

AdeptHEQ-FL: Adaptive Homomorphic Encryption for Federated Learning of Hybrid Classical-Quantum Models with Dynamic Layer Sparing

Md Abrar Jahin^{1*}, Taufikur Rahman Fuad², M. F. Mridha^{3*}, Nafiz Fahad⁴, Md. Jakir Hossen^{4*}

¹University of Southern California

²Islamic University of Technology

³American International University-Bangladesh

⁴Multimedia University

jahin@usc.edu, taufikur@iut-dhaka.edu,
firoz.mridha@aiub.edu, jakir.hossen@mmu.edu.my

Abstract

Federated Learning (FL) faces inherent challenges in balancing model performance, privacy preservation, and communication efficiency, especially in non-IID decentralized environments. Recent approaches either sacrifice formal privacy guarantees, incur high overheads, or overlook quantum-enhanced expressivity. We introduce AdeptHEQ-FL, a unified hybrid classical-quantum FL framework that integrates (i) a hybrid CNN-PQC architecture for expressive decentralized learning, (ii) an adaptive accuracy-weighted aggregation scheme leveraging differentially private validation accuracies, (iii) selective homomorphic encryption (HE) for secure aggregation of sensitive model layers, and (iv) dynamic layer-wise adaptive freezing to minimize communication overhead while preserving quantum adaptability. We establish formal privacy guarantees, provide convergence analysis, and conduct extensive experiments on the CIFAR-10, SVHN, and Fashion-MNIST datasets. AdeptHEQ-FL achieves a $\approx 25.43\%$ and $\approx 14.17\%$ accuracy improvement over Standard-FedQNN and FHE-FedQNN, respectively, on the CIFAR-10 dataset. Additionally, it reduces communication overhead by freezing less important layers, demonstrating the efficiency and practicality of our privacy-preserving, resource-aware design for FL. Our code is publicly available at: <https://github.com/Abrar2652/QML-FL>.

1. Introduction

Federated Learning (FL) has emerged as a transformative paradigm for collaborative Machine Learning (ML), allowing decentralized devices to train a shared model without centralizing sensitive data [20, 26]. This approach is crucial for privacy-sensitive applications, such as personalized medicine, secure finance, and the Internet of Things (IoT),

where data privacy and resource constraints are critical. However, effectively deploying FL is challenged by a triad of issues: statistical heterogeneity from non-Identical and non-Independently Distributed (non-IID) data, privacy vulnerabilities despite data localization, and high communication and computational overheads [20, 26]. Non-IID data among clients often hinders model performance and slows convergence. While FL inherently maintains some privacy by keeping data local, model updates remain vulnerable to attacks that can infer sensitive information. Frequent model exchanges between clients and servers amplify communication costs, especially in bandwidth-constrained environments. These interconnected issues require a unified solution to improve the robustness and scalability of FL.

Existing approaches often address these challenges in isolation, resulting in fragmented solutions. Quantum FL (QFL) utilizes quantum circuits to improve model expressivity [17, 24], but many frameworks overlook formal privacy guarantees or non-IID robustness [17]. Privacy-preserving techniques, such as Differential Privacy (DP) [38] and Homomorphic Encryption (HE) [41], protect data but compromise utility in non-IID settings or incur significant overhead [10]. Efficiency-focused methods, like model compression [13], reduce communication but seldom integrate quantum capabilities or comprehensively address privacy. This gap highlights the need for a unified framework that optimizes performance, privacy, and efficiency in a hybrid classical-quantum context.

To bridge these identified gaps, we propose **AdeptHEQ-FL**, a novel framework designed as a *unified* solution. Where existing QFL approaches often lack formal privacy or non-IID robustness, AdeptHEQ-FL synergistically combines its hybrid classical-quantum architecture with adaptive accuracy-weighted aggregation (utilizing differentially private validation accuracies) to explicitly improve perfor-

mance on non-IID data while improving model expressivity. To counter the significant overhead or utility degradation associated with many privacy-preserving techniques, especially in non-IID settings, AdeptHEQ-FL strategically employs HE (CKKS scheme) on critical final classical layers during aggregation, balancing strong privacy with computational feasibility, and further bolsters utility through its adaptive aggregation that prioritizes more accurate client models. Unlike efficiency-focused methods that typically neglect quantum capabilities or comprehensive privacy, our dynamic layer sparing mechanism is integrated to reduce communication and computation, specifically exempting quantum layers to preserve their crucial adaptability and ensuring the overall privacy-preserving nature of the framework is maintained. By holistically integrating these components, AdeptHEQ-FL provides a more comprehensive approach than existing fragmented solutions, aiming to concurrently optimize performance, privacy, and efficiency within a hybrid classical-quantum FL paradigm.

