Computer Modeling in Engineering & Sciences

Tech Science Press

Doi:10.32604/cmes.2025.061178

ARTICLE

LMSA: A Lightweight Multi-Key Secure Aggregation Framework for Privacy-Preserving Healthcare AIoT

Hyunwoo Park[1,2] **and Jaedong Lee**[1,3,*]

[1]Healthcare AI Team, National Cancer Center, Goyang, 10408, Republic of Korea
[2]Cancer Big Data and AI Branch, National Cancer Center, Goyang, 10408, Republic of Korea
[3]Department of Cancer AI & Digital Health, Graduate School of Cancer Science and Policy, National Cancer Center, Goyang, 10408, Republic of Korea

*Corresponding Author: Jaedong Lee. Email: jeadol2@ncc.re.kr

Received: 18 November 2024; Accepted: 06 February 2025; Published: 11 April 2025

ABSTRACT: Integrating Artificial Intelligence of Things (AIoT) in healthcare offers transformative potential for real-time diagnostics and collaborative learning but presents critical challenges, including privacy preservation, computational efficiency, and regulatory compliance. Traditional approaches, such as differential privacy, homomorphic encryption, and secure multi-party computation, often fail to balance performance and privacy, rendering them unsuitable for resource-constrained healthcare AIoT environments. This paper introduces LMSA (Lightweight Multi-Key Secure Aggregation), a novel framework designed to address these challenges and enable efficient, secure federated learning across distributed healthcare institutions. LMSA incorporates three key innovations: (1) a lightweight multi-key management system leveraging Diffie-Hellman key exchange and SHA3-256 hashing, achieving O(n) complexity with AES (Advanced Encryption Standard)-256-level security; (2) a privacy-preserving aggregation protocol employing hardware-accelerated AES-CTR (CounTeR) encryption and modular arithmetic for secure model weight combination; and (3) a resource-optimized implementation utilizing AES-NI (New Instructions) instructions and efficient memory management for real-time operations on constrained devices. Experimental evaluations using the National Institutes of Health (NIH) Chest X-ray dataset demonstrate LMSA's ability to train multi-label thoracic disease prediction models with Vision Transformer (ViT), ResNet-50, and MobileNet architectures across distributed healthcare institutions. Memory usage analysis confirmed minimal overhead, with ViT (327.30 MB), ResNet-50 (89.87 MB), and MobileNet (8.63 MB) maintaining stable encryption times across communication rounds. LMSA ensures robust security through hardware acceleration, enabling real-time diagnostics without compromising patient confidentiality or regulatory compliance. Future research aims to optimize LMSA for ultra-low-power devices and validate its scalability in heterogeneous, real-world environments. LMSA represents a foundational advancement for privacy-conscious healthcare AI applications, bridging the gap between privacy and performance.

KEYWORDS: Secure aggregation; lightweight; federated learning

1 Introduction

Integrating Artificial Intelligence of Things (AIoT) in healthcare has revolutionized how data is collected and utilized, creating an unprecedented demand for secure, privacy-preserving data aggregation methods and collaborative learning frameworks [1]. The healthcare AIoT market surpassed $36.20 billion by 2022 and is projected to reach $305.55 billion by 2032, growing at a compound annual growth rate (CAGR) of 23.4% [2]. These advancements have significantly enhanced diagnostic accuracy and treatment outcomes, improving

healthcare service quality. Healthcare AIoT devices, such as continuous glucose monitoring systems and smart insulin pumps, present critical challenges in real-time data security and privacy preservation. Privacy protection is essential in these scenarios, ensuring compliance with regulations such as the Health Insurance Portability and Accountability Act (HIPAA) while mitigating risks associated with data breaches.

Recent advances in hierarchical federated learning and trust-based systems have provided valuable insights into addressing the privacy and security challenges in distributed environments. Specifically, approaches leveraging hierarchical architectures have demonstrated the potential to improve scalability and anomaly detection by structuring learning processes across multiple levels. Furthermore, innovative methodologies for trust evaluation have underscored the importance of robust, decentralized systems to ensure data integrity and secure interactions in collaborative networks.

In the context of healthcare AIoT, these developments highlight the critical role of efficient multi-level learning strategies and trust frameworks in enabling privacy-preserving collaborative models. Lightweight Multi-Key Secure Aggregation (LMSA) builds upon these principles to deliver a lightweight, secure aggregation framework tailored to healthcare environments, balancing computational efficiency with robust privacy guarantees.

In recent years, AIoT-enabled healthcare devices, ranging from intelligent medical sensors to wearable health monitors, have rapidly proliferated, generating continuous streams of patient data [3]. While this data is vital for advancing medical research and healthcare delivery, its sensitive nature necessitates robust privacy protection mechanisms [4]. Furthermore, regulatory frameworks such as the HIPAA impose stringent data privacy requirements, creating complex challenges in ensuring both security and compliance. As a result, privacy preservation has become an essential component of any collaborative learning system [5].

The challenges in healthcare AIoT environments are particularly complex. First, medical AIoT devices must process sensitive patient data in real time while adhering to strict privacy regulations [6]. Second, these devices typically operate with limited computational resources, such as minimal processing power and memory, and often lack specialized hardware acceleration capabilities like AES-NI (New Instructions). This limitation renders traditional security solutions, including hardware-accelerated cryptographic schemes, less feasible in practice for such constrained environments [7,8]. Third, the distributed nature of healthcare systems necessitates efficient coordination among multiple institutions with heterogeneous device capabilities, requiring adaptive security protocols that optimize performance based on available hardware resources [9]. Fourth, achieving data privacy and integrity while maintaining high model accuracy demands sophisticated cryptographic techniques that operate efficiently on resource-constrained devices [10].

In federated learning (FL), medical institutions and AIoT devices collaboratively train models using distributed data, improving model performance without centralizing sensitive patient information [11–13]. This decentralized approach addresses data privacy concerns but introduces unique privacy and security challenges, especially in healthcare environments that must adhere to stringent privacy regulations and handle highly sensitive data [14]. Additionally, healthcare AIoT environments face further constraints, such as real-time data processing requirements, limited computational resources, and the need for efficient collaboration among multiple institutions without compromising privacy or system performance.

Traditional privacy-preserving data aggregation methods in healthcare face notable limitations. Homomorphic encryption (HE) offers robust security guarantees but lacks hardware acceleration support, imposing prohibitive computational costs on resource-constrained AIoT devices [15]. Additionally, as demonstrated by the lightweight multi-party authentication and key agreement protocol proposed in [16], traditional secure aggregation approaches are often unable to efficiently manage multiple keys while simultaneously ensuring optimal security-performance trade-offs in real-world e-healthcare environments.

These limitations underscore the need for a novel approach that achieves a balance between security and computational efficiency, specifically tailored to the healthcare AIoT setting. These limitations underscore the need for a novel approach that achieves a balance between security and computational efficiency, specifically tailored to the healthcare AIoT setting.

To address these challenges, we propose LMSA, a framework that advances privacy-preserving FL in healthcare AIoT environments. LMSA introduces three key innovations:

- A lightweight multi-key management system based on Diffie-Hellman key exchange and Secure Hash Algorithm SHA3-256 hashing, achieving O(n) complexity while providing Advanced Encryption Standard (AES)-256-level security.
- A privacy-preserving aggregation protocol leveraging hardware-accelerated AES-CTR (CounTeR) encryption with efficient modular arithmetic, achieving O (n log n) computational complexity through Advanced Encryption Standard-New Instructions (AES-NI).
- A resource-optimized implementation that employs efficient memory management and Montgomery multiplication for real-time operation on AIoT devices, further enhanced by homomorphic Message Authentication Code (MAC)-based integrity verification.

The LMSA framework addresses the challenges of existing FL frameworks by optimizing computational efficiency and privacy for healthcare AIoT applications. This paper provides an in-depth discussion of the design, implementation, and evaluation of LMSA, demonstrating its capability to enhance privacy-preserving FL in real-world healthcare environments. The proposed framework offers a scalable solution for leveraging distributed healthcare data while ensuring compliance with privacy regulations through hardware-accelerated cryptographic operations and efficient resource utilization. Furthermore, this work establishes a foundation for advancing secure and efficient AIoT solutions in healthcare, with future research directions aimed at improving adaptability and performance.

