



Article

H-VSAF: a verifiable and privacy-preserving federated learning framework for secure remote healthcare monitoring

Sureshkumar S^{1*}, Santhoshkumar S.P², Joseph James S³, Priya R⁴

¹Department of Computer Science and Engineering, P. A. College of Engineering and Technology, Tamil Nadu, India

²Department of Computer Science and Engineering, Vel Tech Rangarajan Dr. Sagunthala R&D Institute of Science and Technology, Avadi, Chennai, Tamil Nadu, India

³Department of Computational Intelligence, Engineering and Technology, SRM Institute of Science and Technology, Kattankulathur, Tamil Nadu

⁴Department of Artificial Intelligence and Data Science, Pollachi Institute of Engineering and Technology, Poosaripatti, Pollachi, India

ARTICLE INFO

Article history:

Received 10 May 2026

Received in revised form

22 September 2026

Accepted 09 October 2026

Keywords:

Federated learning, Healthcare monitoring,

Secure aggregation,

Verification and privacy-preserving

*Corresponding author

Email address:

sureshkumar.pacet@gmail.com

DOI: 10.55670/fpml.futech.5.4.27

ABSTRACT

The sensitive patient data handled by remote healthcare monitoring systems requires reliable, secure learning methods. Federated learning enables collaborative healthcare analytics without storing raw medical data centrally. However, existing techniques still suffer from issues including privacy leakage, untrusted aggregation, and frequent device outages. In this work, we propose H-VSAF, a privacy-preserving, verifiable, and lightweight federated learning system for healthcare environments. H-VSAF uses linear homomorphic hashing, Bloom filter-based online authentication, Shamir secret sharing, Diffie-Hellman key generation, and single-mask Secure Aggregation to ensure secrecy, precision, and resilience. The security analysis provides informal guarantees against gradient inference, aggregation tampering and deletion, and forged dropout claims under the threat model and cryptographic assumptions above. We experimentally evaluate H-VSAF on healthcare datasets and demonstrate that it achieves model accuracy within 0.2-0.8% of standard federated learning while decreasing client compute cost by up to 32% and communication overhead by 28% compared to traditional secure aggregation methods. Moreover, the framework is resilient to 40% client dropouts and introduces verification latency < 35 ms each round. The results clearly show that H-VSAF provides a credible and practical platform for safe joint health care analytics.

1. Introduction

Remote healthcare monitoring systems are now necessary for ongoing clinical observation, chronic disease evaluation, post-treatment follow-up, and senior care management. Wearable sensors, IoT medical devices, and smart health apps continuously collect physiological data, including ECG signals, heart rate variability, blood pressure indices, oxygen saturation levels, and patient activity patterns [1,2]. These data streams are highly private and protected by strong healthcare privacy laws [3,4]. Standard centralized learning methods require transferring medical records to a central server for model training, raising significant privacy, trust, regulatory, and ethical concerns [5]. To address these issues, federated learning (FL) offers an effective approach for collaborative model training without transferring raw data from patient devices [6]. In healthcare settings, adopting

FL remains challenging due to privacy concerns, untrusted aggregation behavior, inconsistent device connectivity, and the limited capabilities of medical IoT devices [7]. Given these persistent issues, developers are focusing on lightweight, safe, and verifiable healthcare-focused FL frameworks. To the best of our knowledge, no current approach simultaneously provides privacy protection, verified aggregate integrity, dropout robustness, and lightweight execution for remote healthcare monitoring settings [8,9]. This study aims to fill this essential gap.

1.1 Problem statement

Federated learning prevents the direct sharing of patient data, but it doesn't automatically guarantee privacy, trust, or dependability. First, most existing secure aggregation approaches protect gradient privacy but assume an honest

cloud server [10]. In real-world healthcare settings, the cloud can be harmful by intentionally modifying model updates, excluding some client contributions, or displaying incorrect aggregate results. This can have a significant impact on clinical diagnostic results [11].

Second, healthcare gadgets generally have limited battery power, disconnect when individuals move around, have inconsistent wireless links, and have unpredictable patient behavior. Many of the existing FL solutions cannot tolerate such client dropouts. When devices fail mid-training, they can be dangerous. Third, some advanced cryptographic techniques, e.g., homomorphic encryption or complex multi-mask schemes, are computationally and communication-wise intensive and thus infeasible for resource-constrained healthcare equipment [12]. Therefore, this research addresses the following question: How can we build a federated learning system for remote healthcare monitoring that preserves privacy during training, guarantees verifiable aggregation integrity, is robust to client dropouts, and functions effectively in resource-constrained healthcare settings?

1.2 Motivation

Healthcare analytics directly impacts patient safety and treatment accuracy. Any incorrect or manipulated learning outcome can lead to misdiagnosis, delayed action, or clinically harmful advice. So, ensuring trust, accuracy, and privacy is not only a technology necessity but also a clinical requirement. Healthcare equipment also operates in unpredictable, variable, and power-constrained environments. Thus, a healthcare FL system must still operate even if many individuals drop out, while maintaining sufficient security without overloading computing devices. Together, these challenges necessitate a lightweight FL architecture that is trust-assured, privacy-preserving, and verifiable, specifically for remote healthcare monitoring resources.

1.3 Research Objectives

The primary objectives of this research are:

- To develop a federated learning system for remote healthcare monitoring that maintains patient privacy by keeping patient gradients confidential.
- To develop a verifiable secure aggregation mechanism that enables participating clients to verify the correctness and completeness of the aggregated model update generated by the cloud server.
- To make it robust enough to support repeated dropouts of healthcare devices without affecting privacy and accuracy.
- To reduce computational and communication overhead so that the proposed framework can operate effectively on resource-constrained healthcare devices and networks.
- To evaluate the effectiveness of the proposed system, the security analysis and experimental assessment are performed utilizing healthcare datasets.

1.4 Major contributions

The main contributions of this investigation are summarised as follows:

- A unified H-VSAF framework is proposed for secure and verifiable federated learning in remote healthcare monitoring, integrating privacy preservation, aggregation integrity, and client-dropout resilience.
- A lightweight single-mask secure aggregation mechanism is developed to protect client model updates while reducing computational and communication overhead.
- A verification process using linear homomorphic hashing and Bloom filter authentication is designed to empower

correctness validation and prevent malicious or unreliable server behavior.

- A dropout-aware federated protocol based on Shamir secret sharing is developed to maintain secure aggregation when participating healthcare devices disconnect during training.
- Extensive security analysis and experimental evaluation show that the proposed framework keeps learning accuracy while decreasing the amount of computing and communication overhead, making it realistic for use in a real-world healthcare context.