Our **primary contributions** are: (i) We introduce a novel *adaptive aggregation mechanism* for FL that employs differentially private client validation accuracies and HE to effectively address non-IID data and ensure privacy. (ii) We propose a *hybrid classical-quantum architecture* integrating CNNs for feature extraction with PQCs to improve model expressivity in federated settings. (iii) We develop an *efficient dynamic layer sparing technique* that reduces communication overhead by adaptively freezing less impactful classical layers while preserving the adaptability of quantum layers. (iv) We provide a theoretical convergence analysis for the proposed framework, accounting for adaptive aggregation, layer sparing, and privacy mechanisms.

2. Related Works

FL enables collaborative model training across decentralized devices while prioritizing data privacy, yet faces challenges from non-IID, privacy vulnerabilities, and high communication costs [20, 27]. Recent efforts explore quantum computing, privacy-preserving mechanisms, and efficiency optimizations, often addressing these issues in isolation. We critically review these efforts across four dimensions—quantum-improved FL, privacy preservation, communication efficiency, and adaptive/specialized approaches—identifying gaps that our AdeptHEQ-FL framework addresses through adaptive accuracy-weighted aggregation, classical-quantum hybridization, and formal convergence guarantees.

2.1. Quantum FL

Quantum FL (QFL) leverages quantum circuits to improve model expressivity. FedQNN [17] employs QNNs and discusses secure data handling, but lacks formal privacy mechanisms like DP, leaving potential vulnerabilities unaddressed. Similarly, [32, 38] integrate DP into QFL but overlook com-

munication costs and provide no convergence proofs, limiting their robustness. FHE-FedQNN [10] combines fully HE (FHE) with quantum circuits, reporting results on datasets like CIFAR-10 [23], Brain MRI [29], and PCOS [15]. Its uniform aggregation struggles with non-IID data, and FHE’s complexity leads to high communication overhead, a general concern in FHE-based FL approaches, rendering it impractical for edge devices. Its extension, MQFL-FHE, while leveraging hybrid quantum-FHE operations for multimodal tasks, still faces computational and communication inefficiencies. Theoretical studies like [6] explore Quantum Neural Networks (QNN) for FL without empirical validation, while [24] demonstrates QFL experimentally but omits formal convergence analysis. These works highlight QFL’s potential but fail to unify privacy, efficiency, and theoretical rigor, gaps AdeptHEQ-FL addresses.

2.2. Privacy-Preserving Techniques

Privacy in FL often relies on DP or HE. DP-based methods [8, 38] add noise to updates, degrading accuracy in non-IID settings [20]. HE-based aggregation [35, 41] ensures security but introduces significant computational overhead, limiting scalability. Hybrid approaches like ADPHE-FL [39] and others [1, 2, 42] adaptively combine DP and HE to balance privacy and utility in classical FL, yet neglect quantum improvements and communication efficiency for non-IID data. Comparative analyses [5] evaluate DP versus HE but offer no solutions for non-IID challenges, underscoring the need for AdeptHEQ-FL’s quantum-aware, adaptive privacy framework.

2.3. Efficiency in FL

Efficiency-focused FL methods aim to reduce communication and computational costs. FedSIGN [13] employs sign-based compression to lower bandwidth and provides convergence analysis, but lacks quantum compatibility, restricting its applicability to classical settings. Multi-party computation (MPC) approaches [7, 21, 37] reduce communication overhead, yet often compromise accuracy in non-IID settings and ignore quantum integration, as seen in their classical focus. These methods highlight a trade-off between efficiency and performance that AdeptHEQ-FL mitigates through adaptive layer freezing and quantum-improved aggregation.

2.4. Adaptive and Specialized Approaches

Adaptive FL frameworks like [34] explore functional encryption for security but do not address learning dynamics or non-IID convergence. Quantum-safe FL [41] applies HE without tackling non-IID data or providing guarantees. Quantum-inspired methods [4, 36] optimize computation via tensor networks or Quantum Key Distribution (QKD) but lack formal convergence guarantees for non-IID settings. Application-

specific FL, such as for mental healthcare [14], emphasizes both privacy and scalability, noting trade-offs such as longer training times for improved privacy. AdeptHEQ-FL distinguishes itself by integrating adaptive aggregation, quantum advantages, and rigorous convergence analysis for non-IID settings, offering a comprehensive solution.

3. Methodology

AdeptHEQ-FL is a novel FL framework that integrates classical and quantum neural networks to address challenges such as non-IID, privacy preservation, and communication efficiency. By combining performance-based adaptive aggregation, layer-wise adaptive freezing, and DP, AdeptHEQ-FL improves model performance, reduces communication overhead, and ensures client privacy while maintaining compatibility with HE. A high-level overview of the complete system of AdeptHEQ-FL is presented in Figure 1.

