Computers, Materials & Continua

Tech Science Press

https://doi.org/10.32604/cmc.2026.075495

ARTICLE

Quantum-Resistant Secure Aggregation for Healthcare Federated Learning

Chia-Hui Liu[1] **and Zhen-Yu Wu**[2,*]

[1]Department of Electronic Engineering, National Formosa University, Huwei Township, Yunlin County, Taiwan
[2]Department of Maritime Information and Technology, National Kaohsiung University of Science and Technology, Cijin District, Kaohsiung City, Taiwan

*Corresponding Author: Zhen-Yu Wu. Email: hkwu668@nkust.edu.tw

Received: 02 November 2025; Accepted: 02 February 2026; Published: 12 March 2026

ABSTRACT: Federated Learning (FL) enables collaborative medical model training without sharing sensitive patient data. However, existing FL systems face increasing security risks from post quantum adversaries and often incur non-negligible computational and communication overhead when encryption is applied. At the same time, training high performance AI models requires large volumes of high quality data, while medical data such as patient information, clinical records, and diagnostic reports are highly sensitive and subject to strict privacy regulations, including HIPAA and GDPR. Traditional centralized machine learning approaches therefore pose significant challenges for cross institutional collaboration in healthcare. To address these limitations, Federated Learning was introduced to allow multiple institutions to jointly train a global model while keeping local data private. Nevertheless, conventional cryptographic mechanisms, such as RSA, are increasingly inadequate for privacy sensitive FL deployments, particularly in the presence of emerging quantum computing threats. Homomorphic encryption, which enables computations to be performed directly on encrypted data, provides an effective solution for preserving data privacy in federated learning systems. This capability allows healthcare institutions to securely perform collaborative model training while remaining compliant with regulatory requirements. Among homomorphic encryption techniques, NTRU, a lattice based cryptographic scheme defined over polynomial rings, offers strong resistance against quantum attacks by relying on the hardness of the Shortest Vector Problem (SVP). Moreover, NTRU supports limited homomorphic operations that are sufficient for secure aggregation in federated learning. In this work, we propose an NTRU enhanced federated learning framework specifically designed for medical and healthcare applications. Experimental results demonstrate that the proposed approach achieves classification performance comparable to standard federated learning, with final accuracy consistently exceeding 0.93. The framework introduces predictable encryption latency on the order of hundreds of milliseconds per training round and a fixed ciphertext communication overhead per client under practical deployment settings. In addition, the proposed system effectively mitigates multiple security threats, including quantum computing attacks, by ensuring robust encryption throughout the training process. By integrating the security and homomorphic properties of NTRU, this study establishes a privacy preserving and quantum resistant federated learning framework that supports the secure, legal, and efficient deployment of AI technologies in healthcare, thereby laying a solid foundation for future intelligent healthcare systems.

KEYWORDS: Federated learning (FL); homomorphic encryption; NTRU cryptography; healthcare data privacy; quantum-resistant security

1 Introduction

With the rapid advancement of artificial intelligence (AI), a wide range of AI-driven applications has emerged across various domains, bringing substantial economic benefits and significantly improving the convenience of daily life. In particular, progress in big data technologies has further expanded the scope and impact of AI systems. However, training high-performance AI models typically requires large volumes of data, and conventional machine learning approaches often rely on centralized aggregation of training datasets. In privacy-sensitive domains such as medicine and healthcare, this centralized paradigm poses serious risks of data leakage, which has become increasingly critical in today's highly digitized medical environments.

In the medical domain, patients' personal information, clinical records, and diagnostic reports are highly sensitive and must be strictly protected by healthcare institutions. Such data cannot be shared without explicit patient consent, and healthcare providers are required to clearly inform patients of the intended use of their data during collection and processing. As a result, cross-institutional data sharing becomes a significant challenge, particularly for AI model training, which often benefits from integrating data across multiple healthcare organizations to improve predictive performance. To address this challenge, Google introduced Federated Learning (FL) in 2016, enabling multiple institutions to collaboratively train a global AI model without sharing local datasets. By keeping raw data localized, FL provides an effective mechanism for privacy preservation while facilitating AI adoption in privacy-sensitive domains. For medical centers and healthcare providers that are constrained by strict data-sharing regulations, federated learning enables decentralized collaborative training without direct access to patient-level data, thereby allowing participation without compromising privacy.

Federated learning adopts a decentralized data architecture in which patient data remain securely stored within individual hospitals or clinics. Each participating institution performs local model training using its own edge computing resources or on-premises servers, and only model parameters are transmitted to a central server for aggregation. The aggregated global model is then redistributed to participating institutions for further local optimization, enabling iterative refinement while maintaining compliance with strict data protection regulations such as HIPAA. This privacy-preserving paradigm supports a wide range of medical AI applications, including diagnosis assistance and treatment optimization. Despite these advantages, existing federated learning systems largely rely on classical cryptographic or statistical privacy mechanisms that exhibit inherent limitations in security-critical medical environments. Differential privacy (DP) techniques often degrade model accuracy due to noise injection, while additive homomorphic encryption schemes such as Paillier introduce substantial computational and communication overhead and, more importantly, lack resistance to quantum attacks. Consequently, current FL security solutions face a critical challenge in simultaneously achieving strong privacy protection, acceptable system overhead, and long-term post-quantum security. This challenge is further intensified by the rapid progress of quantum computing, which threatens the security foundations of widely deployed public-key cryptosystems. Encryption schemes based on integer factorization or discrete logarithm assumptions, such as RSA and ECC, are expected to become vulnerable in the presence of large-scale quantum computers. The NIST Post-Quantum Cryptography initiative has explicitly highlighted these risks and emphasized the urgent need for cryptographic primitives that can withstand quantum adversaries [1]. As a result, post-quantum cryptography (PQC) has emerged as a key research direction for ensuring long-term security in critical applications, including healthcare systems.

To address these challenges, federated learning in medical environments requires cryptographic mechanisms that are not only privacy-preserving but also quantum-resistant and compatible with encrypted aggregation. In this work, we propose a privacy-preserving federated learning framework based on NTRU (N-th degree Truncated Polynomial Ring Units) to enhance data encryption and privacy protection throughout the federated learning process. NTRU is a lattice-based cryptographic scheme defined over polynomial rings, whose security relies on the hardness of lattice problems, particularly the Shortest Vector Problem (SVP). Since no known polynomial-time algorithms can efficiently solve this problem, NTRU is widely regarded as a quantum-resistant cryptographic primitive. By integrating NTRU encryption into federated learning, secure transmission of model updates can be ensured, thereby mitigating potential quantum computing threats. In addition, NTRU supports limited homomorphic properties that allow certain algebraic operations to be performed directly on ciphertexts, which is sufficient for secure aggregation in federated learning scenarios [2]. Homomorphic encryption enables computations to be carried out on encrypted data while producing results equivalent to those obtained in the plaintext domain, without revealing sensitive information. Semi-homomorphic encryption schemes, such as Paillier for addition and RSA or ElGamal for multiplication, support only a single type of operation, whereas fully homomorphic encryption (FHE) supports both additive and multiplicative operations [2,3].

These properties make homomorphic encryption particularly suitable for federated learning, as they enable essential computations to be performed without decrypting sensitive data. In summary, this paper explores the application of NTRU-based cryptographic techniques in federated learning systems to enhance privacy protection and security against quantum threats. The proposed approach strengthens the security of federated learning while supporting the deployment of AI technologies in medical and other privacy-sensitive domains by safeguarding sensitive data.

The main contributions of this paper are summarized as follows:

1. A critical security gap in existing healthcare-oriented federated learning frameworks is identified, in which classical cryptography and differential privacy mechanisms fail to simultaneously provide post-quantum security and acceptable system overhead.
2. An NTRU-enhanced federated learning framework is proposed to enable encrypted model aggregation under lattice-based, quantum-resistant security assumptions.
3. A complete system implementation and quantitative experimental evaluation are presented, demonstrating that encrypted federated learning can preserve model accuracy while maintaining predictable and practical computational and communication overhead.
4. Stable convergence behavior under multi-client medical federated learning settings is demonstrated, supporting the feasibility of the proposed framework for real-world healthcare deployment.