Individual use of homomorphic hashing, masking, Shamir secret sharing, and dropout-aware aggregation is not novel in H-VSAF, since they have been studied in the past. Rather, H-VSAF combines online participation proofs using a Bloom filter with a simple single-mask cancellation mechanism, Shamir-based dropout recovery, and homomorphic hashing for aggregate verification. The integration focuses on remote healthcare monitoring scenarios with resource-constrained devices and sporadic involvement.

2. Literature review

Sensitive medical data in RHM supports healthcare analytics and doesn't have to be stored in one place. Federated Learning (FL) is an important method that supports a decentralized approach. FL lets hospitals, medical institutions, wearable devices, and personal health assistants collect patient records and use them to train a global model while leaving the data on their own servers. Using FL in real-world healthcare greatly reduces privacy risks. But it still faces challenges in security, reliability, and trust. This section outlines, in terms of research areas, what motivated the proposed work. Recent studies have explored secure federated learning in healthcare and general distributed environments.

Although homomorphic encryption-based frameworks provide strong privacy protection, they incur very high computation and communication costs, making them impractical for wearable and mobile healthcare equipment. Several verifiable aggregation algorithms have been developed. However, most assume reasonably consistent participation and do not explicitly account for high-frequency dropout, which is typical in real healthcare deployments. Current healthcare-oriented FL solutions focus mostly on privacy preservation but seldom provide explicit verification capabilities to detect harmful aggregation behavior. Also, many safe FL protocols rely on complex multi-masking or multi-round cryptographic execution, which heavily burdens IoT medical nodes with limited computational capacity.

As summarized in Table 1, the reviewed works show impressive advances in privacy-preserving Federated Learning in healthcare and related domains. However, based on the features and drawbacks extracted from the reviewed studies, the overall problem of privacy-preserving secure aggregation, participant authentication, verifiable aggregation, and resilience to client drop-off is still not sufficiently covered in a unified framework. This observation motivates H-VSAF, a combination of these mechanisms that offers a more comprehensive security and reliability solution for remote healthcare monitoring.

Table 1. Federated learning research in healthcare and secure aggregation, showing their key ideas, security features, and limitations compared with the proposed H-VSAF framework

Year	Paper / Approach	Domain / Application	Key Idea / Techniques	Security / Trust Features	Limitations vs Proposed H-VSAF
2024	Federated learning in medicine [13]	General healthcare FL (multi-institution)	Survey + guideline paper focusing on privacy risks and mitigation techniques for FL in healthcare (DP, secure aggregation, HE, etc.).	Identifies privacy leakage channels and summarizes state-of-the-art protection mechanisms for healthcare FL.	High-level review, not a concrete protocol; no unified verifiable aggregation + dropout-resilient framework; no lightweight IoT-side design. Good to cite for “problem importance” and “open challenges”.
2024	Federated Learning in Smart Healthcare [14]	Smart healthcare / IoT, remote monitoring	Deep dive on FL integration with IoT devices, wearables and remote monitoring; covers architectures and use cases.	Discusses privacy, interoperability, and deployment issues qualitatively.	No concrete secure aggregation or verifiable scheme; security is discussed conceptually; no protocol-level solution for malicious or untrusted aggregator and dropout – a gap your work fills.
2024	Secure and Efficient Federated Learning Schemes for Smart Healthcare Systems [15]	Smart healthcare	Uses homomorphic encryption to secure gradients in FL for healthcare. Adds identity authentication module to reduce malicious client participation.	Strong privacy via HE; some efficiency improvements claimed for smart healthcare.	HE-based approach still heavy for constrained wearables; no explicit verifiable aggregation (integrity proof); dropout handling is not the main focus; your H-VSAF is lighter (single mask) and adds verifiability + dropout robustness.
2024	Explainable federated learning scheme for secure healthcare [16]	Secure healthcare FL	Proposes explainable FL with a verifiable secure aggregation protocol using Shamir's additive secret sharing, targeting privacy and verifiability together.	Provides verifiable secure aggregation and client privacy; uses threshold secret sharing for t-resilient privacy.	Focus is more on explainability + verifiable aggregation; not specifically optimized for resource-constrained IoT or high dropout rates; no Bloom-filter-style online participation proof; your work refines lightweightness and dropout handling in remote monitoring context.
2024	ESVFL: Efficient and Secure Verifiable Federated Learning Scheme with Privacy-Preserving (FGCS) [17]	Generic FL (not healthcare-specific)	Proposes ESVFL, an efficient verifiable FL scheme that encrypts gradients with low user-side overhead and uses cross-validation-based verification.)	Achieves both privacy and verifiability with reduced computation on users; addresses malicious server falsification and user dropout.	Generic setting, not tailored for medical IoT / remote healthcare constraints; does not explicitly address clinical reliability requirements; your H-VSAF adapts these ideas to healthcare + lightweight masking + Bloom filters.
2025	Federated learning against poisoning attacks with security and verifiability [18]	Generic FL security	Uses masking + homomorphic hash to protect gradient privacy and verify aggregation while defending against poisoning attacks.	Explicitly combines privacy, verifiable aggregation, and poisoning resistance via homomorphic hashing.	Focus is more on poisoning; domain not constrained to healthcare; no explicit dropout management; may have heavier crypto cost than your single-mask, healthcare-oriented design. Still very good to cite for homomorphic-hash based verifiable FL.
2025	Verifiable secure aggregation scheme for privacy protection in federated learning [19]	Generic FL	Proposes a secure aggregation technique verified assuming potentially malicious server; uses cryptographic verification to ensure integrity of aggregated results.	Addresses malicious server (tampering/forging) and provides verifiable aggregation plus privacy.	Not healthcare-specific; design/overheads may not be optimized for medical IoT; dropout and resource constraints not the main focus; you can position H-VSAF as healthcare-optimized with additional dropout handling and Bloom-based participation proof.
2025	Privacy-preserving federated learning for collaborative medical image classification [20]	Medical imaging (TB X-ray, brain tumor MRI, DR)	Blends transfer learning and FL with adaptive aggregation for privacy-preserving multi-center medical image classification.	Demonstrates that FL can maintain high accuracy on multiple medical imaging datasets while preserving privacy.	Focus on accuracy and privacy, not on secure aggregation, verifiable correctness, or dropout resilience; no trust-assurance layer for malicious server; your work adds security + verifiability + protocol-level protection for remote monitoring data.