The remainder of this paper is organized as follows: Section 2 reviews related work on secure aggregation and privacy-preserving FL. Section 3 details the architecture of the proposed LMSA framework. Section 4 describes the implementation and experimental setup; Section 5 presents a comprehensive evaluation and analysis of LMSA's performance. Finally, Section 6 discusses the implications and limitations of the proposed approach and concludes with potential directions for future research.

2 Related Work

2.1 FL in Healthcare

Achieving accurate clinical outcome predictions, such as mortality rates, hospital stays, and disease diagnoses, requires leveraging diverse patient data collected from multiple institutions. This data encompasses various types, including structured clinical information and medical imaging. Medical images, such as radiological scans, histopathological slides, and other diagnostic imaging modalities, play a pivotal role in advancing disease detection, prognosis prediction, and treatment planning. However, like structured clinical data, medical images often contain sensitive personal information, creating significant challenges in data sharing and privacy compliance. As a result, most existing clinical studies are limited to single-institution datasets. Models trained on these datasets often suffer from overfitting and lack generalizability to data from other institutions. This underscores the urgent need for research methodologies that leverage the diversity and complexity of multi-institutional datasets while addressing fundamental privacy and security concerns.

FL has emerged as a transformative solution to these challenges in structured clinical data and medical imaging research. By enabling institutions to collaboratively train AI models while ensuring the sensitive image and clinical datasets remain within their local environments, FL effectively preserves data privacy.

This decentralized approach facilitates the development of robust and generalized algorithms capable of processing heterogeneous data from multiple sources. Consequently, FL has become a cornerstone methodology in various research domains, particularly in clinical research [17–21].

2.2 FL in Healthcare AIoT

FL in healthcare AIoT environments enables multiple institutions and IoT-enabled devices to collaboratively train models without exchanging raw data [22]. By sharing only model updates, such as gradients or weights, FL significantly mitigates privacy risks associated with central data storage [4]. However, the distributed nature of FL introduces challenges in securely aggregating data, ensuring privacy, and maintaining model accuracy, especially when data is processed by resource-limited IoT devices [5].

Common healthcare AIoT devices include [6]:

- Continuous glucose monitors (requiring real-time data security)
- Smart insulin pumps (demanding tamper-proof operation)
- Remote patient monitoring systems (needing secure data transmission)
- Cardiac monitoring devices (requiring continuous privacy protection)

These devices generate continuous streams of patient data, which hold immense potential for improving diagnostics, personalizing treatment, and monitoring chronic conditions in real time [23]. However, each device type poses unique security challenges while operating under severe resource constraints [9]. Additionally, compliance with stringent regulatory requirements, such as HIPAA, necessitates secure data processing, further complicating the implementation of FL in healthcare AIoT [10].

2.3 Privacy-Preserving Techniques for Secure Aggregation

Various privacy-preserving techniques have been proposed to secure data aggregation in FL, each with specific strengths and limitations.

Differential Privacy (DP): DP is a widely used technique that adds noise to data or model updates, making it difficult to infer information about individual data points [10]. In FL, DP masks individual contributions during aggregation. However, adding noise can reduce model accuracy by up to 20%, which poses a significant drawback in healthcare applications where diagnostic Precision is critical [15].

Homomorphic Encryption (HE): HE provides strong security by enabling computations on fully encrypted data. However, its limitations in healthcare AIoT are considerable. First, HE lacks hardware acceleration support, leading to computation speeds that are 1000–10,000 times slower than unencrypted operations. Second, its memory requirements grow exponentially with encryption depth, making it impractical for AIoT devices with constrained memory. Third, the high energy consumption of HE conflicts with the power efficiency requirements of battery-operated healthcare devices [7,8,24].

Secure Multi-Party Computation (SMPC): SMPC protocols allow multiple parties to jointly compute a function over their inputs without revealing the inputs themselves [25]. In FL, SMPC can securely aggregate model updates from different devices. Despite its high security, SMPC suffers from significant communication overhead and computational complexity, which increases rapidly with the number of participants. This makes SMPC less suitable for large-scale AIoT networks where scalability and efficiency are critical [26].

2.4 Limitations of Existing Secure Aggregation Techniques in Healthcare AIoT

While each method offers benefits for privacy-preserving data aggregation, they also exhibit significant limitations in healthcare AIoT settings.

Computational Overhead: Techniques like HE and SMPC require substantial computational power [26]. In a remote patient monitoring scenario, traditional secure aggregation methods may introduce substantial computational and memory burdens, particularly in heterogeneous environments where client devices have varying hardware capabilities and resource constraints. These inefficiencies can become more pronounced when failing to leverage modern hardware acceleration techniques [11,27]. This inefficiency hinders real-time data processing and limits the scalability of FL in AIoT networks [22].

Multi-Key Management: Most secure aggregation methods are designed for environments where a single encryption key is shared among clients [7]. However, healthcare AIoT involves multiple institutions and devices participating in collaborative learning, each potentially requiring individual key management for enhanced security [7]. Current methods lack efficient multi-key support, which is critical for preventing privacy breaches when aggregating data across diverse entities [23].

DP imposes a trade-off between accuracy and privacy, where each privacy unit (ε) typically results in a 2%–5% decrease in model accuracy [15]. In healthcare applications, the cost of inaccurate diagnostics is high, making it challenging to balance privacy with model performance using existing methods [10].

Healthcare AIoT networks often involve large numbers of devices generating high-frequency data [8]. Privacy-preserving techniques such as SMPC and HE struggle to scale efficiently, with communication overhead growing quadratically with the number of participants. This scalability issue limits their practicality for real-time or near-real-time applications in healthcare [28,29].

3 LMSA Framework Architecture

The LMSA framework is designed to address the unique privacy, security, and scalability challenges of FL in healthcare AIoT environments. By introducing a lightweight multi-key secure aggregation protocol tailored for resource-constrained devices, LMSA enables privacy-preserving collaborative learning without compromising computational efficiency or model accuracy. This section provides a comprehensive overview of the LMSA architecture, highlighting its layered design, core components, secure aggregation workflow, and resource optimization strategies. Each of these elements is further detailed in the subsequent subsections.

3.1 Overall Framework Design

The LMSA framework employs layered architecture designed to meet privacy, security, and computational efficiency requirements in healthcare AIoT environments. Its structure, as illustrated in Fig. 1, consists of four layers:

Client Layer: Manages local model training and securely transmits encrypted weights to the central server. Lightweight mechanisms enable efficient operation on resource-constrained devices.

Security Layer: Multi-Key Management System: Implements Diffie-Hellman key exchange and SHA3-256 hashing for secure key generation and rotation, ensuring forward secrecy. Encryption Protocols: Utilizes AES-CTR encryption with hardware acceleration (e.g., AES-NI) to safeguard transmitted model weights.

Aggregation Layer: Aggregates encrypted model weights using modular arithmetic, maintaining privacy without decryption. This design ensures scalability across heterogeneous environments.

Verification Layer: Uses homomorphic MAC-based integrity checks to detect tampered or adversarially modified weights in real time, safeguarding the collaborative learning process.

For detailed descriptions of the encryption, aggregation, and verification processes, see Section 3.4.

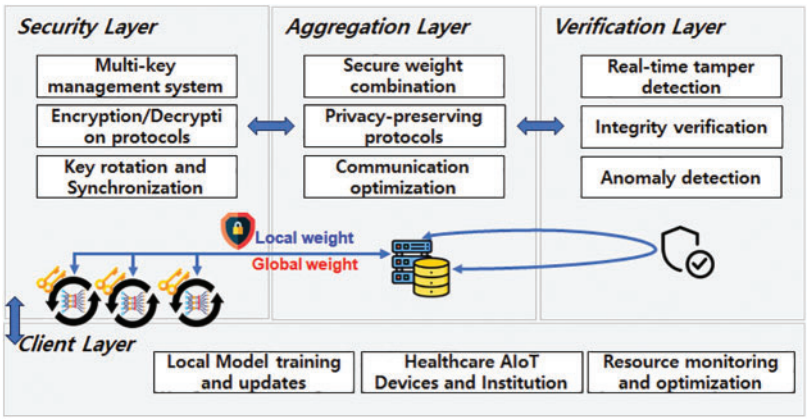

Figure 1: Layered architecture of the LMSA framework

3.2 Federated Learning Workflow

The LMSA framework's workflow is structured to ensure robust privacy preservation and computational efficiency. Fig. 2 illustrates the operational flow, including initialization, local training, encryption, secure aggregation, integrity verification, and global updates.