2 Literature Review

2.1 NTRU Encryption

The NTRU public-key cryptosystem is a post-quantum cryptographic algorithm [4]. It is a cryptographic scheme based on a polynomial ring, with security that can be reduced to the difficulty of solving the Shortest Vector Problem (SVP) in a lattice. This level of computational hardness is known to resist currently known quantum attacks. A lattice refers to a set of points in n-dimensional Euclidean space $\mathbb{R}^n$, formed by integer linear combinations of a given basis. Lattice-based encryption algorithms achieve cryptographic security by leveraging the computational hardness of vector problems within lattices. Among the most common of these problems are the Shortest Vector Problem (SVP) and the Closest Vector Problem (CVP), both of which are central to the security foundations of lattice-based cryptosystems.

1. **Shortest Vector Problem (SVP):** SVP refers to the problem of finding the shortest non-zero vector within a lattice, that is, identifying a non-zero vector $v \in L$ that minimizes the Euclidean Norm $\| v \| = \sqrt{v_1^2 + v_2^2 + \ldots + v_n^2}$. The security of NTRU relies primarily on the hardness of the SVP on lattices, making it resistant to attacks based on large-scale computations or factoring mechanisms. NTRU's security architecture provides strong resistance to quantum computational attacks, a capability that current cryptographic schemes, such as RSA and ECC, lack in the face of quantum computing advancements.

2. **Closest Vector Problem (CVP)**: The security of lattice-based cryptosystems such as NTRU is closely related to the hardness of the Closest Vector Problem (CVP), where the objective is to find a lattice vector $v \in L$ that is closest to a target vector $w \in R^n$, minimizing the distance $\| w - v \|$. NTRU is constructed over the truncated polynomial ring $R = \mathbb{Z}_q [x] / (x^N - 1)$, where any polynomial can be expressed as $f(x) = f_0 + f_1 x + \ldots + f_{N-1} x^{N-1}$, and multiplication is performed modulo $x^N - 1$. In this algebraic structure, polynomial multiplication corresponds to a cyclic convolution, denoted by the operator $\star$, and defined as:

$$f(x) \star g(x) = \sum_{k=0}^{N-1} \left(\sum_{(i+j) \equiv k (mod N)} f_i g_j \right) x^k \tag{1}$$

While both SVP and CVP are fundamental lattice problems, the security of NTRU primarily relies on the hardness of the SVP, which forms the cryptographic basis of the proposed scheme.

2.2 Utilizing Federated Learning in Medical Applications

Federated Learning has drawn significant interest in healthcare, primarily due to its capacity to maintain patient confidentiality. Recently, FL has shown considerable effectiveness in aiding patients with COVID-19 [5] as well as in the management of chronic conditions [6]. NVIDIA has contributed to the field by launching an open-source software development kit (SDK), NVIDIA FLARE, to facilitate FL's application in medical research [7]. Beyond medical imaging, FL has been employed in smart healthcare frameworks, utilizing data from wearable technology [8,9]. Sheller et al. [10] illustrated FL's potential to leverage diverse data sources for brain tumor segmentation without data sharing across institutions. Their FL model achieved a 2.63% increase in Dice similarity coefficient compared to models trained solely on single-site data. Khoa et al. developed the FedMCRNN model to predict sleep quality using multimodal data from wearable devices. Their model, evaluated on many-to-one and many-to-many configurations, achieved an accuracy of 96.77% and 68.72%, respectively, with the FL configuration surpassing the non-FL baseline in the many-to-one case [9,11]. Federated learning has been shown to be effective for collaborative medical data analysis across multiple institutions without sharing sensitive patient data. In particular, prior work has shown that federated learning can achieve competitive model performance in medical applications while reducing data transmission and preserving privacy compared to centralized learning approaches [12].

Recent medical federated learning studies have demonstrated that collaborative training across institutions can achieve competitive performance in tasks such as medical image analysis, biosignal classification, and disease diagnosis, while keeping raw patient data locally within each organization. However, many existing medical FL deployments still assume a trusted aggregator or rely on plaintext aggregation, which leaves the training process vulnerable to inference attacks (e.g., gradient inversion or membership inference) when model updates are exposed. These observations motivate the need for secure aggregation mechanisms that provide strong confidentiality guarantees without sacrificing practicality in healthcare settings.

Beyond security-focused designs, recent federated learning studies have also explored communication-efficient and personalized training strategies to improve scalability and adaptability in practical deployments, including healthcare-related scenarios [13].

2.3 Security and Privacy Threats in Federated Learning

Federated Learning systems have certain vulnerabilities in terms of security. For instance, model distribution, interactive training, and parameter aggregation all present potential risks. During the training process, gradients and communication may inadvertently expose user privacy. This section explores various security defense mechanisms currently employed in FL, including differential privacy, secure multi-party computation, and homomorphic encryption, as detailed below:

1. **Differential Privacy (DP):** Differential Privacy employs a noise-injection mechanism to obfuscate data, ensuring that privacy-sensitive information is not disclosed at any level, thereby enabling lightweight privacy protection. The privacy loss is quantified by parameters (ϵ, δ), where ϵ represents the privacy budget. A smaller ε value indicates a stricter privacy protection standard, while a lower δ signifies higher confidence in the privacy guarantee.
   Noise: Several factors influence the properties of noise in Differential Privacy (DP), such as data type and processing capability. For example, continuous noise is particularly prominent with numerical data, while discrete noise is typically applied to nominal data. Given an acceptable level of utility loss that ensures the indistinguishability of protected subjects, the impact of perturbed noise is relatively independent of noise type (e.g., Gaussian noise, Laplace noise). In Federated Learning, DP adds noise to gradients, providing flexibility in training data protection at the participant level or data record level. This approach can significantly reduce the risk of reverse data reconstruction by perturbing shared gradients, thereby increasing resistance to gradient-based privacy attacks such as Deep Leakage from Gradients (DLG) and Model Inversion Attacks (MIA) [14].
   Central and Local Differential Privacy (CLDP): Differential Privacy (DP) can be categorized based on where noise is introduced, comprising Central Differential Privacy (CDP) [15] and Local Differential Privacy (LDP) [16]. In CDP, noise is added to the aggregated results at a central location, providing privacy guarantees at the participant level. In contrast, LDP employs a randomized perturbation algorithm that adds noise locally on each participant's data before it is sent to the server. LDP is implemented at the sample level to obscure the contribution of specific samples within a participant's dataset, allowing participants to customize their privacy budget. Noise in LDP can be injected into local update parameters or computed gradients. For instance, differential private stochastic gradient descent with LDP provides a higher level of privacy protection at the sample level. By comparison, CDP is often insufficient to prevent sample-level gradient leakage, which could compromise the privacy of participants' training data [17].
   Defense against Attacks: Differential Privacy (DP) can effectively mitigate poisoning attacks and enhance system robustness. Research has shown that CDP is more effective than LDP in reducing the accuracy of backdoor attacks while also offering better utility. To prevent inference of unrelated attributes, Naseri et al. applied differential privacy to mitigate property inference attacks (PIA) [17], conducting experiments on gender classification within the Labeled Faces in the Wild (LFW) dataset. The experimental results indicate that LDP is ineffective against PIA, as it only provides privacy guarantees at the sample level and cannot directly prevent the leakage of population-level attributes.

2. **Secure Multi-party Computation (SMC):** SMC enables multiple participants to collaboratively compute a model or function without revealing their private inputs to one another, thus preserving privacy [18]. In the context of Federated Learning, secure aggregation protocols built on SMC can protect participants' updates [19,20]. In these protocols, each participant masks their local updates through secret sharing and pairwise masking. For example, the protocol may compute the weighted average of update vectors from a random subset of participants. As a result, when the masked updates are aggregated by the aggregator, the additional random factors are canceled out. Consequently, the aggregator can only access the aggregated model information without any details about individual models.
   Secret Sharing: Secret sharing is an encryption technique in which a secret is divided into multiple parts and distributed among parties. The secret can only be reconstructed when all parties combine their respective shares. Secret sharing employs consistency checks to verify correctness, ensuring that

each participant adheres to the protocol. This privacy-preserving mechanism, based on secret sharing, enables secure aggregation of gradient updates, allowing each participant to safely share their local updates with others [21]. In Federated Learning, common secret sharing protocols include additive secret sharing and Shamir's secret sharing. Additive secret sharing utilizes random numbers generated by each party to create respective shares, while Shamir's Secret Sharing employs secure polynomial interpolation within a finite field.