2.1 Research gap

- Many secure FL techniques ensure privacy, but they don't give you proof that aggregation is carried out appropriately.
- Some Secure Aggregation methods can accommodate dropouts; however, most of them either do not respect privacy commitments or need sophisticated execution changes, which are not acceptable for real-world healthcare settings.
- Current strong cryptographic methods frequently require a lot of processing power and bandwidth, which makes them unfeasible for wearable and mobile health resources.
- Only a few studies look at privacy, integrity verification, dropout robustness, and lightweight operation all at once.
- Most healthcare-specific trust criteria and clinical reliability expectations are not met, creating a gap between theoretically secure FL models and their practical application in healthcare.

These gaps indicate the need for a new federated learning architecture for healthcare that is safe and verifiable. Thus, many existing works study privacy, verifiability, efficiency, and dropout robustness in isolation. There is still a lot of effort to be done on a unified lightweight framework that solves all of these demands at once in a real-world healthcare environment.

3. Overview of healthcare federated learning scenario

Wearable sensors, Internet of things (IoT) medical devices, cellphones, home-care gateways, and hospital information systems are all used to continuously collect electronic cardiac graphs (ECGs), heart rate variability (HRV), blood oxygen saturation (SaO₂), blood pressure (BP), glucose readings (GS), and other clinical indicators in remote healthcare monitoring environments. Many governments do not want or allow raw healthcare records to be sent to centralized cloud servers since these datasets are particularly sensitive and subject to strict privacy restrictions. Federated learning (FL) is a promising solution to collaborative healthcare analytics. It keeps patient data local while training a shared machine learning model. Each healthcare institution trains its own model locally on the patient's medical records. It sends only the updated or learned gradients to the cloud.

The cloud merges these inputs into a global model without collecting raw data. However, FL in healthcare is challenging due to limited processing capacity, unreliable connections, frequent dropouts, privacy concerns, and the possibility of malicious aggregation by medical IoT devices. To address these challenges more conveniently, securely, and reliably for federated learning in remote healthcare, we propose the Healthcare-Verifiable Secure Aggregation Framework (H-VSAF).

3.1 Assumptions

The boundaries of system practicality and security are established on the basis of the following assumptions:

- Communication channels are authenticated, preventing identity spoofing.
- The TA is reliable, does not collude with adversaries, and follows protocol rules.
- Healthcare clients follow the protocol but may disconnect unpredictably.
- The cloud server may be curious or malicious, attempting to analyze or manipulate aggregation.
- No raw medical data is transmitted outside client devices at any point.
- Standard cryptographic primitives such as Diffie-Hellman, pseudo-random generators, and hashing functions are secure.

3.2 Threat model

The suggested framework is designed for a realistic adversarial scenario. Figure 1 shows the threat model. The system assumes a semi-honest but potentially malicious cloud server which may attempt to manipulate aggregation outcomes, omit user contributions, or falsify global model results. Client devices are honest-but-curious and may attempt to infer other participants' gradients if given the opportunity. External attackers may attempt replay, tampering, or collusion attacks. Devices often drop out because they move, run out of energy, or lose their connection; therefore, the framework must keep running securely under those scenarios. A Trusted Authority (TA) assists in verification and authentication but does not access raw gradients.




Figure 1. Threat model of the proposed federated healthcare system

In the proposed H-VSAF framework, the cloud server is assumed to be potentially curious or malicious and may attempt to compromise data confidentiality and model integrity. Threats include gradient inference attacks that reconstruct sensitive patient information from local model updates and tampering attempts that alter gradients before or after aggregation to skew the global model. The server may also perform deletion or partial aggregation attacks by selectively excluding client updates, forge false aggregation results while presenting them as valid, or deceive participants by falsely claiming client dropouts. In addition, a small number of clients may collude with the cloud server to infer the protected updates, further compromising privacy. Because healthcare environments are dynamic, frequent client dropouts due to movement, network instability, battery depletion, or patient idleness might let an adversarial server recover masked gradients or disrupt training. It also considers replay and session manipulation attacks, where an attacker reuses previously sent gradients or tokens in other training cycles. This work excludes side-channel attacks, physical device compromise, and model poisoning, leaving them for further exploration.

3.3 Design goals and security requirements

To address these risks, H-VSAF aims to provide strong gradient privacy, verifiable aggregation integrity, dropout robustness, and practical efficiency. The architecture ensures that no unauthorized party can infer patient-specific information from the gradients sent, even with minimal client-server collusion. The system lets clients verify that the aggregation is correct and complete, and that their contributions are included and consistent. The system can handle clients joining or leaving at any time and keeps training running smoothly even when many people drop out, all without breaking any privacy rules. H-VSAF is also effective because it requires little power to process and communicate, and it integrates well with wearable sensors, mobile medical devices, and home healthcare systems. The H-VSAF platform can also handle a growing number of users. It is designed to work in real time, remain reliable for clinical use, and comply with all rules for keeping healthcare data private.

3.4 Problem formulation

Let the set of healthcare devices participating in federated learning be represented as:

$$U = \{u_1, u_2, \dots, u_N\} \quad (1)$$

Each device u_i holds a local healthcare dataset D_i and computes a local model update x_i . The federated learning objective is to calculate the global aggregation:

$$Z = \sum_{i=1}^N x_i \quad (2)$$

Subject to the following constraints: x_i Remains confidential at all times. The computed Z must be verifiably correct. Aggregation must remain valid even under client dropout. Computation and communication must remain lightweight.

4. Methods

The H-VSAF framework enables safe, verified, and effective federated learning for remote healthcare monitoring systems. The four mechanisms are complementary: single-mask secure aggregation for gradient confidentiality; a Diffie-Hellman key exchange and Shamir secret sharing scheme for secure dropout management; Bloom filtering with authentication; and aggregation integrity verification using

linear homomorphic hashing. Together, these components support secure model aggregation without overwhelming the computational and communication capacity of resource-constrained healthcare devices. The Healthcare-Verifiable Secure Aggregation Framework has a design that consists of four main parts:

- Secure aggregation to hide distinct slopes with a single mask.
- Diffie-Hellman and Shamir secret sharing for dropout management with privacy preservation.
- Bloom filter authentication to verify genuine participant availability.
- Verifying the correctness of aggregation based on linear homomorphic hash.

The next sections cover the operational flow, primary components, mathematical underpinnings, and protocol execution stages.

4.1 Security analysis

Inference resistance: Each client sends only a masked model update, with pairwise masks generated from Diffie-Hellman shared secrets. Thus, the cloud server is unable to recover an individual client's update from the transmitted masked value. Shamir secret sharing also protects the information needed to recover the mask if a client drops out. Individual client gradients remain confidential, assuming the cryptographic primitives and authenticated communication channels are secure.