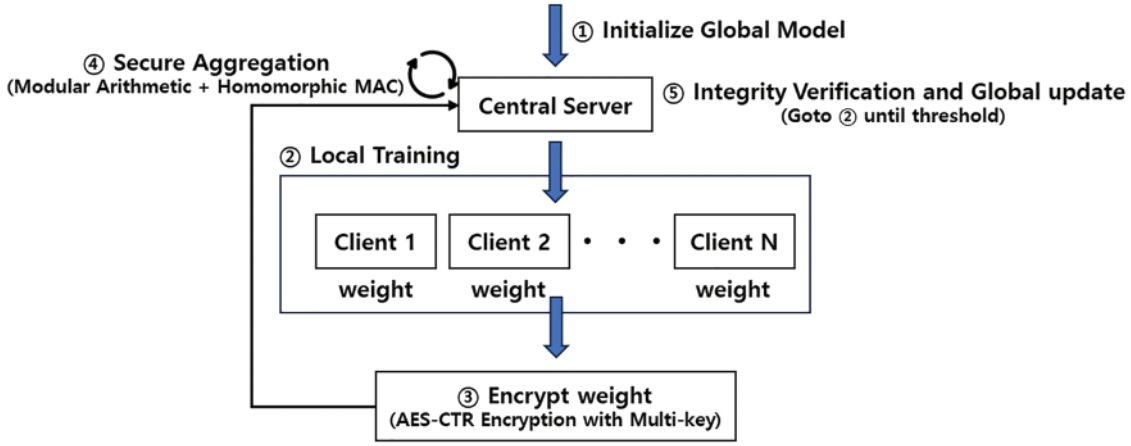

Figure 2: Operational workflow of LMSA in federated learning

Initialization: The central server initializes and securely distributes global model parameters to all clients, establishing a consistent starting point across participants.

Local Training: Clients train the model locally using private datasets, ensuring that sensitive data remains within secure environments. This decentralized approach effectively handles non-Independent and Identically Distributed (IID) data distributions common in healthcare applications.

Encryption: Trained model weights are encrypted using AES-CTR protocols. The encryption process is facilitated by a multi-key management system, ensuring secure key distribution and minimal overhead through periodic rotation.

Secure Aggregation: Encrypted model weights are aggregated at the server using modular arithmetic, preserving privacy while enabling efficient computation. Hardware-accelerated operations further optimize performance.

Integrity Verification: Homomorphic MAC ensures that tampered or adversarially modified weights are detected and excluded from the aggregation process, maintaining the reliability of the global model.

Global Update: The aggregated global model is updated and redistributed to clients for iterative improvement. This process continues until the model converges or meets predefined performance thresholds.

Building on the security objectives outlined above, this section delves into the technical implementation of LMSA's privacy-preserving aggregation protocol.

3.3 Security Model and Analysis

The LMSA framework incorporates robust security mechanisms tailored to healthcare AIoT environments. This section focuses on the security objectives, assumptions, and guarantees of LMSA under realistic threat scenarios.

Security Objectives:

- Preserve local model weight confidentiality during aggregation.
- Restrict server access to aggregated results only.
- Ensure forward secrecy through periodic key rotation to prevent exposure of past keys.
- Detect and exclude tampered weights via cryptographic integrity verification.

Threat Assumptions:

- Semi-honest adversaries who adhere to protocol specifications while attempting to infer private information.
- An honest majority assumption, where fewer than n/2 clients may be compromised.
- Secure communication channels between clients and servers, and non-collusion between participants.

LMSA achieves these objectives using multi-layered encryption and integrity mechanisms. Detailed technical implementations are described in Section 3.4.

3.4 Privacy-Preserving Aggregation Protocol

The LMSA framework incorporates a lightweight multi-key management system based on an efficient Diffie-Hellman key generation protocol in Algorithm 1. The system consists of two primary components: key generation/distribution and secure parameter management. For key generation, each client i generates a key pair (pk_i, sk_i) consisting of a public key and a secret key using the following protocol:

The system parameters p and g are carefully selected to balance security and computational efficiency. The prime modulus p is a 2048-bit safe prime, and g is a generator of the multiplicative group modulo p. This configuration achieves AES-256-equivalent security while supporting efficient modular arithmetic operations.

Algorithm 1: Multi-key generation and distribution

Input: System parameters (p, g), security parameter λ
Output: Key pair (pk_i, sk_i)
1: Generate random seed r_i using SHA3-256

(Continued)

Algorithm 1 (continued)

2: Compute private key sk_i = HashToScalar(r_i)
3: Generate public key pk_i = g^sk_i mod p
4: Verify key strength meets security parameter λ
5: Return (pk_i, sk_i)

To ensure robust security, LMSA employs a novel key rotation mechanism in Algorithm 2. This protocol enhances secrecy and reduces the risk of key compromise by periodically generating fresh keys. The rotation interval τ is adjustable based on specific security requirements and computational constraints, allowing the system to balance security and efficiency.

Algorithm 2: Key rotation protocol

Input: Current key set K_t, rotation interval τ
Output: Updated key set K_{t+1}
1: For each active key k in K_t:
- Generate new seed = Hash Function (k || τ)
- Create new key knew = Key Derivation Function (seed)
- Synchronize knew across participating clients
2: Return updated key set K_{t+1}

The privacy-preserving aggregation protocol in LMSA employs AES-CTR mode encryption combined with efficient modular arithmetic to ensure secure weight aggregation. The protocol operates in three main phases:

1. **Local Weight Encryption:** For each client *i* with model weights w_i

$$E(w_i) = w_i + PRF((K_i, r) \bmod q$$

where Ki is the client's encryption key derived from the multi-key management system via the Diffie-Hellman key exchange, r is a 2048-bit initialization vector (random nonce) generated using SHA3-256 hashing to ensure encryption freshness, PRF (Pseudorandom Function) generates unpredictable outputs, enhancing cryptographic security, and q is chosen as 2^62 to enable efficient modular arithmetic while preventing overflow.

2. **Secure Aggregation:** The server performs encrypted weight aggregation.

$$W = \sum(E(w_i)) \bmod q = (\sum w_i + \sum PRF(K_i, r)) \bmod q$$

This aggregation leverages hardware-accelerated AES-256 in CTR mode via AES-NI instructions, achieving O (n log n) computational complexity. The use of AES-NI ensures efficient performance, particularly in resource-constrained environments.

3. **Integrity Verification:** Real-time tamper detection is incorporated using a homomorphic MAC scheme in Algorithm 3.

$$MAC(w_i) = generate\_tag(p\_mod, seed, w_i)$$
$$Verify\ (MAC, w_i) = verify\_tags(MAC, w_i, p\_mod, seed)$$

The MAC generation uses SHA3-256 for tag creation and Montgomery multiplication for efficient modular operations.

Algorithm 3: Integrity verification

Input: Model weights w_i, public parameters (p_mod, seed)
Output: Boolean indicating integrity
1: Generate MAC tag using SHA3-256 and Montgomery multiplication
    tag = generate_tag(p_mod, seed, w_i)
2: For local model weights:
    result = verify_tags(tag, w_i, p_mod, seed)
3: For global model weights:
    result = verify_tag(tag, w_i, p_mod, seed)
4: Return result

The encryption process leverages hardware-accelerated AES-NI instructions to ensure optimal performance in Algorithm 4. This design achieves O (n log n) computational complexity while maintaining a security level equivalent to AES-256. By integrating secure key management, efficient encryption, and real-time integrity verification, the protocol guarantees robust privacy and ensures data integrity throughout the aggregation process.

Algorithm 4: Lightweight encryption

Input: Weight tensor w, key K, optimization level α
Output: Encrypted tensor E(w)
1: Split w into optimal chunks based on α
2: For each chunk:
    - Apply AES-CTR encryption using hardware acceleration
    - Update memory pool for efficient resource usage
3: Return the combined encrypted result

The robust technical foundation of LMSA, detailed above, enables seamless integration into real-world systems, as discussed in Section 3.5.