3.1. Problem Formulation

Consider an FL scenario with N clients, each possessing a local dataset $\mathcal{D}_i \sim p_i(x, y)$ that may exhibit non-IID distributions. The global model parameters $\theta = [\theta^c, \theta^q]$ consist of classical (θ^c) and quantum (θ^q) components. We define the learning objective as:

$$\min_{\theta} \sum_{i=1}^N \underbrace{\frac{|\mathcal{D}_i|}{\sum_{j=1}^N |\mathcal{D}_j|}}_{\text{fixed } w_i} \mathcal{L}_i(\theta; \mathcal{D}_i) \quad (1)$$

where the fixed aggregation weights w_i proportionally reflect each client's dataset size, ensuring that clients with larger datasets contribute more significantly to the global model. While this formulation establishes a stable baseline objective using static weights, the actual aggregation process (detailed in Section 3.4) employs dynamic weights $w_i^{(t)}$ derived from privatized validation accuracies to address non-IID challenges. The hybrid architecture simultaneously optimizes both quantum and classical parameters, maintaining regularization stability in classical components while allowing quantum layers to adapt freely to complex data patterns.

3.2. Model Architecture

The proposed AdeptHEQ-FL model integrates both classical and quantum computing techniques, leveraging the strengths of each. The following subsections discuss each of these components.

3.2.1. Classical Component

The classical component of the architecture is implemented as a CNN, which is a type of DL model that is particularly effective for analyzing grid-like data such as images. CNNs

function by applying multiple layers of convolutional filters that extract localized features from the input image, including edges, textures, and shapes.

In this implementation, the CNN is composed of three sequential convolutional blocks. Each block consists of multiple convolutional layers, followed by a Rectified Linear Unit (ReLU) activation function and a max-pooling layer. The convolutional layers perform a mathematical operation known as a discrete convolution:

$$s(i, j) = (I \times K)(i, j) = \sum_m \sum_n I(i + m, j + n) K(m, n) \quad (2)$$

where $I(i, j)$ represents the input image and $K(m, n)$ is a learnable filter (also called a kernel), this operation slides the kernel across the input image, producing a feature map that highlights the presence of specific patterns detected by the filter. The ReLU activation function is then applied element-wise to the feature maps, transforming the values according to: $\text{ReLU}(x) = \max(0, x)$. This introduces non-linearity into the model, enabling it to learn complex representations of the data. Following the activations, a max-pooling operation is applied to downsample the feature maps, reducing their spatial dimensions and controlling overfitting by summarizing the most prominent features. After the final convolutional block, the feature maps are flattened into a one-dimensional vector and passed through fully connected (dense) layers to produce a final classical feature representation: $f_{\text{CNN}}(x; \theta^c) \in \mathbb{R}^{2^n}$, where n is the number of qubits in the quantum circuit (here, $2^4 = 16$).

3.2.2. Quantum Component

The core innovation in AdeptHEQ-FL is the incorporation of a PQC, which serves as a QNN for feature processing. Unlike classical networks that manipulate continuous or discrete numerical values, quantum circuits process data encoded into quantum states. In this work, the PQC operates on 4 qubits and consists of 2 layers of Strongly Entangling Layers — a widely-used ansatz in variational quantum algorithms (Figure 2) [18, 19]. The quantum circuit performs the following steps:

Amplitude Embedding The output from the CNN, denoted by $f_{\text{CNN}}(x; \theta^c)$, is first encoded into the quantum circuit through amplitude embedding¹. This encoding maps a normalized classical vector $x \in \mathbb{R}^{2^n}$ into the amplitudes of a quantum state:

$$|\psi_x\rangle = \sum_{i=0}^{2^n-1} x_i |i\rangle \quad (3)$$

where $|i\rangle$ represents the computational basis states of the qubit system. Amplitude embedding ensures that the sum of

¹<https://docs.pennylane.ai/en/stable/code/api/pennylane.AmplitudeEmbedding.html>