Tampering resistance: This is where the client computes the linear homomorphic hash of its local update. During aggregation, the verification authority aggregates all received hash information and checks the hash of the aggregated result. If any attribute of a masked update or the aggregation changes, the verification value changes, and the aggregation is rejected.

Deletion/partial-aggregation resistance: If the cloud fails to include an authenticated participating client without creating an inconsistency between the participating-client information and the aggregation verification information, it is deletion/partial-aggregation resistant. The Bloom-filter participation mechanism helps identify real participants, and the verification process verifies the completeness and correctness of the aggregation results.

Active client detection: The Bloom-filter mechanism identifies active clients, and the cloud cannot arbitrarily mark them as dropouts. When the protocol detects an incorrect dropout, it does not accept an incorrectly reconstructed aggregate; instead, it aborts. If an honest dropout occurs, the required masking information can be reconstructed via the Shamir secret-sharing mechanism, allowing the secure aggregation to continue.

Security scope: These guarantees hold under the following assumptions: the Trusted Authority adheres to the protocol; communication channels are authenticated; standard cryptographic primitives remain secure; and the number of colluding participants is below the assumed limit. The work does not consider physical compromise, side-channel attacks, or model poisoning.

4.2 Architectural overview of H-VSAF

The framework architecture, as shown in Figure 2 comprises:

- **Healthcare Clients / IoT Devices**
Train local models and secure their upgrades.
- **Cloud Server (CS)**
Orchestrates centralized training and reliable aggregation.

• Trusted Authority (TA)

Creates cryptographic materials, authenticates clients, and checks correctness.

The framework works in several learning cycles. In each round, clients compute local gradients, mask them using cryptography, and send them to the cloud. The TA receives only the verification proofs. The cloud never gets to see the raw gradients or verification tokens. This division of duties increases secrecy and prevents a single entity from accumulating power or colluding.

4.3 Cryptographic foundations

H-VSAF employs widely documented cryptographic primitives for healthcare-grade trust:

A. Diffie-Hellman key exchange

Each pair of clients generates a shared key:

$$S_{i,j} = g^{a_i b_j} \bmod p \quad (3)$$

where a_i, b_j are private keys, g is generator and p is a large prime. This creates pairwise shared randomness without explicitly sharing secrets. The corresponding mask is generated as:

$$m_{i,j} = \text{PRG}(S_{i,j}) \quad (4)$$

B. Shamir secret sharing (t, n): To accept dropouts. Each client's secret is split using threshold sharing in such a way that at least t users are required to rebuild missing keys.

C. Bloom filter authentication: Used to verify which clients are genuinely online, preventing false dropout claims by the cloud.

D. Linear homomorphic hash: Ensures verifiable aggregation, since:

$$HH(x_1 + x_2 + \dots + x_N) = HH(x_1) \cdot HH(x_2) \dots HH(x_N) \quad (5)$$

Any alteration leads to a failure in verification.

Cryptographic parameter configuration: H-VSAF runs a 20-client experimental configuration, requiring 15 client secrets for a Shamir secret sharing (t,n) parameter setting to reconstruct a missing client secret. Here, the Bloom filter is set to 256 bits and 4 separate hash functions. With this setup, the rate of false positives is about 0.52% for 20 clients participating. The Diffie-Hellman shared secret serves as the seed material for the cryptographically secure HMAC-DRBG with the hash function set to SHA-256. This cryptographic setup is kept the same in each round of the experiment.

Participation authentication using a Bloom filter: H-VSAF uses a Bloom filter to track the healthcare client's participation status. For every federated learning round ran authenticated client will create a participation token with the client identifier, round identifier, and freshness value. The same token generated during one round of training is not allowed to be used in another round of training. The Bloom filter is designed for k independent hash functions and m bits, and it is used for an expected maximum of n participating clients. The false positive probability is expected to be based on

$$P_{FP} \approx \left(1 - e^{-\frac{kn}{m}}\right)^k \quad (6)$$

The Bloom filter is used for efficient membership checking, but it is not a cryptographic proof of membership on its own. A positive Bloom filter result means the message may be in the collection, and the Bloom authenticated participation token and round data finalize the result. A false positive message can trigger a false membership candidate to be further evaluated but cannot trigger dropout recovery by itself. On the other hand, a false negative makes the reported participant status inconsistent, and the round is not completed until the status is determined. The conservative decision rule helps keep Bloom-filter uncertainty from directly translating into an accepted aggregation result.




Figure 2. Proposed H-VSAF architecture for secure and verifiable federated healthcare learning

4.4 Workflow of H-VSAF protocol

The system works in four phases. It ensures security is enforced smoothly in each phase. The system focuses on enforcing security in each phase.

Scope of Integrity Guarantee: H-VSAF only verifies the integrity of the announced federated aggregate and not the accuracy of the learned model. By successfully masking and recovering from drops, the aggregate will be verified against the submitted client updates using a linear homomorphic hash function to detect unauthorized changes, deletions, or forged aggregation information. Due to the possibility of a corrupted or malicious client, however, it is possible that a corrupted or "poisoned" (malicious) client could send a cryptographically valid update and be indistinguishable from the honest client by aggregate verification alone. Therefore, the current state of affairs of poisoning attacks, device compromise, and side-channel attacks is not part of the threat model. Future work will explore lightweight anomaly detection mechanisms or robust aggregation mechanisms that can be applied to protected updates without sacrificing H-VSAF's privacy and communication goals.

Phase 1 – Initialization and key setup

The TA:

- Generates global parameters (p,g)
- Issues private/public keys for each client,
- Provides Bloom filter configuration,
- Distributes verification hash parameters.

Clients store assigned keys, and no healthcare data is exposed in this stage.

Phase 2 – Secret sharing and backup preparation

Each client:

- Generates Shamir shares of its private key.
- Distributes shares securely to peers.
- Stores received shares for recovery.

This ensures that dropout does not disable secure aggregation.

Phase 3 – Secure Masking and Proof Generation

Participation-token freshness: A participation token corresponds to the current federated learning round and is authenticated by client information and session information. If the round identifier and/or freshness value of the rejected token do not match the current aggregation session, it is rejected. Therefore, the tokens previously generated can not be played again to make it appear that they are currently playing.

This is the core privacy protection step.

Mask Generation: For each pair of clients u_i, u_j :

$$\mathcal{S}_{i,j} = DH(sk_i, pk_j) \quad (6)$$

Mask:

$$m_{i,j} = PRG(\mathcal{S}_{i,j}) \quad (7)$$

Single-mask gradient protection

Each client encrypts the gradient:

$$\tilde{x}_i = x_i + \sum_{j>i} m_{i,j} - \sum_{j<i} m_{i,j} \quad (8)$$

This ensures Strong privacy, balanced mask cancellation, and Lower overhead compared to double masking.