Pairwise Masking: Pairwise masking is a method where mutually untrusting participants use Diffie–Hellman key exchange to generate pairwise masks. When an attack compromises a participant using secret sharing, other honest participants may provide all their shares to the aggregator to cancel out the masked updates intended for the victim. However, this approach risks revealing the victim's data. Pairwise masking prevents model updates from being disclosed when the adversarial aggregator attempts to reconstruct participant noise [20]. Unfortunately, due to the large number of participants, the aggregation result often fails to offer sufficient privacy protection for individual participants. Therefore, pairwise masking is typically unsuitable for cross-device federated learning (CDFL) scenarios. Multi-party computation (MPC) protocols generally incur high transmission costs. For example, in secret sharing, each participant must interact with others to distribute shares, leading to communication costs that increase exponentially with the number of participants in FL. Consequently, optimizing the communication costs of secure multi-party computation in FL is a key area of focus for researchers.

3. **Homomorphic Encryption (HE):** Homomorphic encryption is a cryptographic defense technique that enables secure computations without direct access to plaintext, as the result of operations on ciphertexts corresponds to that of operations on plaintexts.

   Homomorphic encryption schemes may support either additive operations (e.g., Paillier [3]) or multiplicative operations (e.g., RSA, ElGamal), forming the class of semi-homomorphic encryption (semi-HE). When both operations are supported, the scheme becomes fully homomorphic encryption (FHE) [2]. In federated learning, homomorphic encryption is widely adopted to protect the transmission of encrypted model updates between participants and the central aggregator [22–24].

Several different security defense techniques for Federated Learning have been discussed above. A summary of these defense methods and their respective characteristics is presented in Table 1.

Table 1: Analysis of security defense techniques for federated learning systems.

Technique	Method	Generative Adversarial Network	Model Inversion Attack	Property Inference Attack	Data/Model Poisoning	Model Inversion	Cost
Defense Against Attacks (DP)	Increase the noise level by adjusting noise parameters	✓	✓	✗	✓	✓	Accuracy reduction due to noise injection
Secure Multi-party Computation (SMC)	Adds a secure aggregation protocol to prevent data leakage	✗	✓	✓	✗	✓	Increased transmission cost
Homomorphic Encryption (HE)	Encrypt the locally updated model before uploading	✓	✓	✓	✗	✓	Reduced training efficiency due to encryption overhead

Summary and Limitations of Existing Defense Mechanisms: The above defense techniques provide important protection for federated learning systems; however, each approach exhibits inherent limitations in privacy-sensitive medical scenarios. Differential privacy-based methods often introduce accuracy degradation due to noise injection, which is undesirable for clinical decision-making. Secure multi-party computation and secret sharing protocols typically incur significant communication overhead and scalability challenges as the number of participants increases. In addition, many homomorphic-encryption-based solutions rely on classical public-key cryptosystems, which may become vulnerable in the presence of quantum adversaries. These limitations highlight the need for a secure aggregation framework that simultaneously offers confidentiality, scalability, and post-quantum resistance, motivating the integration of lattice-based cryptography in this work which leads to the NTRU-enhanced system design presented in Section 3.

Recent studies have further investigated verifiable and collusion-resistant federated learning frameworks to strengthen aggregation integrity and privacy guarantees. For example, PriVeriFL introduces aggregation verifiability under privacy constraints, while other works address robustness against collusion attacks and enhance secure aggregation efficiency in distributed learning environments [25,26].

Discussion and Motivation: However, these techniques are often evaluated in isolation, making it difficult to understand their suitability for medical federated learning systems under post-quantum threat models. While Table 1 summarizes representative defense mechanisms against privacy and security threats in federated learning, it primarily focuses on the technical characteristics of individual protection strategies. To further clarify the research landscape and highlight the positioning of this work, a higher-level comparison of existing federated learning approaches is provided in Table 2. This comparison emphasizes system-level properties, including medical applicability, secure aggregation capability, and post-quantum security considerations.

Table 2: Outcome comparison of related work in secure medical federated learning.

Category	Representative Approach	Medical FL	Secure Aggregation	Post-Quantum Secure	Key Limitation
Plain FL	(medical FL applications)	✓	✗	✗	Plaintext updates may leak sensitive information
DP-based FL	(DP/LDP-based methods)	✓	✗	✗	Privacy–utility trade-off (accuracy degradation)
SMC/Secret Sharing FL	(SMC/secret sharing protocols)	✓	✓	✗	High communication cost; limited scalability
Classical HE-based FL	(Paillier/other classical HE)	✓	✓	✗	Computational overhead; not post-quantum secure

(Continued)

Table 2 (continued)

Category	Representative Approach	Medical FL	Secure Aggregation	Post-Quantum Secure	Key Limitation
Proposed Scheme	This work	✓	✓	✓	Requires careful parameter selection to balance security and efficiency

Based on the above analysis, existing federated learning approaches for healthcare either lack strong cryptographic guarantees, suffer from scalability issues, or fail to address post-quantum security concerns. In the next section, we introduce an NTRU-enhanced federated learning framework that provides efficient encrypted aggregation together with formal security guarantees under standard lattice hardness assumptions.

3 The Proposed NTRU-Enhanced Federated Learning System across Multiple Medical Institutions

This section presents the proposed NTRU-enhanced federated learning system designed to support secure model collaboration among multiple medical institutions. The system allows healthcare institutions to jointly train a global learning model while ensuring that raw patient data and individual model updates remain private at all times. By embedding a lattice-based NTRU cryptographic mechanism into the FL workflow, the system provides post-quantum confidentiality, encrypted aggregation capability, and robust protection against reconstruction or inference attacks. The overall architecture, cryptographic foundation, and workflow integration are described in the following subsections.

Section 3.1 introduces the system architecture and the interactions among the trusted center (TC), the aggregator, and the participating medical institutions. Section 3.2 details the NTRU operations, including key generation, encryption, and decryption. Section 3.3 explains how these cryptographic primitives are integrated into the federated learning workflow, while Section 3.4 presents the implementation-ready algorithmic description.

Before presenting the system workflow and cryptographic operations, all notations used throughout Section 3 are summarized in Table 3.

Table 3: Notations used in the proposed system.

Symbol	Description
N	Degree of the truncated polynomial ring used in NTRU
p, q	Small and large moduli for plaintext and ciphertext operations
$f(x), g(x)$	Small polynomials selected during key generation
$f_p(x), f_q(x)$	Modular inverses of $f(x)$ modulo p and q
$h(x)$	Public key polynomial
$sk = (f(x), f_p(x))$	Private key of an institution
$r_j(x)$	Random polynomial sampled by institution P_j
$m_j(x)$	Plaintext polynomial encoding model update of institution P_j

(Continued)

Table 3 (continued)

Symbol	Description
$e_j(x)$	Ciphertext sent by institution P_j
$H(\cdot)$	Secure cryptographic hash function (e.g., SHA-256), assumed to be collision-resistant and preimage-resistant
$\alpha_j = H(\mathrm{id}_j)$	Authentication tag for institution P_j
S	Set of authenticated institutions in a given round
$E(x)$	Aggregated ciphertext computed by the aggregator
$a(x)$, $b(x)$	Intermediate polynomials used in NTRU decryption
T	Maximum number of communication rounds

These notations will be used consistently in the subsequent subsections, including the encryption scheme, algorithmic flow, and complexity analysis.

3.1 System Flow and Architecture

The proposed system adopts a horizontal federated learning paradigm in which multiple medical institutions collaborate to train a global model while keeping all patient data strictly local. Fig. 1 illustrates the overall architecture of the proposed framework, showing how the trusted center (TC), the aggregator server, and the participating institutions interact through secure communication channels. In this architecture, each institution independently trains a local model, encrypts its model update using the NTRU public key, and transmits only encrypted updates to the aggregator. The aggregator performs ciphertext-domain aggregation and broadcasts the encrypted global update back to all institutions, ensuring that no plaintext model parameters are ever exposed at the server side.