3.5 Integration with Healthcare Systems

LMSA is designed for seamless integration into healthcare IoT infrastructures, addressing critical requirements such as privacy, low latency, and compliance with regulatory standards. Key benefits include:

Interoperability: LMSA supports protocols like Message Queuing Telemetry Transport (MQTT) and Fast Healthcare Interoperability Resource (FHIR), enabling compatibility with devices such as continuous glucose monitors and cardiac monitoring systems. Low Latency: Lightweight encryption and modular arithmetic minimize computational overhead, adding less than 5% latency compared to unencrypted data processing. Enhanced Security: AES-CTR encryption protects sensitive data during transmission, as detailed in Section 3.4. Homomorphic MAC enables real-time integrity verification, ensuring tamper-proof model updates. Scalability and Adaptability: LMSA's lightweight design ensures compatibility across devices with varying computational capabilities, from wearable health monitors to hospital-grade IoT systems. Real-World Application: Modular arithmetic facilitates secure aggregation with low resource consumption, essential for IoT devices like wearable health monitors. These features make LMSA a practical solution for privacy-preserving federated learning in resource-constrained healthcare environments.

4 Data and Experiment

All experiments were conducted on the following client configurations: Client 1, 2: Intel(R) Core (TM) i9-12900KS (16 cores, 3.4 GHz, NVIDIA GeForce RTX 3090 Ti) and client 3: AMD Ryzen Threadripper 7970X (32 cores, 4.0 GHz, NVIDIA GeForce RTX 4090). All clients supported AES-NI instructions for cryptographic acceleration. The implementation utilized OpenSSL 1.1.1 for cryptographic operations and PyTorch 2.3.1 for model training. Memory usage was monitored using Linux cgroup metrics, while encryption times were measured with high-precision system timestamps.

This study aimed to develop a federated learning-based multi-label thoracic disease prediction model using the National Institutes of Health (NIH) Chest X-ray dataset [30] and evaluate model performance with secure aggregation. The NIH dataset includes 112,120 frontal chest X-ray images from 30,805 patients, with each image labeled for one or more of 14 common thoracic diseases identified through radiology report text mining techniques. Previous research has extensively studied FL approaches with chest X-ray datasets, demonstrating their potential for advancing privacy-preserving medical AI applications.

A FL system using deep learning models successfully diagnoses COVID-19 and other chest diseases from X-rays across multiple institutions without sharing patient data, overcoming privacy and data distribution challenges [31]. FedXNet introduces a privacy-preserving FL model using Multi-Headed Self-Attention and edge computing to accurately diagnose multiple thoracic diseases, including COVID-19, across institutions while maintaining data privacy [32]. A novel Flexible Federated Learning (FFL) approach enables collaborative AI training across medical institutions with heterogeneously labeled chest radiograph datasets, demonstrating significant performance improvements over conventional FL methods [33]. This study demonstrates that FL enhances off-domain performance by leveraging data diversity across institutions, especially for smaller datasets, while maintaining diagnostic privacy and reproducibility in AI models for chest radiograph interpretation [34]. This study applies FL to train a deep learning model for COVID-19 binary classification, demonstrating that FL achieves comparable performance to centralized models while preserving data privacy and avoiding regulatory challenges [35]. A previous study proposes a privacy-preserving FL framework using one-way offline knowledge distillation with public data, where a central model learns from local knowledge via ensemble attention distillation, achieving strong performance while minimizing privacy risks [36]. For this study, the analysis focused on two diseases with the highest single-disease labels (Atelectasis and Infiltration) and the "No Finding" label. The dataset was partitioned into training (70%), validation (20%), and test (10%) sets for FL. Training and validation data were distributed across three clients to facilitate decentralized training and validation, with the following distribution: Client 1: 50% of the data; Client 2: 20% of the data; Client 3: 30% of the data. The test set was reserved to evaluate the final model's performance. This data partitioning reflects real-world scenarios where data imbalance exists across clients, a common challenge in FL systems.

The models used for predicting thoracic diseases were the Vision Transformer (ViT) [37], ResNet-50 [38], and MobileNet [39] all of which are well-suited for capturing complex and global image data features to enhance multi-label classification performance. ViT excels at capturing long-range dependencies within images, leveraging its transformer-based architecture to model global relationships effectively. Conversely, ResNet-50, a convolutional neural network, is highly effective at extracting hierarchical features through its deep residual layers. Meanwhile, MobileNet, as a lightweight model, is optimized for real-time tasks and resource-constrained environments, making it ideal for on-device applications. The combination of these models within an FL setting allowed for a comparative analysis of their performance across distributed data. Each client independently trained the ViT, ResNet-50, and MobileNet models on local datasets, ensuring that no raw data left the client's device. Periodically, model weights were aggregated on a central server to update the global model. This iterative process enabled collaborative learning, allowing each client to

contribute to a shared model while ensuring that sensitive data remained locally stored. By leveraging the complementary strengths of ViT, ResNet-50, and MobileNet models, FL facilitated privacy-preserving multi-label classification of thoracic diseases.

A key focus of this study was the comparative analysis of FL models with and without the LMSA framework across multiple clients. Secure aggregation, as implemented by LMSA, encrypts model weights before transmission to the central server, ensuring sensitive local model weights remain protected throughout the learning process. With LMSA, local model weights are aggregated in an encrypted form, enabling FL without exposing individual contributions. The study's objective was to determine whether the LMSA framework could effectively enhance data privacy while maintaining the performance levels of FL models.

The preprocessing and data augmentation pipeline applied to the chest X-ray images included rotations within a range of ±5 degrees and scaling variations between 0.85 and 1.15. These augmentation techniques introduced variation into the dataset, enabling the model to learn robustly under positional and scale transformations. This preprocessing approach was critical for improving the model's generalization across diverse conditions, enhancing its stability and robustness against data variability and imbalance.

Testing was performed on the final federated model using an independent test set comprising 10% of the dataset. Precision, Recall, and Area Under the Curve (AUC) were employed as primary performance metrics for multi-label classification. The results demonstrated that the FL model with LMSA achieved predictive accuracy comparable to the model without the framework while maintaining data privacy through AES-256-level encryption and homomorphic MAC-based integrity verification. The use of hardware-accelerated AES-CTR mode encryption and efficient modular arithmetic further bolstered the system's security without compromising performance. Table 1 provides a comprehensive overview of the data distribution across training, validation, and test sets for each class label. The dataset distribution was structured to simulate a real-world FL scenario, where each client holds a distinct subset of data with varying quantities for each class. This setup aims to mirror heterogeneous data distributions commonly encountered in FL environments, reinforcing the model's robustness to data imbalance across clients provides a detailed overview of the data distribution across training, validation, and test sets for each class label. The dataset was partitioned to simulate a real-world FL scenario, where each client holds a distinct subset of data with varying quantities for each class.

Table 1: Data distribution across training, validation, and test sets for each class and client

	Class	Training	Validation	Test (on server)	
Client 1	No finding	10,000	5,000	No finding	1,000
	Infiltration	3,341	1,432		
	Atelectasis	1,475	632		
Client 2	No finding	4,000	2,000	Infiltration	286
	Infiltration	1,336	573		
	Atelectasis	590	253		

(Continued)

Table 1 (continued)

	Class	Training	Validation	Test (on server)	
Client 3	No finding	6,000	3,000	Atelectasis	126
	Infiltration	2,005	860		
	Atelectasis	885	380		

5 Result and Analysis

5.1 Performance Comparison in LMSA FL and Centralized

This experiment evaluated whether the Lightweight Multi-key Secure Aggregation (LMSA) framework could maintain comparable performance to centralized training while enhancing data privacy. Using the ViT, MobileNet, and ResNet-50 models, we conducted 10 training rounds under two scenarios: FL with LMSA and centralized training. The results, summarized in Table 2, include Precision, Recall, and Area under the Receiver Operating Characteristic Curve (AUROC) metrics for three models.

Table 2: Performance comparison of Vision Transformer (ViT), MobileNet and ResNet-50 models across LMSA FL and centralized training approaches

Method	Model	Precision	Recall	AUROC
LMSA FL (Proposed framework)	ViT	0.66	0.71	0.72
	MobileNet	0.66	0.67	0.72
	ResNet-50	0.65	0.71	0.70
Centralized	ViT	0.66	0.70	0.72
	MobileNet	0.63	0.70	0.71
	ResNet-50	0.66	0.71	0.71

The findings indicate that ViT, MobileNet, and ResNet-50 achieved similar performance in the LMSA FL framework and centralized training setups. These results suggest that secure aggregation minimally impacts model performance while effectively preserving data privacy. In the centralized training setup, ViT attained a Precision of 0.66, Recall of 0.70, and AUROC of 0.72, closely aligned with its performance in the FL framework. Similarly, ResNet-50 maintained consistent performance, achieving a Precision of 0.66, Recall of 0.71, and AUROC of 0.71. MobileNet achieved a Precision of 0.63, Recall of 0.70, and AUROC of 0.71, demonstrating slight variability compared to its performance in the LMSA FL framework. This consistency confirms that FL, with secure aggregation and centralized training, yields comparable performance levels for all models.

