y. Sun, B. Li, S. Ji, P. Cheng, and J. Chen, "Verifi: Towards verifiable federated unlearning," in *IEEE Transactions on Dependable and Secure Computing*, vol. 21. IEEE, 2024, pp. 5720–5736.

[11] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," in *Foundations and trends® in theoretical computer science*, vol. 9, 2014, pp. 211–407.

[12] A. Sekhari, J. Acharya, G. Kamath, and A. T. Suresh, "Remember what you want to forget: Algorithms for machine unlearning," in *Advances in Neural Information Processing Systems*, vol. 34, 2021, pp. 18 075–18 086.

[13] C. Guo, T. Goldstein, A. Hannun, and L. V. D. Maaten, "Certified data removal from machine learning models," in *arXiv preprint arXiv:1911.03030*, 2019.

[14] J. Zhang, D. Chen, J. Liao, W. Zhang, H. Feng, G. Hua, and N. Yu, "Deep model intellectual property protection via deep watermarking," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44. IEEE, 2021, pp. 4005–4020.

[15] S. Mohammadi, V. Tsouvalas, I. Symeonidis, A. Balador, T. Ozcelebi, F. Flammini, and N. Meratnia, "Efu: Enforcing federated unlearning via functional encryption," in *arXiv preprint arXiv:2508.07873*, 2025.

[16] M. K. Islam, A. Biswas, and H. Hu, "Adaptive real-time gap detection system: A multi-algorithm approach integrating ultrasonic sensing and machine learning for robust structural analysis," in *2025 IEEE 4th International Conference on Computing and Machine Intelligence (ICMI)*. IEEE, 2025.

[17] K. Alex and G. Hinton, "Learning multiple layers of features from tiny images," *University of Toronto*, vol. 7, 2009.

[18] S. Caldas, S. M. K. Duddu, P. Wu, T. Li, J. Konečný, H. B. McMahan, V. Smith, and A. Talwalkar, "Leaf: A benchmark for federated settings," in *arXiv preprint arXiv:1812.01097*, 2018.

[19] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, and G. M. Summers, "Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2097–2106.