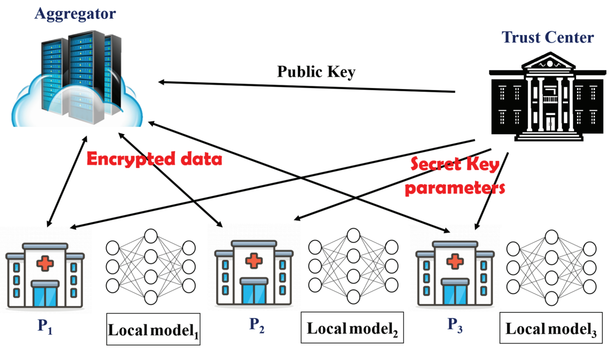

Figure 1: Overview of the proposed NTRU-enhanced federated learning system. The trusted center distributes key materials, institutions train local models, and the aggregator performs encrypted aggregation.

During the system initialization phase, the TC selects the NTRU parameters (N, p, q) and generates the corresponding key pairs. Each institution receives its private key, while the public key and the institution identifiers are shared with the aggregator. This setup establishes a secure authentication mechanism ensuring that only legitimate institutions may participate in the training process.

In the model training stage, every participating institution performs local training on its private dataset. The dataset never leaves the institution's environment. After the local model is trained, the model update is encoded into a polynomial representation compatible with the NTRU encryption scheme.

Before transmitting updates, each institution encrypts its encoded model update using the NTRU public key and attaches an authentication tag generated by a cryptographic hash function. The aggregator server receives these encrypted updates, verifies the authentication tags, and performs aggregation directly in the ciphertext domain using the homomorphic properties of NTRU. The aggregated ciphertext is then broadcast to all institutions. Each institution decrypts the ciphertext using its private key to recover the global model update and continues the next round of training until convergence.

3.2 NTRU-Based Encryption Scheme for the Federated Learning System

To guarantee post-quantum confidentiality and support encrypted aggregation across distributed medical institutions, this system incorporates an NTRU-based public-key cryptosystem. The NTRU operations used in the system, including key generation, encryption, authentication, aggregation, and decryption, correspond to Eqs. (2)–(13), which formally define the underlying mathematical mechanisms.

Key Generation:

Key generation is performed by the trusted center (TC). The TC selects the NTRU parameters (N, p, q), where N denotes the degree of the truncated polynomial ring, q is the large modulus used for ciphertexts, and p is a smaller modulus for plaintext model coefficients. NTRU operates in the polynomial ring $R_q = \mathbb{Z}_q\left[x\right]/(x^N - 1)$, where polynomial multiplication is implemented through cyclic convolution.

The TC selects two small polynomials $f(x)$ and $g(x)$, ensuring that $f(x)$ is invertible modulo both p and q. Their inverses $f_p(x)$ and $f_q(x)$ are computed according to Eqs. (4) and (5). The public key is obtained using Eq. (6), while the private key $(f(x), f_p(x))$ is securely distributed to each institution. The TC also defines a secure cryptographic hash function $H(\cdot)$ (e.g., SHA-256), which is assumed to be collision-resistant and preimage-resistant, for authentication.

$$f(x) = \sum_{i=0}^{N-1} f_i x^i \tag{2}$$

$$g(x) = \sum_{i=0}^{N-1} g_i x^i \tag{3}$$

$$f_p(x) = f(x)^{-1}\ (modp) \tag{4}$$

$$f_q(x) = f(x)^{-1}\ (modq) \tag{5}$$

$$h\left(x\right) = f_q\left(x\right) \star g\left(x\right)\left(modq\right) \tag{6}$$

$$sk = \left(f(x), f_p(x)\right) \tag{7}$$

Encryption:

After an institution completes its local training for a communication round, the model update is encoded into a polynomial $m_j(x)$ with coefficients bounded by $\left[-\frac{p}{2}, \frac{p}{2}\right]$, ensuring that plaintext coefficients remain small relative to q. A small random polynomial $r_j(x)$ is sampled, and the ciphertext is computed using the NTRU encryption rule, which introduces a scaling factor p to ensure correct decryption:

$$e_j\left(x\right) = pr_j\left(x\right) \star h\left(x\right) + m_j\left(x\right)\left(mod\ q\right) \tag{8}$$

An authentication tag is generated for each institution:

$$\alpha_j = H\left(\mathrm{id}_j\right) \tag{9}$$

Aggregation:

Upon receiving authenticated ciphertexts from all participating institutions, the aggregator verifies each tag, discards unauthenticated updates, and aggregates the remaining ciphertexts using polynomial addition:

$$E(x) = \sum_{j=1}^{n} e_j(x)\,(mod\,q) \tag{10}$$

The aggregated ciphertext $E(x)$ is then broadcast back to all participating institutions.

Decryption:

Upon receiving the aggregated ciphertext $E(x)$, each institution uses its private key to apply the NTRU decryption process. The decryption steps correspond to Eqs. (11)–(13), which include computing an intermediate polynomial, reducing coefficients modulo p, and applying the modular inverse $f_p(x)$ to recover the plaintext aggregated update. This decrypted update is then incorporated into the local model.

Upon receiving the aggregated ciphertext, each institution uses its private key to perform the NTRU decryption procedure. The steps correspond to Eqs. (11)–(13), which involve computing an intermediate polynomial, performing coefficient reduction modulo p, and applying the modular inverse $f_p(x)$ to recover the aggregated plaintext update.

$$a(x) = f\,(x) \star E\,(x)\,(modq) \tag{11}$$

$$b\,(x) = a\,(x)\,(modp) \tag{12}$$

$$m\,(x) = f_p(x) \star b(x)\,(modp) \tag{13}$$

This plaintext result $m(x)$ is then incorporated into the local model update for the next federated learning round.

Lemma 1: *SVP-Based Hardness of Plaintext Recovery in the NTRU Scheme.*

Given a ciphertext polynomial $e_j\,(x) = pr_j\,(x) \star h\,(x) + m_j\,(x)\,(modq)$, recovering the plaintext polynomial $m_j\,(x)$ without knowledge of the private key $f\,(x)$ can be reduced to solving a shortest vector problem (SVP) in the corresponding NTRU lattice:

Proof: The public key $h\,(x)$ defines an NTRU lattice constructed from the polynomial ring $\mathbb{Z}_q\,[x]\,/(x^N - 1)$. Any adversary attempting to recover $m_j\,(x)$ from $e_j\,(x)$ must distinguish the term $pr_j\,(x) \star h\,(x)$ from noise without access to the secret polynomial $f\,(x)$. This task can be reduced to finding a short vector in the NTRU lattice, which is known to be computationally hard under both classical and quantum adversarial models. Therefore, plaintext recovery without the private key implies an efficient solution to the SVP, contradicting its assumed hardness. □

Theorem 1: *IND-CPA Security of the Proposed NTRU-Based Encryption Scheme.*

The proposed NTRU-based encryption scheme used in the federated learning system is indistinguishable under chosen-plaintext attacks (IND-CPA secure), assuming the hardness of the shortest vector problem (SVP) in NTRU lattices.

Proof: The encryption of model updates incorporates a randomized polynomial $r_j\,(x)$, which ensures that identical plaintexts produce computationally indistinguishable ciphertexts. Under the assumption that distinguishing valid ciphertexts from uniformly random elements in the polynomial ring is as hard as solving the SVP in the associated lattice, no probabilistic polynomial-time adversary can distinguish encryptions of chosen plaintexts with non-negligible advantage. Consequently, the proposed scheme satisfies IND-CPA security under the standard NTRU hardness assumptions. □

Remark:

The above lemma and theorem establish that the confidentiality of encrypted model updates in the proposed federated learning framework is guaranteed by the well-studied hardness of lattice problems, providing resistance against both classical and quantum attacks.

These theoretical guarantees form the basis of the security analysis presented in Section 4.

3.3 NTRU-Enhanced Federated Learning Workflow

This subsection explains how the NTRU cryptographic operations described in Section 3.2 integrate into the federated learning process executed across multiple medical institutions. The workflow consists of four major phases: local training, NTRU-based encryption, encrypted aggregation at the aggregator, and collaborative decryption.

During each communication round, every institution trains a local model using its private dataset. After training, the model update is encoded into polynomial form and encrypted using the NTRU public key. A hash-based authentication tag ensures that only legitimate institutions can submit model updates.