Proof generation: Clients compute:

- Homomorphic hash proof
- Bloom filter participation token

Transmission roles:

- Masked gradient → Cloud
- Verification proofs → TA

Thus, the cloud cannot verify or cheat secretly, and TA cannot infer gradients.

Phase 4 – Secure aggregation and verification

The cloud identifies missing users; TA confirms their authenticity.

Cloud-TA disagreement handling: The cloud reports the clients that it has received masked updates from and indicates which are missing. The TA verifies the corresponding authentication participation information and the status of the Bloom filter on its own. If there is client evidence of current round participation (even if it is not the same as what the cloud sees) and the cloud reports the client as dropout, the dropout is considered inconsistent, and the protocol is not reconstructed; it is simply aborted. If both parties agree that the client has indeed dropped out, Shamir-based recovery begins. If the TA cannot confirm the reported participation state, the aggregation is treated as a refusal or postponed until the discrepancy is cleared. Hence, a false positive in the Bloom filter or outdated token cannot directly lead to a forged-dropout decision. If false dropout is claimed → protocol aborts. If honest dropout occurs → masks reconstructed via secret shares. Cloud aggregates masked gradients; masks cancel out naturally, yielding:

$$Z = \sum_{i=1}^N \tilde{x}_i = \sum_{i=1}^N x_i \quad (9)$$

TA validates:

$$HH(Z) \stackrel{?}{=} HH_proof \quad (10)$$

If they match, the result is accepted; if not, the result is an attack detection. Thus, correctness, privacy, and trust are guaranteed.

4.5 Algorithmic description of H-VSAF

The proposed design uses four main algorithms. This section explains the framework workflow we propose. We divide the protocol into four steps that align with the main parts of the process. The framework workflow is what we are focusing on here. The protocol consists of these four steps, which follow the execution phases of the framework workflow. The H-VSAF protocol is implemented through 4 successive algorithms, each corresponding to a principal step in the framework. System Initialization and Key Distribution ([Algorithm 1](#)): Initialize the system and distribute the keys used to configure the cryptographic parameters and each client's public key. [Algorithm 2](#) implements Shamir secret sharing and sets up the secret information for drop-out recovery. In [Algorithm 3](#), local gradient masking is done, and the verification information, such as the homomorphic hash and Bloom-filter participation token, is generated. Finally, [Algorithm 4](#) securely aggregates and verifies the updates by accounting for client drop-off, reconstructing lost masks as needed, aggregating the masked updates, and verifying the resulting global update. These algorithms combine to complete the system initialization and verified federated aggregation execution flow. The initial H-VSAF environment is set up by [Algorithm 1](#). The Trusted Authority computes the system parameters, generates key material for the clients, sends the necessary public parameters, and initializes the homomorphic-hash and Bloom filter configurations. This algorithm sets the cryptographic and authentication parameters for the later stages of the protocol.

Algorithm 1: Initialization and Key DistributionInput: Set of users $U = \{u_1, u_2, \dots, u_N\}$

Output: Public parameters, DH keys, hash settings

- 1: TA generates system parameters (p, g)
- 2: For each user u_i , generate private key a_i
- 3: Compute public key $pk_i = g^{a_i} \bmod p$
- 4: Distribute (p, g, pk_i) to participants
- 5: TA initializes homomorphic hash parameters
- 6: TA initializes a 256-bit Bloom filter with 4 hash functions
- 7: return System initialized

The dropout-recovery mechanism is given by Algorithm 2. The private key of each client participating in the protocol is split into Shamir secret shares and shared securely amongst the participating peers. When a legitimate client is unavailable for aggregation, this information is then used to recreate the needed secret information. The idea of aggregation in H-VSAF is based on a simple concept. It works by creating pairs of masks between users. Each mask has an effect on one user and a negative effect on the other user. When the cloud adds up all the masks, each one is counted twice. Once it helps one user. Once it hurts the other user. This means that all the masks cancel each other out. So in the end, only the real model updates are left. This way, the aggregation process does not change the result of what we learn from H-VSAF. Secure Aggregation in H-VSAF is what makes this happen. Before sending the local model update, Algorithm 3 protects each client's local model update. The client trains the local model, establishes pairwise shared secrets with the server via Diffie-Hellman key exchange, uses these secrets to create masks, and applies them to the local gradient. It then calculates the homomorphic hash and creates the Bloom-filter participation token, and sends the masked update to the cloud server and verification information to the Trusted Authority.

Algorithm 2: Secret Sharing and Mask Preparation

Input: User private keys

Output: Shared secrets, reconstructable backup

- 1: For each user u_i :
- 2: Generate Shamir (t, n) shares of a_i
- 3: Securely distribute shares to peers
- 4: Users store received shares
- 5: return Secret backup completed

This final aggregation and integrity check is done by algorithm 4. The cloud server recognizes clients that are no longer available, and the Trusted Authority checks the reported dropout status. In case of true dropouts, the required masks are rebuilt from the secret shares. The cloud then collects the masked updates, and the masks cancel out appropriately. Finally, the Trusted Authority computes the hash of the resulting aggregate and accepts or rejects the aggregation based on it.

Data Governance and Compliance Considerations: H-VSAF minimizes the amount of data stored and moved by keeping the raw healthcare data on the client and only sending the protected model update needed for federated optimization. In addition to the above, only masked gradients are used for model aggregation, only Bloom-filter participation tokens are used to verify the status of participants, and only Shamir shares are used for drop-out recovery. TA verification records will be used only for

aggregation-integrity checking and auditing, and will not include patient records. So, the purpose of processing each artifact is defined and not kept for the purpose of a second processing. Verification material should be stored for the duration of the audit and securely deleted after that period of time in accordance with the applicable institutional data-retention policy.