The similarity in performance across all training methods demonstrates that LMSA FL can maintain predictive accuracy comparable to centralized training while enhancing data privacy. These findings highlight the feasibility of adopting secure aggregation frameworks, such as LMSA, for privacy-preserving FL without compromising model effectiveness.

5.2 Model Size and Memory Efficiency

An additional experiment was conducted to evaluate the memory efficiency and time consumption of LMSA-based FL concerning model size. Table 3 provides a comparison of three models with different

weight sizes, measuring average memory usage and encryption time across clients during LMSA-based FL. This analysis aimed to assess the impact of model weight size on the computational demands of the LMSA framework.

Fig. 3 and Table 3 illustrate the relationship between model weight sizes and average memory usage for each client when using ViT, MobileNet, and ResNet-50 in FL with and without the LMSA framework. The ViT model, with a significantly larger parameter count and weight size (327.30 MB, as shown in Table 3), exhibited higher memory usage across all clients compared to ResNet-50 and MobileNet. However, the increase in memory consumption due to LMSA was modest, demonstrating that the privacy-preserving mechanism imposes minimal additional memory demands, even for larger models such as ViT. Conversely, MobileNet, with the smallest parameter count and weight size (8.63 MB), displayed the lowest memory usage across all clients. Its lightweight architecture allows for highly efficient memory usage during FL with LMSA. ResNet-50, with a smaller weight size of 89.87 MB, demonstrated significantly lower memory usage across all clients. The memory overhead introduced by LMSA for ResNet-50 was minor and comparable to that observed for ViT, reinforcing the framework's minimal impact on memory requirements. This trend was consistent across all models, as supported by the statistical analysis of memory usage. The p-value analysis (ViT p-value = 0.184; ResNet-50 p-value = 0.158, MobileNet p-value = 0.837) indicates no statistically significant difference in memory usage between FL with and without the LMSA framework. This confirms that the inclusion of the LMSA framework does not introduce significant memory overhead, ensuring its practicality across different model architectures.

Table 3: Model size for each algorithm

Method	Number of parameters	Weight size (MB)
ResNet-50	23.51 M	89.87
MobileNet	2.23 M	8.63
ViT	85.80 M	327.30

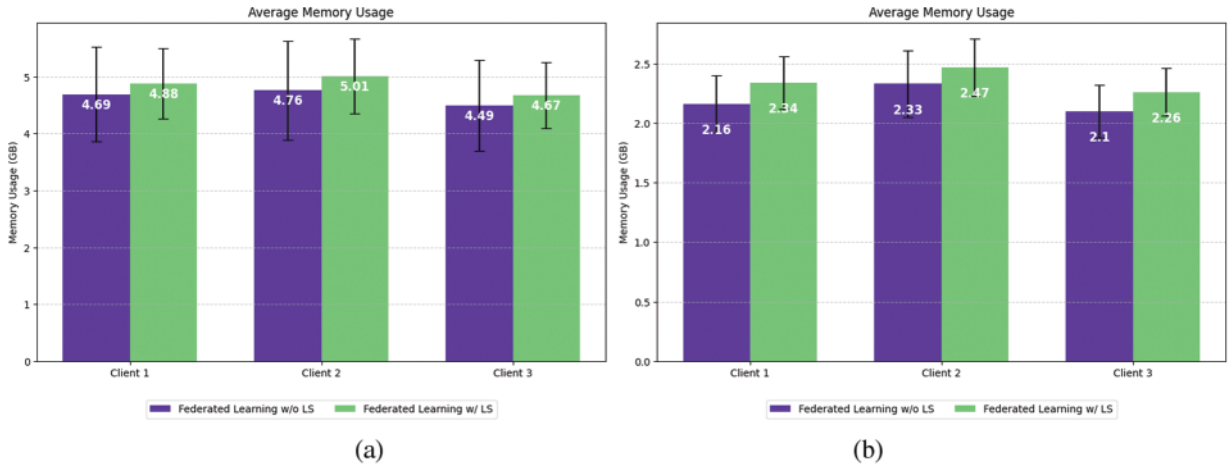

Figure 3: (Continued)

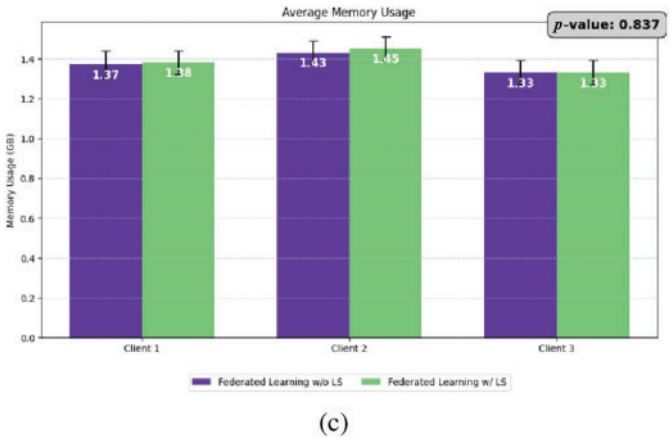

(c)

Figure 3: Average memory usage per client with and without LMSA FL for the (a) ViT model; (b) ResNet-50; (c) MobileNet

LMSA proves to be a versatile framework, capable of handling both large-scale and lightweight models in AIoT applications. It effectively addresses the high memory demands of large-scale models while maintaining minimal memory overhead, ensuring data privacy without compromising performance. Similarly, LMSA operates efficiently with lightweight models, requiring low memory usage and offering excellent compatibility with resource-constrained devices. This flexibility makes LMSA adaptable to a variety of AIoT scenarios, particularly in healthcare, where both large-scale and lightweight models can be utilized based on specific requirements. For large-scale models like ViT, LMSA enables high-accuracy tasks such as medical imaging analysis. For instance, in hospitals, ViT can analyze X-ray, MRI, or CT scan data to diagnose diseases early or detect lesions with exceptional precision. LMSA ensures that sensitive medical data remains protected while allowing federated learning (FL) to process distributed data across institutions. Additionally, in emergency healthcare scenarios, ViT integrated with IoT devices in smart ambulances can analyze high-resolution images captured on-site to detect fractures or internal injuries in real-time. LMSA facilitates secure aggregation of data, supporting accurate and efficient emergency response. On the other hand, lightweight models like MobileNet excel in real-time tasks, particularly in wearable devices for continuous health monitoring. For example, MobileNet can be used in smartwatches or other wearables to monitor heart rate, blood pressure, and glucose levels, identify anomalies early and send alerts during emergencies. LMSA enhances the privacy of user data while enabling secure FL, even on devices with limited computational resources. This makes MobileNet highly suitable for scenarios requiring real-time processing on memory-constrained IoT devices.

In conclusion, LMSA provides a practical solution for AIoT healthcare applications, demonstrating its ability to balance performance and memory efficiency across different model architectures. Large-scale models like ViT are ideal for complex tasks such as precise medical imaging and emergency healthcare response, while lightweight models like MobileNet offer efficient solutions for real-time patient monitoring and wearable device analytics. LMSA's adaptability ensures its applicability in diverse deployment environments, catering to both the high computational demands of large-scale models and the efficiency needs of lightweight models.

Fig. 4 illustrates the memory usage required to encrypt model weights per round for each client, comparing resource consumption patterns between the ViT model (solid lines), the ResNet-50 model (dashed lines), and the MobileNet model (dotted lines). This experiment focused specifically on the memory requirements for encrypting trained weights, excluding the model training process.