The aggregator receives the encrypted updates, verifies their authentication tags, and aggregates the ciphertexts using the additive homomorphic property of NTRU. Because aggregation is performed entirely in the ciphertext domain, the aggregator never observes any plaintext values. The aggregated ciphertext is then broadcast back to all participating institutions.

Each institution decrypts the aggregated ciphertext using its private key to recover the global model update. This guarantees that all institutions obtain the same global parameters while maintaining full privacy of intermediate local updates. The training process continues for multiple rounds until convergence or until a predefined stopping criterion is reached.

The complete operational flow, which includes initialization, encryption, authentication, ciphertext aggregation, and joint decryption, is summarized later in Algorithm 1.

3.4 Algorithm Description

Algorithm 1 summarizes the complete operational workflow of the proposed NTRU-enhanced federated learning system. The algorithm integrates the cryptographic computations defined in Section 3.2 with the federated learning workflow described in Section 3.3. It consists of four key phases: (i) system initialization, (ii) local training and NTRU-based encryption, (iii) ciphertext authentication and homomorphic aggregation at the aggregator, and (iv) collaborative decryption across institutions.

During initialization, the trusted center generates the NTRU public and private keys and distributes the required parameters to all participating institutions. In each communication round, institutions train local models, encode their updates into polynomial form, and encrypt them using the NTRU public key. Authentication tags ensure that only legitimate updates are incorporated into the aggregation process.

The aggregator verifies all received ciphertexts, discards invalid ones, and performs homomorphic addition to compute the aggregated ciphertext without accessing any plaintext information. The aggregated ciphertext is then broadcast back to the participating institutions. Each institution decrypts the ciphertext using its private key to recover the aggregated model update, ensuring consistency of the global model across all participants.

The complete procedure is formally presented in Algorithm 1.

Algorithm 1: NTRU-enhanced federated learning across multiple medical institutions

```
Input:
    •Local datasets $\{D_j\}_{j=1}^{n}$
    •NTRU parameters $(N, p, q)$
    •Maximum communication rounds T
Output:
    •Global model parameters $w^{(T)}$

System Initialization (Trusted Center)
    Select NTRU parameters $(N, p, q)$.
    Sample small polynomials $f(x)$ and $g(x)$.
    Compute inverses:
        $f_p(x) = f(x)^{-1}\ (mod p)$,
        $f_q(x) = f(x)^{-1}\ (mod q)$.
    Compute public key:
        $h(x) = f_q(x) \star g(x)\ (mod q)$.
    Distribute private key $(f(x), f_p(x))$
        and publish $h(x)$ and hash $H(\cdot)$.

For each communication round $t = 0, \dots, T-1$ do

Phase 1: Local Training (each institution $P_j$)
    Train local model on $D_j \rightarrow w_j^t$.
    Encode update into polynomial $m_j(x)$.
    Sample random polynomial $r_j(x)$.
    Encrypt:
        $e_j(x) = pr_j(x) \star h(x) + m_j(x)\ (mod q)$.
    Compute authentication tag $\alpha_j = H(\mathrm{id}_j)$.
    Send $(e_j(x), \alpha_j)$ to aggregator.
Phase 2: Authentication and Homomorphic Aggregation
    For each $(e_j(x), \alpha_j)$,
        verify $\alpha_j = H(\mathrm{id}_j)$.
    Discard unauthenticated updates.
    Aggregate:
        $E(x) = \sum_{j \in S} e_j(x)\ (mod q)$.
    Broadcast $E(x)$ to institutions in $S$.
Phase 3: Collaborative Decryption (each $P_j \in S$)
    Compute $a(x) = f(x) \star E(x)\ (mod q)$.
    Compute $b(x) = a(x)\ (mod p)$.
    Recover plaintext:
        $m(x) = f_p(x) \star b(x)\ (mod p)$.
    Update local model using $m(x)$.
Stopping Condition
    If convergence or stopping criterion is met: break;
        otherwise continue.
```

3.5 Complexity Analysis

The computational and communication complexity of the proposed NTRU-enhanced federated learning system is analyzed in this section. Because all cryptographic operations operate on truncated polynomials of degree N, the dominant computational cost arises from polynomial convolution, which can be efficiently implemented using FFT-based techniques with a complexity of $O(N\log N)$. This property ensures that the cryptographic mechanisms scale well with the model size. A detailed summary of these complexity results is provided in Table 4 for clarity and ease of reference.

Table 4: Complexity summary of the NTRU-enhanced FL system.

Operation	Complexity	Description
Key Generation	$O(N\log N)$	Polynomial sampling + modular inversion
Encryption (Per Client)	$O(N\log N)$	Dominated by convolution $\boldsymbol{r_j}\,(\boldsymbol{x}) \star \boldsymbol{h}\,(\boldsymbol{x})$
Aggregation (Server)	$O(\lvert S\rvert N)$	Elementwise ciphertext addition over all authenticated clients
Decryption (Per Client)	$O(N\log N)$	Two convolutions + coefficient reduction
Communication Per Client Per Round	$O(N\log q)$	Upload + download of ciphertexts

During key generation, the trusted center samples two small polynomials and computes their modular inverses with respect to p and q. Polynomial sampling incurs a cost of $O(N)$, while polynomial inversion using the extended Euclidean algorithm requires $O(N\log N)$. Hence, the overall complexity of key generation is $O(N\log N)$. For encryption, each participating institution samples a random polynomial, computes the convolution $r_j\,(x) \star h\,(x)$, and adds the plaintext polynomial. As convolution dominates the computation, encryption incurs a total cost of $O(N\log N)$ per client.

On the server side, authenticated ciphertexts are aggregated through elementwise addition to obtain the final aggregated ciphertext $E(x)$. This operation does not involve convolution and scales linearly with the number of authenticated clients $\lvert S\rvert$, resulting in a complexity of $O(\lvert S\rvert N)$. For decryption, each institution computes the convolution $f(x) \star E(x)$, performs coefficient reduction modulo p, and subsequently applies the convolution $f_p(x) \star b(x)$. These steps yield a total decryption cost of $O(N\log N)$ per client.

The communication complexity is primarily determined by the exchange of ciphertext polynomials. Each institution uploads a ciphertext $e_j\,(x)$ of size $O(N\log q)$ bits and receives a ciphertext $E(x)$ of similar size from the aggregator. The authentication tag α_j contributes negligible overhead compared to the ciphertext size. Therefore, the communication cost per client per round is $O(N\log q)$, which grows linearly with the model dimension and remains practical for federated learning deployments.

To summarize, the proposed system maintains computational efficiency comparable to state-of-the-art lattice-based cryptographic schemes while offering secure aggregation and scalability with respect to both model dimension and the number of participating institutions.

3.6 Quantum Security Comparison

The emergence of large-scale quantum computers poses a fundamental threat to classical public-key cryptosystems based on integer factorization and discrete logarithm problems. As a result, it is essential to assess the security of the proposed federated learning framework within the broader landscape of post-quantum cryptography. The proposed scheme adopts NTRU-based encryption, whose security relies on the hardness of lattice problems that are widely believed to be resistant to both classical and quantum attacks.

Rather than competing with standardized post-quantum primitives such as CRYSTALS-Kyber or CRYSTALS-Dilithium in terms of key encapsulation or digital signatures, the proposed NTRU-enhanced federated learning scheme is designed to provide secure aggregation functionality under post-quantum assumptions. Its role is therefore complementary to existing post-quantum standards. For completeness, Table 5 compares the quantum security properties of the proposed scheme with representative post-quantum and classical cryptosystems, highlighting their underlying hardness assumptions, parameter settings, and cryptographic roles.

Table 5: Quantum security comparison of the proposed scheme.