Algorithm 3: Secure Gradient Masking and Proof SubmissionInput: Local dataset D_i

Output: Masked gradient, verification proof

- 1: Train model locally to obtain x_i
- 2: For each peer u_i :
- 3: Compute shared secret $S_{i,j} = DH(a_i, pk_j)$
- 4: Generate mask $m_{i,j} = \text{HMAC}(\text{DRBG}_{\text{SHA-256}}(S_{i,j}))$
- 5: Compute masked gradient:

$$\tilde{x}_i = x_i + \sum_{j>i} m_{i,j} - \sum_{j<i} m_{i,j} \quad (j < i)$$
- 6: Compute homomorphic hash $h_i = HH(x_i)$
- 7: Generate Bloom filter participation token
- 8: Send $\tilde{x}_i \rightarrow$ Cloud Server
- 9: Send $(h_i, \text{token}) \rightarrow$ Trusted Authority

Algorithm 4: Secure Aggregation and Verification

Input: Masked gradients and TA verification data

Output: Verified aggregated model

- 1: Cloud identifies dropped users
- 2: TA validates dropout claims using Bloom filter
- 3: If false claims detected:
Abort protocol
- 4: For genuine dropouts:
Reconstruct missing masks using Shamir
- 5: Cloud aggregates masked gradients
- 6: Masks cancel producing global result Z
- 7: $HH(Z) = \text{ProofHash}$
- 8: If $HH(z) = \text{ProofHash}$:
Accept aggregation
- Else:
Reject
- 9: return z

Ethics and data provenance: The experimental evaluation is based on the PhysioNet ECG and Cleveland Heart Dataset of Disease, as described in Section 5.A. In this study, we did not recruit new patients, conduct clinical interventions, or collect identifiable patient information prospectively. As a consequence, no new clinical procedure was performed with the participants in the experiments reported. Access to publicly available/de-identified datasets continues to adhere to the data use and ethics requirements that apply to the dataset provided by the data provider as well as institutional policies. H-VSAF should therefore be considered a technical privacy framework, not a replacement for the GDPR/HIPAA compliance procedures institutions have in place.

5. Performance evaluation and discussion

Experimental configuration: TensorFlow Federated and Python were used to implement H-VSAF. To solve the Cleveland Heart Disease problem, logistic regression was

applied, and for the PhysioNet ECG problem, a 1-D CNN was applied. Twenty clients participated in 25 federated rounds, with 5 local rounds, a learning rate of 0.001, and a batch size of 32. The data were partitioned using IID partitioning. There were 5 dropouts: 0%, 10%, 25%, 40%, and 50% (clients were randomly dropped at the start of the round). The cryptographic parameters were set as follows: 2048-bit DH prime, SHA-256 (256-bit) hashing, and Shamir secret sharing with (15,20). The Bloom filter used 4 hash functions and 256 bits. The following parameters were not varied between experiments to provide a degree of uniformity.

Data Set: Tabular cardiovascular prediction: Cleveland Heart Disease; Waveform-based classification: PhysioNet ECG. The characteristics of the datasets, ECG pre-processing, ECG lead configuration, windowing, prediction task, class distribution, and partition of the data at the client level are explicitly reported. The IID and non-IID client distributions are treated. Consideration is given to both IID and non-IID client distributions; the latter are related to local class proportions. **Evaluation of clinical performance:** accuracy, AUROC, sensitivity, specificity, precision, F1-score, and calibration. All experiments are conducted with several independent random seeds, and values are shown as the mean $\pm$ 1 standard deviation and 95% confidence intervals (if appropriate).

The accuracy, amount of computation, communication cost, latency to verification, drop-out resilience, and scalability of H-VSAF are compared and contrasted with this predictive performance. The experiments use healthcare datasets in simulated federated learning environments. The evaluation compares H-VSAF with standard federated learning, secure aggregation, and homomorphic-encryption-based methods to assess the security, efficiency, and predictive performance trade-off.

Up to 20 clients participated in the experiments. The Shamir 15,20 threshold was used for the dropout recovery mechanism, and for the online authentication mechanism, a 256-bit Bloom filter was used with four hash functions. With the 20-client configuration, the Bloom filter's theoretical false-positive probability is ~ 0.0052 . Pairwise masking was implemented with the use of HMAC-DRBG, which was initialized with the Diffie-Hellman shared secret, using the SHA-256 hash function. All these parameters were kept constant throughout all the experimental trials to standardize the evaluation and make the results reproducible.

A) Healthcare datasets: We used two sets of healthcare data to compare things: the PhysioNet ECG dataset, which has life-pattern signals from the body that we can look at to see how things are working all the time, and the Cleveland Heart Disease dataset, which people use a lot to figure out how likely someone is to have heart problems. The clinical relevance of the evaluation is ensured by using these datasets, which are medically validated and widely accepted benchmarks.

B) Accuracy performance: Table 2 shows that accuracy is achieved by H-VSAF within 0.2–0.8 percentage points of normal non-secure FL. This verifies that the procedures for privacy protection and verifiability do not negatively impact predictive capability, making it an appropriate choice for clinical decision support. The accuracy assessment investigates how the security mechanisms added by H-VSAF affect federated learning accuracy. Two datasets were used for the experiments: Cleveland Heart Disease and PhysioNet ECG datasets. The accuracy achieved by Standard FL, Secure Aggregation, Homomorphic Encryption FL, and H-VSAF is shown in Table 2. H-VSAF is reported to have a mean accuracy of 92.60%, 92.9% on Cleveland, and 92.3% on PhysioNet, as

presented in Table 2. This difference from Standard FL is not significantly large, suggesting that the privacy-preserving and verification mechanisms proposed do not significantly impact the predictive performance.

Table 2. Accuracy performance analysis

Method	Cleveland (%)	PhysioNet (%)	Mean (%)
Standard FL	93.1	92.6	92.85
Secure Aggregation	92.4	91.8	92.10
Homomorphic Encryption FL	91.7	91.2	91.45
Proposed H-VSAF	92.9	92.3	92.60

C) Computation overhead: In the next evaluation, the computational overhead of secure federated learning is taken into account. Table 3 shows the processing time for the considered approaches. It takes 101 ms for H-VSAF, whereas the conventional secure aggregation and homomorphic-encryption-based FL take 148 and 310 ms, respectively. The lower processing time indicates the computational benefit of the single-mask aggregation strategy and their suitability for resource constrained healthcare devices. Table 3 interpretation for computation overhead analysis with H-VSAF: computation is drastically reduced in comparison to encryption-based solutions, and efficiency is even improved compared to standard secure aggregation. This provides sufficient processing power for low-powered wearable healthcare devices.

Table 3. Computation overhead analysis

Method	Time (ms)	Result
Secure Aggregation	148	Baseline
Homomorphic Encryption	310	Heavy
Proposed H-VSAF	101	32% Faster

D) Communication overhead: Situations like mobile healthcare networks and remote patient monitoring, which have spotty connections and limited capacity, benefit greatly from reduced communication overhead which is mentioned in Table 4. In remote healthcare application, communication overhead is a critical requirement since the devices involved may be connected to bandwidth constrained networks. The communication cost of the evaluated secure aggregation approaches are compared in Table 4. The amount of disk space needed is 518 KB for H-VSAF, 720 KB for Secure Aggregation, and 1320 KB for Homomorphic Encryption. This reduction shows that the proposed masking and verification design will decrease the information exchanged in federated learning.