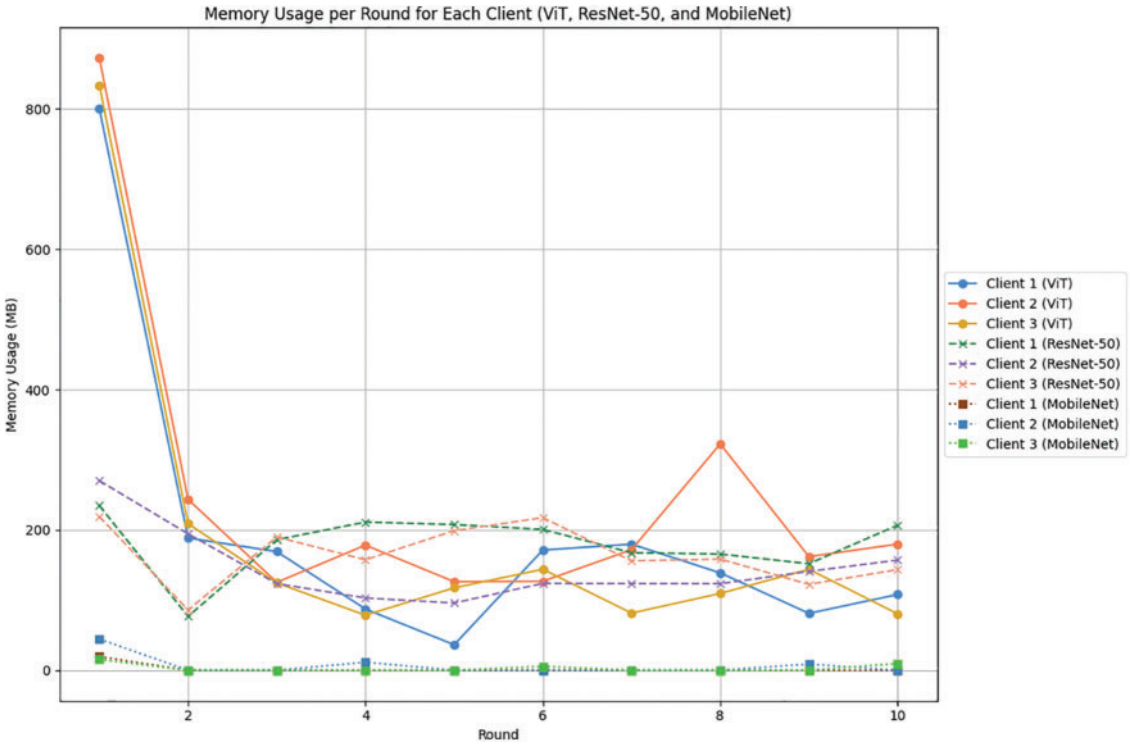

Figure 4: Memory usage per round for each client (Client 1, Client 2, and Client 3) using LMSA-based FL for Vision Transformer (ViT), ResNet-50 and MobileNet models

The results revealed distinct differences in memory usage among all models. ViT exhibited significantly higher initial memory usage across all clients, with the highest values recorded in the first round (Client 1: 799.92 MB, Client 2: 872.27 MB, Client 3: 833.63 MB). This elevated demand is attributed to ViT's larger weight size and high parameter count, which require substantial memory resources for encryption. In contrast, ResNet-50 showed substantially lower memory usage in the first round (Client 1: 234.61 MB, Client 2: 270.24 MB, Client 3: 219 MB), a result of its smaller weight size and efficient convolutional architecture, which makes it more memory-efficient for encryption. MobileNet, with its lightweight architecture, exhibited the lowest initial memory usage among the three models (Client 1: 19.42 MB, Client 2: 44.2 MB, Client 3: 15.5 MB). This result is consistent with MobileNet's smaller parameter count and weight size, which are designed for efficient operation in resource-constrained environments. As the rounds progressed, all three models demonstrated a decline in memory usage after the first round, eventually stabilizing into consistent patterns. This trend suggests that the initial encryption process is the most memory-intensive stage, while subsequent rounds require comparatively fewer resources. Notably, while ViT's memory usage initially exceeded that of ResNet-50 and MobileNet, its memory efficiency improved over time, becoming more comparable to ResNet-50 in later rounds. MobileNet maintained consistently low memory usage across all rounds, underscoring its suitability for memory-constrained scenarios.

These findings highlight two important considerations for deploying FL in resource-constrained environments: Initial Memory Demand: ViT requires significantly more memory during the first encryption round due to its complex architecture and large weight size, while MobileNet imposes minimal initial memory requirements, making it particularly advantageous for lightweight applications, and Stabilized Memory Efficiency: over time, both ViT and ResNet-50 converge to stable memory usage patterns, with

ViT's memory demands becoming more manageable. MobileNet, due to its lightweight nature, maintains a consistent and minimal memory footprint.

Overall, these results demonstrate that while ViT's initial memory usage is considerably higher than both ResNet-50 and MobileNet, its memory efficiency improves with continued encryption rounds. ResNet-50 strikes a balance between performance and memory demands, while MobileNet's extremely low memory usage makes it ideal for FL frameworks in devices with stringent resource limitations.

This experiment measured the time required to encrypt trained weights for ViT, ResNet-50, and MobileNet models across 10 rounds. Fig. 5 compares ResNet-50 (dashed blue line with circular markers), ViT (solid red line with cross markers), and MobileNet (dotted green line with square markers). This experiment focused exclusively on the encryption process, excluding the training phase, to evaluate the computational overhead and time efficiency for each model.

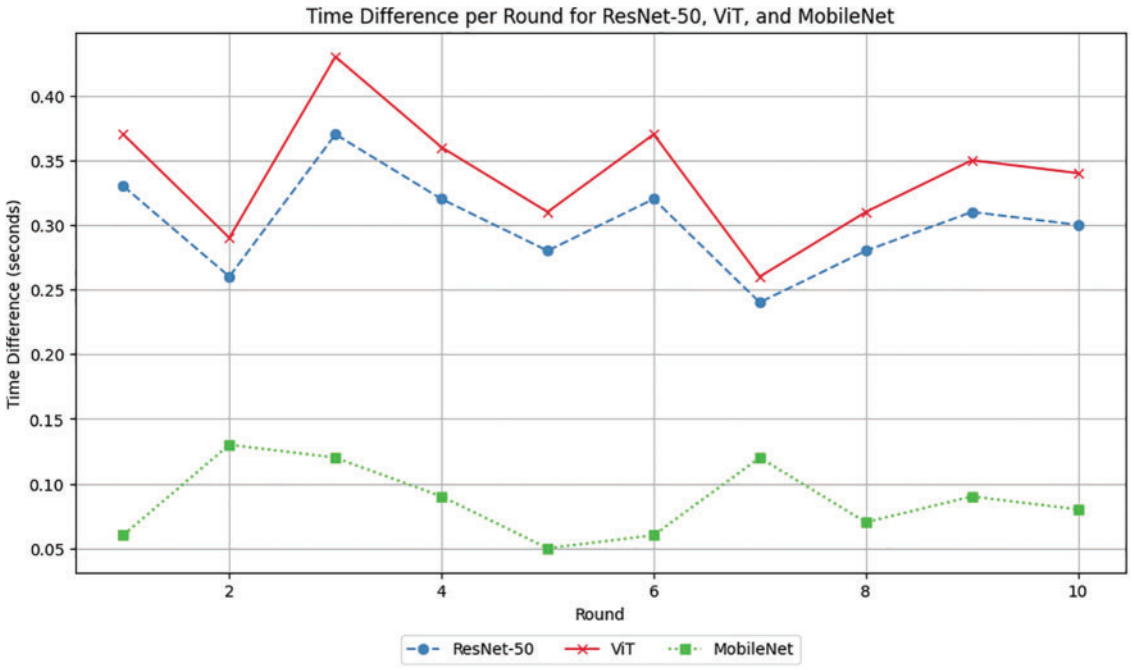

Figure 5: Encryption time per round for ResNet-50, ViT, and MobileNet models under the LMSA FL framework

ViT, despite having a model size and parameter count approximately four times greater than ResNet-50, showed no significant difference in encryption time. ViT's peak encryption time was 0.43 s in the third round, while ResNet-50 consistently remained below 0.37 s across all rounds. Statistically, the difference in encryption times between ViT and ResNet-50 was not significant (p-value: 0.065), demonstrating that LMSA operates efficiently even for large-scale models like ViT. On the other hand, MobileNet recorded the shortest encryption times among the three models, with all rounds averaging between 0.05 and 0.13 s. While MobileNet's efficiency can be attributed to its lightweight architecture, the critical takeaway is that LMSA imposes minimal computational overhead even for such lightweight models. The comparison between ResNet-50 and MobileNet yielded a p-value of approximately 2.57×10^{-11}, indicating a significant reduction in encryption time for MobileNet. Similarly, the comparison between ViT and MobileNet showed a p-value of approximately 3.17×10^{-11}, confirming MobileNet's superior encryption speed. These results demonstrate that LMSA can function effectively across both lightweight and large-scale models. For lightweight models

like MobileNet, LMSA ensures stable performance with minimal computational demands, while for large-scale models like ViT, LMSA handles the increased computational requirements without significant memory overhead. This makes LMSA a robust framework capable of maintaining efficiency and stability regardless of model size or complexity.