Scheme	Underlying Hard Problem	Lattice Dimension /Parameter	Security Level (Approx.)	Cryptographic Role
Proposed NTRU-Enhanced FL	NTRU-SVP	N (e.g., 701)	≥128-bit PQ security	Secure aggregation in FL
Standard NTRU	NTRU-SVP	$N = 701$	≥128-bit PQ security	Public-key encryption
CRYSTALS-Kyber-768	Module-LWE	$k = 3$	NIST Level 3	Key encapsulation
CRYSTALS-Dilithium-3	Module-LWE/SIS	$k = 4$	NIST Level 3	Digital signature
RSA-2048	Integer factorization	2048-bit modulus	Not PQ-secure	Legacy encryption
ECC-P256	Elliptic curve DLP	256-bit curve	Not PQ-secure	Legacy public key

4 Security Analysis

This section presents a comprehensive security analysis of the proposed NTRU-enhanced federated learning framework. Building upon the formal cryptographic foundations established in Section 3, including the system model, Algorithm 1, Lemma 1 (SVP-based hardness of plaintext recovery), and Theorem 1 (IND-CPA security), we analyze the resilience of the proposed framework against representative threat models in federated learning. In particular, we consider inference attacks, model poisoning attacks, collusion involving an honest-but-curious aggregator, and quantum adversaries.

4.1 Resistance to Inference Attacks

Inference attacks aim to extract sensitive information about local training data by observing transmitted model updates or aggregated results. In conventional federated learning systems, unencrypted gradients or model parameters may leak substantial information, enabling gradient inversion or membership inference attacks.

In the proposed framework, local model updates are never transmitted in plaintext. Instead, each medical institution encodes its local update into polynomial form and encrypts it using the NTRU public key before transmission, as described in Section 3.2 and Algorithm 1. The aggregator receives only encrypted model updates and performs aggregation exclusively in the ciphertext domain.

By Lemma 1, recovering the plaintext polynomial from a ciphertext without knowledge of the private key can be reduced to solving a shortest vector problem (SVP) in the corresponding NTRU lattice, which is widely believed to be computationally intractable. Furthermore, Theorem 1 establishes that the encryption scheme satisfies IND-CPA security, ensuring that ciphertexts corresponding to different plaintext updates are computationally indistinguishable. As a result, an adversary observing encrypted updates or aggregated ciphertexts cannot infer individual gradients, model parameters, or sensitive training data with non-negligible probability.

Proposition 1: *Inference Resistance under Encrypted Aggregation.*

Under the proposed NTRU-enhanced federated learning framework, an adversary that observes encrypted local updates and aggregated ciphertexts cannot infer any participant's local training data with non-negligible probability, assuming the hardness of the NTRU-SVP.

4.2 Robustness against Model Poisoning Attacks

Model poisoning attacks seek to manipulate the global model by injecting malicious updates during the aggregation process. In the proposed framework, each encrypted update is accompanied by a hash-based authentication tag that binds the ciphertext to a registered institution identity. The aggregator verifies these authentication tags prior to aggregation and discards any unauthenticated or malformed submissions.

Although the aggregator does not decrypt individual updates, the authentication mechanism prevents external adversaries from injecting arbitrary ciphertexts into the aggregation process. Moreover, since aggregation is performed entirely in the ciphertext domain, attackers cannot adaptively craft poisoned updates based on observed plaintext information.

It is noted that the proposed framework primarily addresses security at the cryptographic and protocol levels. While it does not explicitly detect semantic poisoning behavior within legitimate updates, it provides a secure foundation upon which complementary robustness techniques, such as anomaly detection or robust aggregation rules, can be integrated without compromising data confidentiality.

4.3 Collusion and Honest-but-Curious Aggregator

A common threat model in federated learning assumes an honest-but-curious aggregator that faithfully follows the protocol while attempting to learn private information from received messages. In the proposed system, the aggregator observes only encrypted model updates and performs aggregation solely in the ciphertext domain. At no point does it gain access to plaintext model parameters.

Even in scenarios involving collusion between the aggregator and a subset of participating institutions, confidentiality of honest participants is preserved. Without access to the private key components, colluding adversaries cannot decrypt individual ciphertexts or the aggregated ciphertext. As established by Lemma 1 and Theorem 1, any attempt to recover plaintext information from ciphertexts reduces to solving a hard lattice problem.

Therefore, the proposed framework significantly reduces trust assumptions placed on the aggregator and provides strong protection against collusion-based privacy leakage.

4.4 Security against Quantum Adversaries

The advent of quantum computing poses severe threats to classical public-key cryptosystems based on integer factorization and discrete logarithm problems. In contrast, the proposed framework relies on the hardness of lattice-based problems, specifically the NTRU shortest vector problem, which are widely believed to remain secure against both classical and quantum adversaries.

Shor's algorithm, which efficiently breaks RSA and elliptic curve cryptography, does not apply to lattice-based constructions. While Grover's algorithm offers a quadratic speedup for brute-force search, its impact on lattice-based schemes can be mitigated through appropriate parameter selection. As summarized in Table 5, the proposed NTRU-enhanced federated learning scheme achieves post-quantum security levels comparable to standardized lattice-based cryptosystems such as CRYSTALS-Kyber and CRYSTALS-Dilithium.

These properties make the proposed framework particularly suitable for long-term deployment in security-critical domains, including collaborative medical data analysis, where resistance to future quantum attacks is essential.

4.5 Summary

In summary, the proposed NTRU-enhanced federated learning framework provides strong security guarantees across multiple threat models. Formal cryptographic results established in Section 3 ensure confidentiality under standard lattice hardness assumptions, while the system design mitigates inference attacks, limits the impact of poisoning attempts, reduces trust in the aggregator, and offers resilience against quantum adversaries. Together, these properties demonstrate that the proposed framework achieves a high level of security and robustness suitable for privacy-sensitive, large-scale federated learning applications.

5 Experimental Evaluation

5.1 Dataset Description

We evaluate the proposed NTRU-enhanced federated learning framework using the Wisconsin Diagnostic Breast Cancer (WDBC) dataset, a widely adopted benchmark in medical diagnosis and privacy-preserving learning research. The dataset was originally collected from Fine Needle Aspiration (FNA) examinations of breast masses and is publicly available through the UCI Machine Learning Repository. It can also be accessed via the breast cancer dataset interface in the scikit-learn library.

The WDBC dataset consists of 569 patient-level samples, each represented by a 30-dimensional numerical feature vector derived from digitized FNA images. These features are computed from ten fundamental cell nucleus characteristics, including radius, texture, perimeter, area, smoothness, compactness, concavity, concave points, symmetry, and fractal dimension. For each characteristic, three statistical measurements, namely mean, standard error, and worst value, are provided. The classification task is binary, aiming to distinguish malignant tumors (label 0) from benign tumors (label 1). The dataset contains 357 benign samples (62.8%) and 212 malignant samples (37.2%), reflecting a mildly imbalanced but realistic clinical data distribution. No missing values are present, making the dataset well suited for controlled evaluation of federated learning and cryptographic mechanisms.

5.2 Cryptographic Parameter Settings

To ensure post-quantum security while maintaining practical efficiency, the proposed framework adopts the NTRU-HPS-2048-401 parameter set, which conforms to NIST Post-Quantum Cryptography Level 1 security, corresponding to approximately 128-bit symmetric security strength. The adopted NTRU parameters are summarized in Table 6.

Table 6: Hyper parameters of NTRU.

N	p	q	d_f	d_g	d
401	3	2048	113	113	60

In this configuration, the polynomial ring dimension is set to $N = 401$, with the small modulus $p = 3$ and the large modulus $q = 2048$. The parameters d_f and dg specify the sparsity of the secret and public key polynomials, respectively, by controlling the number of non-zero coefficients. Higher sparsity improves computational efficiency during polynomial convolution, while maintaining sufficient entropy to resist lattice-based attacks. The parameter d determines the sparsity of the random masking polynomial used during encryption, further enhancing semantic security.

This parameter selection provides a balanced trade-off between cryptographic robustness and computational feasibility, making it well suited for medical federated learning scenarios where both security and efficiency are critical.

5.3 *Model Performance Evaluation*

Table 7 presents the final test performance of Plain-FL and the proposed NTRU-FL framework under different numbers of participating users (3, 5, 10, and 20 users). Experimental results show that when the number of users is 5 or 10, both achieve an accuracy of 0.956 and an F1 score of 0.9650. Even with different user sizes, the accuracy remains stable between 0.9473 and 0.9474, while the F1 score ranges from 0.9589 to 0.9577. This observation demonstrates that integrating the NTRU key exchange mechanism does not cause any perceptible damage to model performance. Compared to Plain-FL, NTRU-FL achieves post-quantum security while precisely preserving the training efficiency of standard federated learning, with zero performance difference between the two.