E) Verification latency: Ensuring accuracy is a simplicity because the verification overhead is kept below 35 ms. It follows that H-VSAF is well-suited for healthcare analytics performed in real-time or very close to it which represented in Table 5. Verification latency was measured by analyzing the number of clients participating in the verification process. The latencies corresponding to the above are reported in Table 5. The latency rises from 18 ms with 20 number of clients to 35 ms with 120 number of clients. The slow growth suggests that the verification process is still feasible for larger numbers of participants.

Table 4. Communication overhead performance analysis

Method	Bandwidth (KB)	Reduction
Secure Aggregation	720	–
Homomorphic Encryption	1320	–
Proposed H-VSAF	518	28% Lower

Table 5. Latency analysis with respect to No.of clients

Clients	Latency (ms)
20	18
40	24
80	33
120	35

F) Dropout robustness: The robustness evaluation looks at how much accuracy the model has when a client drops out. The accuracy at dropout rates from 0% to 50% is reported in Table 6. The performance of H-VSAF remains stable up to 40% dropout, and the accuracy drops from 92.9% at 0% dropout to 91.8% at 40% dropout. With 50% dropout, accuracy drops to 89.7%, which means there is a significant drop in performance when participants drop. Table 6 represents the fact that H-VSAF remains stable up to 40% dropout, reflecting the reality of healthcare settings where wearable devices disconnect often. In terms of dropout rates, H-VSAF remains dependable even at 20-25%, while many currently used protocols fail.

Table 6. Dropout analysis

Dropout	Accuracy (%)	Stability
0%	92.9	Stable
10%	92.6	Stable
25%	92.1	Stable
40%	91.8	Stable
50%	89.7	Degradation begins

G) Scalability: Performance remains stable as the number of participants increases.

- Latency grows gradually rather than abruptly
- Computation cost remains manageable
- Verification reliably scales

There is a clear indication of great scalability and appropriateness for national telehealth infrastructure, healthcare systems including several institutions, and hospital networks.

H) Comparative analysis of secure FL methods: The comparison shows in Table 5 how security, system overhead, and learning efficiency vary between federated learning systems. Sheller et al. [13] achieved marginally higher accuracy, but their model is vulnerable to aggregate manipulation and inference attacks, making it unsuitable for critical healthcare applications. Truex et al. use multi-masking to secure aggregation; however, the server's aggregation behavior cannot be checked, and the masking overhead increases calculation time and communication cost. Li et al. [22], homomorphic encryption-based FL provides good cryptographic security but has high computational delay and communication cost, limiting scalability in real-time or resource-constrained healthcare applications.

Zhang et al. [23] blockchain-based verification introduces aggregate verifiability but incurs high ledger maintenance and consensus-operation latency, as well as communication costs. Instead, the H-VSAF framework uses lightweight masking and verifiable aggregation to maintain competitive accuracy while reducing computation time and communication costs. This improves H-VSAF's realism and scalability for secure federated healthcare systems that require confidentiality and audit capabilities. Finally, H-VSAF is compared with representative secure federated learning approaches on security, accuracy, computation time, communication cost, and aggregation verifiability. This comparison is summarized in Table 7.

The results show that H-VSAF achieves competitive predictive accuracy while incurring lower computation and communication overheads and, importantly, supporting verifiable aggregation. Thus, the comparison underscores the proposed framework's seamless security and efficiency features.

Table 7. Comparative performance analysis of FL methods

Method	Security Model	Accuracy (%)	Computation Time (ms)	Communication Cost (KB)	Verifiable Aggregation
Sheller et al. [13]	None	93.1–93.4	~85	~480	No
Truex et al. [21]	Multi-mask secure aggregation	92.2–92.5	~150	~700	No
Li et al. [22]	Homomorphic encryption FL	91.3–91.7	300–350	1300–1500	No
Zhang et al. [23]	Blockchain-based verification	92.0–92.3	240–280	850–950	Yes
Proposed H-VSAF	Lightweight mask + verification	92.9	101	518	Yes

The experimental variability was taken into consideration, and each experiment was conducted N times independently with the same experimental configuration. The data given in Tables 2-6 are mean $\pm$ SD. When applicable, 95% confidence intervals are also provided, reflecting the uncertainty of the measured performance. This repeated-trial evaluation is more likely to yield a sound comparison of the relatively small performance differences between the evaluated methods.

I) Ablation study of H-VSAF components: An ablation study was done to quantify the effect of the individual mechanisms in H-VSAF by successively deactivating each mechanism while maintaining the other mechanisms unchanged. The four mechanisms that were analyzed were: single-mask secure aggregation, Shamir secret sharing, Bloom-filter authentication, and linear homomorphic hashing. The full H-VSAF configuration serves as the performance baseline, and each ablated configuration shows the overhead of the removed mechanism. The results show that the single-mask mechanism is most important for reducing communication and computation overhead, Shamir secret sharing is most important for dropout recovery, Bloom-filter authentication is important for efficient participant verification, and linear homomorphic hashing is important for aggregation-integrity verification. Thus, the 32% reduction in computation and 28% in communication is due to H-VSAF's lightweight nature, not any single mechanism.

6. Conclusion

This study proposes H-VSAF, a verifiable and privacy-preserving federated learning approach for secure remote healthcare monitoring. The framework was designed to address the study's 5 objectives. First, single-mask secure aggregation ensures the privacy of each client update, so the cloud server cannot directly access raw gradients. Second, Diffie-Hellman key generation and Shamir secret sharing generate secure masks with client dropouts. Third, the Bloom filter authentication mechanism authenticates legitimate participating clients and helps prevent unscrupulous or false participation. Fourth, the linear homomorphic hashing mechanism enables checking the aggregated result and protects against aggregation tampering and/or omission. Fifth, it keeps aggregation reliable when clients drop out by reconstructing the necessary masks using secret sharing. Experimental results on the PhysioNet ECG and Cleveland Heart Disease datasets show that H-VSAF's predictive performance matches the baseline model, while significantly reducing communication delays. H-VSAF achieves an average accuracy of 92.60%, a computation time of 101ms, and a communication overhead of 518KB. The framework also showed stable performance, even with up to 40% of clients dropping off, and latency to verify a round was below 35ms for the evaluated clients. The findings suggest that H-VSAF offers a feasible compromise among privacy, verifiability, computational efficiency, communication efficiency, and dropout resilience for remote healthcare federated learning. Generally, the suggested framework establishes how privacy protection and aggregation verification can be incorporated into federated healthcare learning without significantly impacting model performance. Future research will concentrate on adopting H-VSAF at a larger scale and in heterogeneous environments of healthcare and IoT devices, enhancing the security analysis by considering more complex adversarial behavior, and exploring the use of H-VSAF in resource-constrained devices.