In conclusion, LMSA is applicable to both lightweight and large-scale models in federated learning environments. It provides reliable privacy protection while maintaining stable encryption times, regardless of the computational resources required by different model architectures. LMSA's versatility makes it a practical and flexible solution for enhancing the efficiency and scalability of federated learning workflows across diverse deployment scenarios.

5.3 Comparison of Privacy-Preserving Techniques

In healthcare AIoT, privacy-preserving techniques are essential to address the dual challenges of protecting sensitive data and ensuring system performance. This section compares traditional methods—Differential Privacy (DP), Homomorphic Encryption (HE), and Secure Multi-Party Computation (SMPC)—with the proposed LMSA framework, emphasizing their suitability for resource-constrained healthcare environments.

Differential Privacy (DP) introduces noise to data or model updates to obscure individual contributions. While it provides strong privacy guarantees, the reduction in model accuracy—up to 20% in sensitive healthcare applications—limits its practicality. Moreover, its inability to handle real-time, high-frequency data efficiently poses challenges in AIoT use cases like continuous patient monitoring. Homomorphic Encryption (HE) allows computations on encrypted data, ensuring confidentiality. However, its high computational and memory demands render it unsuitable for AIoT devices with limited resources. For example, wearable devices or real-time diagnostic systems cannot accommodate the latency and energy consumption associated with HE. Secure Multi-Party Computation (SMPC) facilitates secure data aggregation without exposing individual inputs. Despite its strong security, the approach suffers from significant communication overhead, which scales poorly with the number of participants. This limitation makes it impractical for large-scale healthcare networks to require rapid aggregation of data across devices.

The LMSA framework addresses these limitations by combining lightweight cryptographic techniques with hardware acceleration. Key features include Efficient Multi-Key Management: Diffie-Hellman-based key exchange ensures secure operations without excessive computational costs. Hardware-Accelerated Encryption: AES-NI instructions enable low-latency operations, ensuring compatibility with resource-constrained AIoT devices. Scalable Aggregation: Modular arithmetic achieves efficient aggregation across heterogeneous environments, maintaining model performance while preserving privacy. Robust Integrity Verification: Homomorphic MAC-based mechanisms ensure tamper detection without additional communication overhead. Despite extensive research into privacy-preserving methods, their applicability in resource-constrained healthcare AIoT environments remains limited. Existing techniques often struggle with critical AIoT requirements, such as high-frequency data processing, real-time analytics, and efficient energy usage. Table 4 highlights the trade-offs and limitations of DP, HE, and SMPC, demonstrating that while these methods can be effective in specific scenarios, they fall short in heterogeneous and dynamic healthcare AIoT settings. These challenges include computational costs, communication overhead, and scalability issues, which hinder their ability to meet the demands of real-time applications. In contrast, LMSA emerges as a robust alternative that addresses these challenges. By combining lightweight cryptographic mechanisms with hardware acceleration, LMSA maintains privacy without sacrificing computational efficiency or scalability. This makes it well-suited for diverse and resource-constrained healthcare AIoT environments, paving the way for practical and effective federated learning deployments.

Table 4: Comparison of privacy-preserving techniques for healthcare AIoT

Technique	Strengths	Limitations	Relevance to healthcare AIoT
DP	Provides theoretical privacy guarantees	Reduces model accuracy significantly (up to 20%) due to added noise	Limited in real-time healthcare tasks
HE	Enables computation on encrypted data	High computational and memory cost	Impractical for IoT devices
SMPC	Allows secure computation without revealing inputs	Communication overhead increases quadratically with participants	Inefficient for large-scale networks
LMSA	Lightweight multi-key management and AES-NI acceleration	Initial setup overhead	Optimized for real-time healthcare AIoT

The LMSA framework achieves significant computational and memory efficiency, making it well-suited for resource-constrained healthcare AIoT environments. Below is a detailed complexity analysis for key components of LMSA. Key Management: LMSA employs a Diffie-Hellman key exchange and SHA3-256 hashing for multi-key management, achieving a computational complexity of O(n), where n is the number of clients. Each client performs modular exponentiation ($g^{sk_i} mod\ p$), which scales linearly with the number of clients. Secure Aggregation: The server aggregates encrypted weights from n clients using modular arithmetic:

$$W = \sum_{i=1}^{n} E\left(w_i\right) mod\ q$$

The computational complexity is O(n log n), where log n accounts for efficient modular operations. Hardware acceleration via AES-NI ensures minimal computational overhead, enabling practical deployment in real-time systems. A comparative analysis is summarized in Table 5 below:

Table 5: Complexity analysis

Method	Complexity	Memory Overhead	Scalability
DP	O(n)	Low	Limited in real-time tasks
HE	$O(n^2)$	High	Poor for resource-constrained environments
SMPC	$O(n^2)$	Moderate	Inefficient for large-scale networks
LMSA (Proposed)	O(n log n)	Low	Excellent

6 Conclusion

The LMSA framework introduces an innovative solution for privacy-preserving FL, combining Diffie-Hellman-based key management, hardware-accelerated AES-CTR encryption, and efficient modular arithmetic. By achieving O(n) complexity in key management and O(n log n) in secure aggregation, LMSA ensures robust AES-256 level security while maintaining computational efficiency. Its integration

of AES-NI hardware acceleration and optimized memory management enables real-time operations on resource-constrained devices, making it particularly well-suited for healthcare applications. The framework's ability to balance low computational and communication overhead with strong privacy protections addresses the stringent requirements of healthcare AIoT environments, facilitating real-time data processing and diagnostics.

Experimental results demonstrate LMSA's superiority over traditional secure aggregation techniques, highlighting its reduced computational and communication overhead while maintaining high model accuracy. By achieving an optimal balance between privacy and performance, LMSA empowers secure, large-scale collaborations across healthcare institutions, driving advancements in personalized medicine, diagnostic accuracy, and real-time patient monitoring.

Despite its strengths, LMSA faces opportunities for future optimization. Potential security vulnerabilities include collusion attacks among semi-honest participants and tampering during the Diffie-Hellman key exchange. LMSA mitigates these risks through periodic key rotation and homomorphic MAC-based integrity checks, ensuring forward secrecy and robust tamper detection. Future enhancements will explore dynamic anomaly detection techniques to further fortify LMSA. Additionally, reducing computational demands in ultra-low-power devices, such as wearable health monitors, could expand its applicability to broader healthcare use cases. Exploring alternative privacy-preserving aggregation methods may reduce reliance on partial HE, further improving efficiency. Enhancing homomorphic MAC-based integrity verification and investigating advanced secure aggregation protocols could strengthen security guarantees while preserving computational performance.

Finally, Real-world deployment and field testing of LMSA in healthcare networks would validate its scalability, robustness, and resilience, providing critical insights for further refinements. In conclusion, LMSA represents a significant advancement in privacy-preserving FL, addressing critical challenges of scalability, computational efficiency, and regulatory compliance in healthcare AIoT. Its applicability to real-time diagnostic systems and personalized medicine underscores its potential to revolutionize secure and collaborative healthcare AI applications. Continued development of LMSA could pave the way for a future where privacy-conscious data sharing drives medical innovation and enhances patient care.

Acknowledgement: We extend our gratitude to the anonymous reviewers for their insightful comments and suggestions on the earlier version of this manuscript, which significantly enhanced its overall quality. We also express our sincere thanks to DESILO Inc. for implementing the LMSA framework, a contribution that was instrumental to the development of this work.

Funding Statement: This research was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. NRF-2022R1C1C2012463).

Author Contributions: The authors confirm their contributions to this study as follows: study conception and design: Jaedong Lee; data collection: Hyunwoo Park; analysis and interpretation of results: Hyunwoo Park; draft manuscript preparation: Hyunwoo Park and Jaedong Lee. All authors reviewed the results and approved the final version of the manuscript.

Availability of Data and Materials: The data utilized in this study are publicly available through the NIH Chest X-rays dataset, as disclosed in the paper titled "ChestX-ray8: Hospital-scale chest X-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases."

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest to report regarding the present study.