Table 7: Final testing metric of plain-FL and NTRU-FL.

Client	Plain-FL				NTRU-FL			
	3 Clients	5 Clients	10 Clients	20 Clients	3 Clients	5 Clients	10 Clients	20 Clients
Accuracy	0.9473	0.9561	0.9561	0.9474	0.9473	0.9561	0.9560	0.9474
F1-score	0.9589	0.9650	0.9650	0.9577	0.9589	0.9650	0.9650	0.9577

5.4 *Computational and Communication Overhead*

Table 8 summarizes the average cryptographic computation and communication overhead of Plain-FL and NTRU-FL after 30 rounds of federated learning at different user scales. Plain-FL has zero cryptographic computation latency because it does not perform encryption or decryption operations during model update transmission. In contrast, NTRU-FL introduces additional encryption and decryption computations to ensure the security of model updates.

Table 8: Cumulative crypto-communication overhead over 30 rounds across different client counts.

30 Rounds	Plain-FL				NTRU-FL			
	3 Clients	5 Clients	10 Clients	20 Clients	3 Clients	5 Clients	10 Clients	20 Clients
Avg Enc. Latency (ms)	–	–	–	–	71,192.7	116,903.4	235,089.3	353,247.5

(Continued)

Table 8 (continued)

30 Rounds	Plain-FL				NTRU-FL			
	3 Clients	5 Clients	10 Clients	20 Clients	3 Clients	5 Clients	10 Clients	20 Clients
Avg Dec. Latency (ms)	–	–	–	–	135,120.7	210,747.9	408,053.6	599,768.2
Avg Total Latency (ms)	–	–	–	–	206,313.4	327,651.3	643,142.9	953,015.7
Avg Transmission Size (KB)	10.1182	10.1182	10.1182	10.1182	215.2529	214.2376	215.2715	215.4152

Regarding computational efficiency, as the number of users increases from 3 to 20, the average total encryption latency increases from 71,192.7 ms to 353,247.5 ms; simultaneously, the average total decryption latency shows a similar upward trend, rising from 135,120.7 ms to 599,768.2 ms. This increase reflects the cumulative cryptographic cost across all participating clients as the number of encrypted model updates grows with the scale of the federation, rather than an increase in per-client encryption or decryption time, thereby demonstrating the scalability of the proposed framework.

Regarding communication overhead, the average transmission size of NTRU-FL is approximately 214.23 KB to 215.41 KB per user, which is significantly higher than Plain-FL's 10.1182 KB. However, this transmission size remains stable across different user settings, indicating that the communication cost is primarily determined by the fixed NTRU parameters rather than by the number of participants. Overall, the experimental results show that the proposed NTRU-enhanced federated learning framework introduces acceptable computational overhead while providing robust post-quantum security. Although encryption and decryption operations introduce latency compared to Plain-FL, these costs remain manageable in non-real-time offline or asynchronous training scenarios such as medical settings. It should be noted that the encryption and decryption latency listed in the table represents the cumulative cost over 30 training rounds; when amortized to a single round, the latency remains within a practical range for real-world deployment.

5.5 Scalability and Latency Variability Analysis

To further analyze client-side heterogeneity, Table 9 reports the variability of per-client cryptographic latency in the proposed NTRU-FL framework, measured in terms of maximum and minimum encryption and decryption latency across clients. The results show that both encryption and decryption latency exhibit increasing variability as the number of participating clients grows. For instance, under the 20-client setting, encryption latency ranges from 341.0 to 735.8 ms, while decryption latency ranges from 591.9 to 1350.7 ms.

Table 9: Variability of per-client cryptographic latency.

	NTRU-FL			
	3 Clients	**5 Clients**	**10 Clients**	**20 Clients**
Enc. Latency (max/min ms)	73,998.7/67,580.1	121,351.2/113,690.2	259,182.8/226,243.9	735,793.7/341,035.6
Dec. Latency (max/min ms)	141,687.4/126,154.0	223,039.1/198,242.3	429,305.1/374,975.7	1350,722.0/591,928.4

This variability is primarily attributed to differences in client-side computational resources and the asynchronous execution of cryptographic operations. Despite these variations, the latency values remain within the same order of magnitude for each client configuration, indicating stable and predictable cryptographic behavior. These findings demonstrate that the proposed NTRU-FL framework maintains robust performance across heterogeneous clients, which is essential for practical deployment in real-world medical federated learning environments.

5.6 *Convergence Behavior*

Fig. 2 illustrates the convergence behavior of the proposed NTRU-FL framework under different numbers of participating clients. Fig. 2a presents the testing loss curves, while Fig. 2b shows the corresponding classification accuracy over federated learning rounds. As shown in Fig. 2a, the testing loss decreases steadily across all client settings, indicating stable convergence under encrypted model updates.

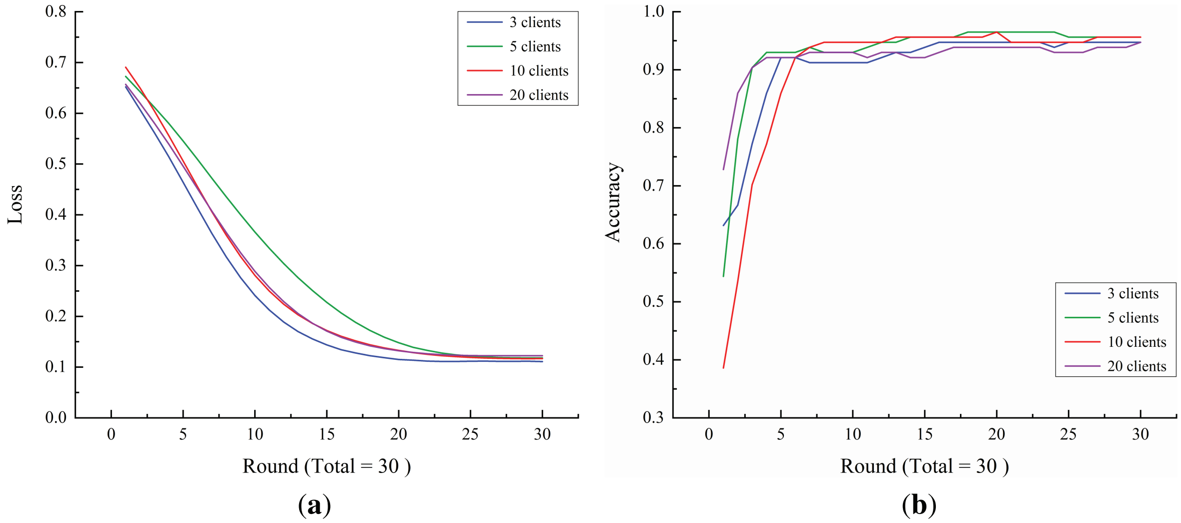

Figure 2: Testing performance of the proposed NTRU-FL framework under different numbers of participating clients: (**a**) testing loss and (**b**) classification accuracy. The results show stable convergence behavior across all client configurations, indicating that the integration of NTRU-based encrypted aggregation preserves the convergence behavior of federated learning.

Fig. 2b demonstrates that classification accuracy increases consistently and converges to a similar level regardless of the number of participating clients. These results confirm that the introduction of NTRU-based encryption does not impede convergence speed or final predictive performance, further validating the practicality of the proposed framework.

6 Conclusion

This study addresses critical security and privacy challenges in healthcare-oriented federated learning by integrating NTRU-based encryption into the collaborative training process. Medical data are inherently sensitive and distributed across multiple institutions, making traditional centralized learning paradigms vulnerable to privacy leakage and regulatory risks. By leveraging lattice-based cryptography, particularly the hardness of the Shortest Vector Problem (SVP), the proposed framework provides strong protection against both classical and quantum adversaries, offering a forward-looking solution for secure and privacy-preserving medical data analytics.

From a post-quantum security perspective, the adoption of NTRU encryption effectively mitigates threats posed by quantum algorithms such as Shor's and Grover's algorithms, thereby ensuring long-term security under increasingly powerful computational capabilities. In addition, the integration of randomized polynomials and encrypted aggregation mechanisms prevents the disclosure of local model updates, strengthening resistance against inference attacks, model inversion attacks, data and model poisoning, brute-force attacks, and man-in-the-middle attacks throughout the federated learning process.