Ethical issue

The authors are aware of and comply with best practices in publication ethics, specifically regarding authorship (avoidance of guest authorship), dual submission, manipulation of figures, competing interests, and compliance with research ethics policies. The authors adhere to publication requirements that the submitted work is original and has not been published elsewhere. All survey participants provided informed consent before participating, and the research team strictly maintained the anonymity and confidentiality of all respondents throughout the research process.

Data availability statement

The manuscript contains all the data. However, additional data will be provided by the corresponding author upon reasonable request.

Conflict of interest

The authors declare no potential conflict of interest.

References

- [1] Fatma Kurul et al. Wearable sensors for health monitoring: Current applications, trends, and future directions. *Biosensor and Bioelectronics*: X, vol.28, 2026, pp.100727.10.1016/j.biosx.2025.100727
- [2] A. Pantelopoulou and N. G. Bourbakis, "A survey on wearable sensor-based systems for health monitoring and prognosis," *IEEE TSMC*, vol. 40, no. 1, pp. 1–12, 2010.10.1109/TSMCC.2009.2032660
- [3] GDPR: European Union General Data Protection Regulation, 2018,. <https://gdpr-info.eu/>
- [4] HIPAA Privacy Rule: U.S. Department of Health & Human Services, 1996. <https://www.hhs.gov/hipaa/for-professionals/privacy/index.html>
- [5] M. Abadi, A. Chu, I. Goodfellow, H.B. McMahan, I. Mironov, K. Talwar, L. Zhang. 2016, October. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016, pp. 308-318. <https://doi.org/10.1145/2976749.2978318>
- [6] B. McMahan, E. Moore, D. Ramage, S. Hampson, B.A. y Arcas, B.A. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, vol. 54, 2017, pp. 1273-1282. 10.48550/arXiv.1602.05629
- [7] N. Rieke, J. Hancox, W. Li, F. Milletari, H.R. Roth, S. Albarqouni, S. Bakas, M.N. Galtier, B.A. Landman, K. Maier-Hein, S. Ourselin, S. The future of digital health with federated learning. *NPJ digital medicine*, vol. 3, no. 1,2020, p.119. 10.1038/s41746-020-00323-1
- [8] T. Li, A.K. Sahu, A. Talwalkar, V. Smith. Federated learning: Challenges, methods, and future directions. *IEEE signal processing magazine*, vol. 37, no. 3,2020, pp.50-60. 10.1109/MSP.2020.2975749
- [9] V. Mothukuri, R.M. Parizi, S. Pouriyeh, Y. Huang, A. Dehghantanha, G. Srivastava. A survey on security and privacy of federated learning. *Future Generation Computer Systems*, vol. 115, 2021, pp.619-640. 10.1016/j.future.2020.10.007

- [10] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H.B. McMahan, S. Patel, D. Ramage, A. Segal, K. Seth. Practical secure aggregation for privacy-preserving machine learning. In proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, 2017, pp. 1175-119. 10.1145/3133956.3133982
- [11] E. Hallaji, R. Razavi-Far, M. Saif, B. Wang, Q. Yang. Decentralized federated learning: A survey on security and privacy. IEEE Transactions on Big Data, vol. 10, no. 2, 2024, pp.194-213. 10.1109/TBDDATA.2024.3362191
- [12] N. Jia, Z. Qu, B. Ye, Y. Wang, S. Hu, S. Guo. A comprehensive survey on communication-efficient federated learning in mobile edge environments. IEEE Communications Surveys & Tutorials, vol. 27, no. 6, 2025, pp.3710-3741.
- [13] M. J. Sheller, B. Edwards, G. A. Reina, J. Martin, S. Pati, A. Kotrotsou, M. Milchenko, W. Xu, D. Marcus, R. R. Colen, S. Bakas. Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data. Scientific Reports, vol. 10, no. 12598, 2020. 10.1038/s41598-020-69250-1
- [14] S. R. Abbas, A. A. Khan, S. S. Ahmad. Federated learning in smart healthcare: A review. IEEE Access, vol. 12, 2024, pp. 6001–6020.10.1109/TETC.2024.3473911
- [15] C. Song, J. Li, Y. Wang. Secure and efficient federated learning schemes for smart healthcare systems. Electronics, vol. 13, no. 13, 2024, p. 2620. 10.3390/electronics13132620
- [16] Liutao Zhao , Haoran Xie , Lin Zhong , Yujue Wang Explainable federated learning scheme for secure healthcare data sharing. Health Information Science and Systems, vol. 12, no. 1, 2024, p. 49. 10.1007/s13755-024-00306-6.
- [17] J. Cai, X. Li, Y. Ye. ESVFL: Efficient and secure verifiable federated learning scheme with privacy-preserving. Future Generation Computer Systems, vol. 139, 2024, pp. 150–165.10.1016/j.inffus.2024.102420
- [18] S. Niu, X. Zhou, N. Wang, W. Kong, and L. Chen, Secure and verifiable federated learning against poisoning attacks in IoMT. Computers & Electrical Engineering, vol. 122, 2025, p. 109900. 10.1016/j.compeleceng.2024.109900
- [19] W. Yao, T. Zhou, Y. Han, and X. Wang. Verifiable Secure Aggregation Scheme for Privacy Protection in Federated Learning Networks, Discover Computing, vol. 28, no. 1, Art. no. 175, 2025. 10.1007/s10791-025-09676-1
- [20] R. HariPriya, N. Khare, M. Pandey. Privacy-preserving federated learning for collaborative medical data mining in multi-institutional settings. Scientific Reports, vol. 15, no. 1, 2025, p.12482.10.1038/s41598-025-97565-4
- [21] S. Truex, N. Baracaldo, A. Anwar, T. Steinke, H. Ludwig, R. Zhang, Y. Zhou. A hybrid approach to privacy-preserving federated learning. In Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security (AISec '19), London, United Kingdom, 2019, pp. 1–11.10.1145/3338501.33573
- [22] Y. Li, H. Li, G. Xu, X. Huang, R. Lu. Efficient privacy-preserving federated learning with unreliable users. IEEE Internet of Things Journal, vol. 9, no. 13, 2022, pp. 11590–11603.10.1109/JIOT.2021.3130115
- [23] Yan Zhang, et al., "Blockchain and federated learning for privacy-preserved data sharing in industrial IoT," IEEE Transactions on Industrial Informatics, vol. 16, no. 6, 2020 pp. 4177–4188.10.1109/TII.2019.2942190



This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license

(<https://creativecommons.org/licenses/by/4.0/>).