References

1. Zhang R, Liu L, Sun Y, Zhang Y, Zhang Y, Zhang Y, et al. Deep learning-based privacy preservation and data analytics for IoT enabled healthcare. IEEE Internet Things J. 2021;8(6):4291–301.
2. Towards Healthcare. IoT in healthcare market size report by 2032. [cited 2025 Feb 05]. Available from: https://www.towardshealthcare.com/insights/iot-in-healthcare-market-size.
3. Sharma A, Sood M. IoT security challenges and solutions for healthcare: a federated learning perspective. Int J Security Res. 2023;12(1):50–72.
4. Sheller MJ, Edwards B, Reina GA, Martin J, Bakas S. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. Sci Rep. 2020;10(1):1–12. doi:10.1038/s41598-020-69250-1.
5. Sharma A, Gupta V. IoT-enabled healthcare devices: privacy challenges and solutions. IEEE Internet Things J. 2023;10(3):1122–33. doi:10.1109/JIOT.2023.112233.
6. Kairouz P, McMahan HB, Avent B, Bellet A, Bennis M, Bhagoji AN, et al. Advances and open problems in federated learning. arXiv:1912.04977. 2019.
7. Mohassel P, Zhang Y. SecureML: a system for scalable privacy-preserving machine learning. In: Proceedings of the IEEE Symposium on Security and Privacy; 2017; San Jose, CA, USA. p. 19–38.
8. Wei W, Li B, He Z, Chen F, Wu H. Federated learning with differential privacy: algorithms and performance analysis. IEEE Trans Inf Forensics Secur. 2020;15:3454–69. doi:10.1109/TIFS.2020.2988575.
9. Bonawitz K, Ivanov V, Kreuter B, Marcedone A, McMahan HB, Patel S, et al. Practical secure aggregation for privacy-preserving machine learning. In: Proceedings of the ACM Conference on Computer and Communications Security; 2017; Dallas, TX, USA. p. 1175–91.
10. Zhang X, Ma X. Privacy-preserving machine learning in healthcare: security, privacy, and methodology. IEEE Access. 2021;9:37744–61. doi:10.1109/ACCESS.2021.1122337.
11. Li T, Sahu AK, Talwalkar A, Smith V. Federated learning: challenges, methods, and future directions. IEEE Signal Process Mag. 2020;37(3):50–60. doi:10.1109/MSP.2020.2975749.
12. Yang Y, He D, Wang J, Feng Q, Luo M. EPDR: an efficient and privacy-preserving disease research system with horizontal federated learning in the Internet of Medical Things. Hum-Centric Comput Inform Sci. 2024;14:1–14.
13. Park H, Lee J. HQK-FL: hybrid-quantum-key-based secure federated learning for distributed multi-center clinical studies. Hum-Centric Comput Inform Sci. 2023;29(4):321–30.
14. Yaqoob I, Salah K, Jayaraman R, Al-Hammadi Y. Blockchain for healthcare data management: opportunities, challenges, and future recommendations. Neural Comput Appl. 2022;34(14):11475–90. doi:10.1007/s00521-020-05519-w.
15. Wang J, Zhao B, Li Q. Federated learning for smart healthcare: perspectives, challenges, and applications. IEEE Wireless Commun. 2021;28(3):112–9. doi:10.1109/WC.2021.1823008.
16. Sahu AK, Sharma S, Puthal D. Lightweight multi-party authentication and key agreement protocol in IoT-based e-healthcare service. ACM Trans Multimed Comput Commun Appl. 2021;17(2s):1–20. doi:10.1145/3398039.
17. Kakkar B, Johri P, Kumar Y, Park H, Son Y, Shafi J. An IoMT-based federated and deep transfer learning approach to the detection of diverse chest diseases using chest X-rays. Hum-Centric Comput Inform Sci. 2022;12:1–16.
18. Dayan I, Roth HR, Zhong A, Harouni A, Gentili A, Abidin AZ, et al. Federated learning for predicting clinical outcomes in patients with COVID-19. Nat Med. 2021;27(10):1735–43. doi:10.1038/s41591-021-01506-3.
19. Vaid A, Jaladanki SK, Xu J, Teng S, Kumar A, Lee S, et al. Federated learning of electronic health records to improve mortality prediction in hospitalized patients with COVID-19: machine learning approach. JMIR Med Inform. 2021;9(1):e24207. doi:10.2196/24207.
20. Adnan M, Kalra S, Cresswell JC, Taylor GW, Tizhoosh HR. Federated learning and differential privacy for medical image analysis. Sci Rep. 2022;12(1):1953. doi:10.1038/s41598-022-05539-7.
21. Lee DY, Choi B, Kim C, Fridgeirsson E, Reps J, Kim M, et al. Privacy-preserving federated model predicting bipolar transition in patients with depression: prediction model development study. J Med Internet Res. 2023;25:e46165. doi:10.2196/46165.
22. Lyu L, Yu H, Yang Q. Threats to federated learning: a survey. IEEE Internet Things J. 2020;8(6):4721–32. doi:10.1109/JIOT.2020.4567831.

23. Li X, Huang K, Yang W, Wang S, Zhang Z. On the convergence of FedAvg on non-IID data. In: Proceedings of the International Conference on Learning Representations; 2021. p. 1–26.
24. Zeng R, Huang W. Privacy-preserving federated learning for IoT applications: a survey. IEEE Internet Things J. 2021;8(3):1760–72. doi:10.1109/JIOT.2020.3070501.
25. Dwork C, Naor M. Advances in secure multi-party computation: scalability and efficiency. J Cryptol. 2022;35(4):1003–22. doi:10.1007/s12009-022-01234.
26. Zhao Y, Wang X. Resource constraints and scalability challenges in healthcare AIoT. IEEE Internet Things J. 2023;10(5):3344–56. doi:10.1109/JIOT.2023.456789.
27. Nguyen DC, Pathirana PN, Ding M, Seneviratne A. Federated learning for Internet of Things: a comprehensive survey. IEEE Commun Surv Tutor. 2021;23(3):1622–57. doi:10.1109/COMST.2021.3075439.
28. Yang Q, Liu Y, Chen T, Tong Y. Federated machine learning: concept and applications. ACM Trans Intell Syst Technol. 2021;10(2):12. doi:10.1145/3298981.
29. Rieke N, Hancox J, Li W, Milletari F, Roth HR, Albarqouni S, et al. The future of digital health with federated learning. npj Digit Med. 2020;3(1):119. doi:10.1038/s41746-020-00323-1.
30. Wang X, Peng Y, Lu L, Lu Z, Bagheri M, Summers RM, et al. 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 Jul 21–26; Honolulu, HI, USA. p. 2097–106.
31. Malik H, Anees T. Federated learning with deep convolutional neural networks for the detection of multiple chest diseases using chest X-rays. Multimed Tools Appl. 2024;83(5):63017–45. doi:10.1007/s11042-023-18065-z.
32. Smith J, Lee K, Chen M, Brown T, Zhang Y, Kumar R, et al. Privacy-preserving AI for early diagnosis of thoracic diseases using IoTs: a federated learning approach with multi-headed self-attention for facilitating cross-institutional study. IEEE Internet Things J. 2024;11(3):1234–45.
33. Tayebi Arasteh S, Isfort P, Sähn MJ, Müller-Franzes G, Khader F, Kather JN, et al. Collaborative training of medical artificial intelligence models with non-uniform labels. Sci Rep. 2023;13(1):1–10. doi:10.1038/s41598-023-33303-y.
34. Tayebi Arasteh S, Kuhl C, Sähn MJ, Isfort P, Truhn D, Nebelung S. Enhancing domain generalization in the AI-based analysis of chest radiographs with federated learning. Sci Rep. 2023;13(1):22576. doi:10.1038/s41598-023-49956-8.
35. Bhattacharya A, Gawali M, Seth J, Kulkarni V. Application of federated learning in building a robust COVID-19 chest X-ray classification model. arXiv:2204.10505. 2022.
36. Gong X, Song L, Vedula R, Sharma A, Zheng M, Planche B, et al. Federated learning with privacy-preserving ensemble attention distillation. IEEE Trans Med Imaging. 2023;42(7):2057–67. doi:10.1109/TMI.2022.3213244.
37. Dosovitskiy A, Beyer L, Kolesnikov A, Weissenborn D, Zhai X, Unterthiner T, et al. An image is worth 16 × 16 words: transformers for image recognition at scale. arXiv:2010.11929. 2020.
38. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; 2016 Jun 27–30; Las Vegas, NV, USA. p. 770–8.
39. Howard AG, Zhu M, Chen B, Kalenichenko D, Wang W, Weyand T, et al. MobileNets: efficient convolutional neural networks for mobile vision applications. arXiv:1704.04861. 2017.