From a regulatory and ethical standpoint, the proposed federated learning framework is designed to align with major healthcare data protection principles, including HIPAA and GDPR. No raw patient data or intermediate gradients are exchanged between participating institutions, and all shared information is restricted to cryptographically protected model updates. This design substantially reduces data exposure risks while enabling privacy-preserving collaboration among multiple medical entities.

Overall, the proposed framework demonstrates practical applicability across a wide range of healthcare scenarios, including clinical diagnosis, drug discovery, personalized treatment, and cross-regional public health monitoring. By enabling secure, compliant, and efficient collaboration without exposing sensitive medical data, this work lays a solid foundation for the deployment of advanced AI-driven healthcare applications and contributes to the development of trustworthy and intelligent medical systems.

Acknowledgement: Not applicable.

Funding Statement: The authors received no specific funding for this study.

Author Contributions: Conceptualization: Chia-Hui Liu; Methodology: Chia-Hui Liu; Formal analysis: Chia-Hui Liu; Software: Chia-Hui Liu; Validation: Chia-Hui Liu and Zhen-Yu Wu; Investigation: Chia-Hui Liu; Writing—Original Draft: Chia-Hui Liu; Writing—Review & Editing: Zhen-Yu Wu; Visualization: Chia-Hui Liu; Supervision: Zhen-Yu Wu; Project administration: Zhen-Yu Wu. The authors solely completed the conception, design, analysis, and writing of this manuscript. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The Wisconsin Diagnostic Breast Cancer (WDBC) dataset used in this study is publicly available through the UCI Machine Learning Repository and is also accessible via the scikit-learn library. No new data were collected for this study.

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Chen L, Jordan S, Liu YK, Moody D, Peralta R, Perlner R, et al. NIST IR, 8105 report on post-quantum cryptography. Gaithersburg, MD, USA: National Institute of Standards and Technology; 2016. p. 15. doi:10.6028/NIST.IR.8105.
2. Acar A, Aksu H, Uluagac AS, Conti M. A survey on homomorphic encryption schemes: theory and implementation. ACM Comput Surv. 2019;51(4):1–35. doi:10.1145/3214303.
3. Paillier P. Public-key cryptosystems based on composite degree residuosity classes. In: Advances in cryptology—EUROCRYPT'99. Berlin/Heidelberg, Germany: Springer; 1999. p. 223–38. doi:10.1007/3-540-48910-X_16.
4. Hoffstein J, Pipher J, Silverman JH. NTRU: a ring-based public key cryptosystem. In: Algorithmic number theory. Berlin/Heidelberg, Germany: Springer; 1998. p. 267–88. doi:10.1007/bfb0054868.
5. Qayyum A, Ahmad K, Ahsan MA, Al-Fuqaha A, Qadir J. Collaborative federated learning for healthcare: multi-modal COVID-19 diagnosis at the edge. IEEE Open J Comput Soc. 2022;3(12):172–84. doi:10.1109/OJCS.2022.3206407.
6. Nguyen DC, Pham QV, Pathirana PN, Ding M, Seneviratne A, Lin Z, et al. Federated learning for smart healthcare: a survey. ACM Comput Surv. 2023;55(3):1–37. doi:10.1145/3501296.
7. Roth HR, Cheng Y, Wen Y, Yang I, Xu Z, Hsieh YT, et al. NVIDIA FLARE: federated learning from simulation to real-world. arXiv: 2210.13291. 2022.
8. Chen Y, Qin X, Wang J, Yu C, Gao W. FedHealth: a federated transfer learning framework for wearable healthcare. IEEE Intell Syst. 2020;35(4):83–93. doi:10.1109/MIS.2020.2988604.
9. Lim WYB, Luong NC, Hoang DT, Jiao Y, Liang YC, Yang Q, et al. Federated learning in mobile edge networks: a comprehensive survey. IEEE Commun Surv Tutor. 2020;22(3):2031–63. doi:10.1109/COMST.2020.2986024.

10. Sheller MJ, Edwards B, Reina GA, Martin J, Pati S, Kotrotsou A, et al. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. Sci Rep. 2020;10(1):12598. doi:10.1038/s41598-020-69250-1.
11. Li Q, Wen Z, Wu Z, Hu S, Wang N, Li Y, et al. A survey on federated learning systems: vision, hype and reality for data privacy and protection. IEEE Trans Knowl Data Eng. 2023;35(4):3347–66. doi:10.1109/TKDE.2021.3124599.
12. Rieke N, Hancox J, Li W, Milletarì 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.
13. Zhou J, Huang F, Wang S, Chen P. FedFIP: a personalized federated learning optimization method with differential privacy protection. In: Advanced intelligent computing technology and applications. Singapore: Springer Nature; 2025. p. 363–74. doi:10.1007/978-981-96-9872-1_30.
14. Zhu L, Liu Z, Han S. Deep leakage from gradients. In: Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS 2019); 2019 Dec 8–14; Vancouver,BC, Canada. p. 14747–56.
15. Kim M, Gunlu O, Schaefer RF. Federated learning with local differential privacy: trade-offs between privacy, utility, and communication. In: Proceedings of the ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP); 2021 Jun 6–11 Toronto, ON, Canada. p. 2650–4. doi:10.1109/icassp39728.2021.9413764.
16. Wei W, Liu L, Wu Y, Su G, Iyengar A. Gradient-leakage resilient federated learning. In: Proceedings of the 2021 IEEE 41st International Conference on Distributed Computing Systems (ICDCS); 2021 Jul 7–10; Washington, DC, USA. p. 797–807. doi:10.1109/icdcs51616.2021.00081.
17. Naseri M, Hayes J, De Cristofaro E. Toward robustness and privacy in federated learning: experimenting with local and central differential privacy. arXiv:2009.03561. 2020.
18. Li Y, Zhou Y, Jolfaei A, Yu D, Xu G, Zheng X. Privacy-preserving federated learning framework based on chained secure multiparty computing. IEEE Internet Things J. 2021;8(8):6178–86. doi:10.1109/JIOT.2020.3022911.
19. So J, Güler B, Avestimehr AS. Turbo-aggregate: breaking the quadratic aggregation barrier in secure federated learning. IEEE J Sel Areas Inf Theory. 2021;2(1):479–89. doi:10.1109/JSAIT.2021.3054610.
20. Xu G, Li H, Liu S, Yang K, Lin X. VerifyNet: secure and verifiable federated learning. IEEE Trans Inf Forensics Secur. 2020;15:911–26. doi:10.1109/TIFS.2019.2929409.
21. Wang F, Zhu H, Lu R, Zheng Y, Li H. A privacy-preserving and non-interactive federated learning scheme for regression training with gradient descent. Inf Sci. 2021;552:183–200. doi:10.1016/j.ins.2020.12.007.
22. Li T, Li J, Chen X, Liu Z, Lou W, Hou YT. NPMML: a framework for non-interactive privacy-preserving multi-party machine learning. IEEE Trans Dependable Secure Comput. 2021;18(6):2969–82. doi:10.1109/TDSC.2020.2971598.
23. Phong LT, Aono Y, Hayashi T, Wang L, Moriai S. Privacy-preserving deep learning via additively homomorphic encryption. IEEE Trans Inf Forensics Secur. 2018;13(5):1333–45. doi:10.1109/TIFS.2017.2787987.
24. Wang L, Polato M, Brighente A, Conti M, Zhang L, Xu L. PriVeriFL: privacy-preserving and aggregation-verifiable federated learning. IEEE Trans Serv Comput. 2025;18(2):998–1011. doi:10.1109/TSC.2024.3451183.
25. Chen Y, He S, Wang B, Feng Z, Zhu G, Tian Z. A verifiable privacy-preserving federated learning framework against collusion attacks. IEEE Trans Mob Comput. 2025;24(5):3918–34. doi:10.1109/TMC.2024.3516119.
26. Cai J, Shen W, Qin J. ESVFL: efficient and secure verifiable federated learning with privacy-preserving. Inf Fusion. 2024;109(1):102420. doi:10.1016/j.inffus.2024.102420.