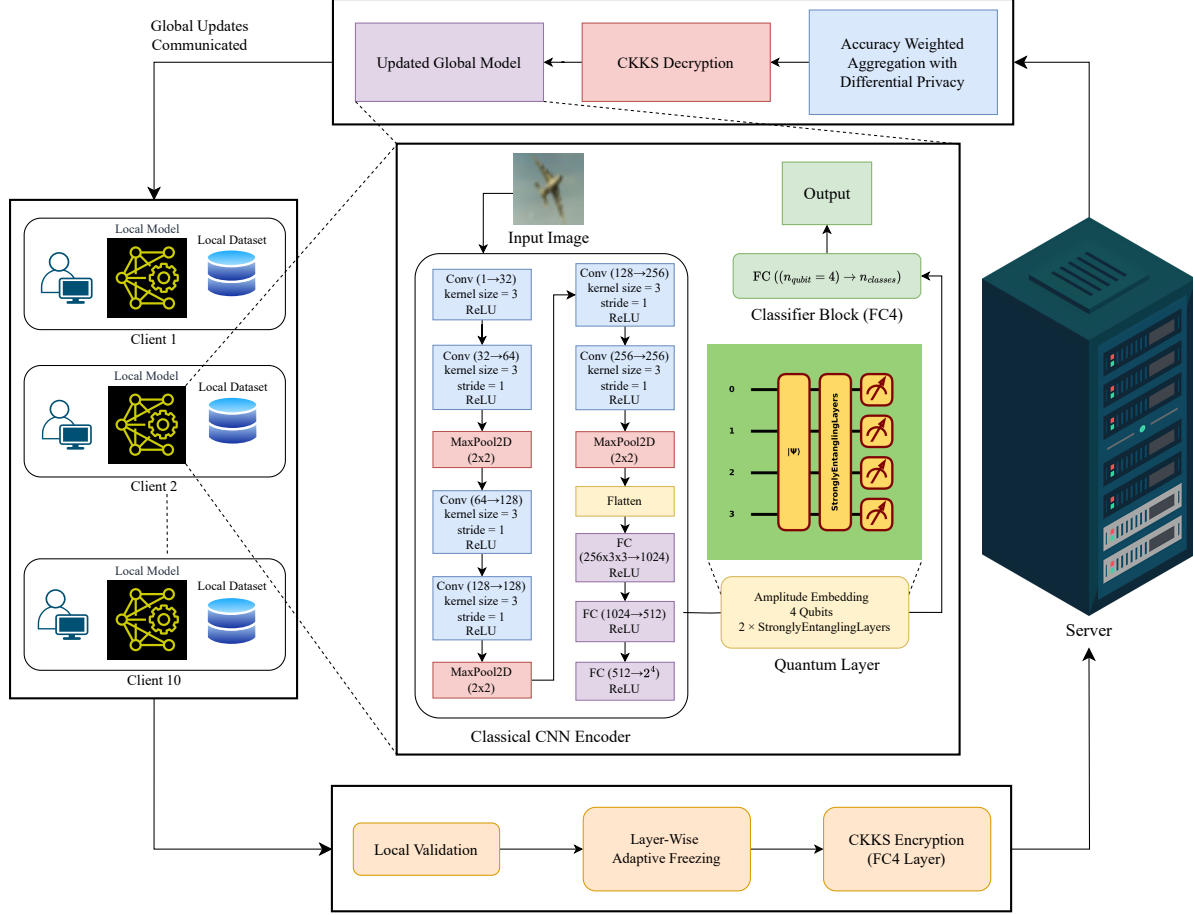


Figure 1. This flowchart provides an overview of the AdeptHEQ-FL framework and illustrates the multi-stage process, detailing the activities conducted on both the client and server sides. Each client independently executes a local training phase using a local dataset on the classical–quantum neural network. This is followed by local validation, adaptive layer freezing, and encryption of the final classifier layer. The encrypted local models are then sent to a central server for global aggregation, resulting in an improved federated model. After aggregation, the updated global model is communicated back to each client, where it is used as the local model for the next round. The **dotted line** here indicates a more detailed version of the blocks in the diagram.

the squared amplitudes equals 1, maintaining a valid quantum state:

$$\sum_{i=0}^{2^n-1} |x_i|^2 = 1 \quad (4)$$

Strongly Entangling Layers After the amplitude embedding step, the encoded quantum state undergoes a sequence of parameterized transformations and entangling operations, collectively termed as Strongly Entangling Layers. These layers are crucial for introducing both individual qubit rotations and inter-qubit correlations, allowing the quantum circuit to model complex feature interactions.

Each Strongly Entangling Layer² comprises two primary components. The first component involves a series of param-

eterized single-qubit rotation gates applied independently to each qubit. Specifically, for each qubit, a cascade of three rotations is performed in the order: $R_z(\theta^1) \rightarrow R_y(\theta^2) \rightarrow R_z(\theta^3)$. $R_z(\theta_{k,i}^1)$ rotates the qubit about the Z-axis by an angle $\theta_{k,i}^1$, $R_y(\theta_{k,i}^2)$ rotates about the Y-axis by $\theta_{k,i}^2$, followed again by $R_z(\theta_{k,i}^3)$. The rotation operations are defined as:

$$R_y(\theta) = \begin{pmatrix} \cos(\theta/2) & -\sin(\theta/2) \\ \sin(\theta/2) & \cos(\theta/2) \end{pmatrix}, \quad R_z(\phi) = \begin{pmatrix} e^{-i\phi/2} & 0 \\ 0 & e^{i\phi/2} \end{pmatrix} \quad (5)$$

Here, each rotation angle θ^j is a trainable parameter, dynamically updated through the optimization process during model training to learn an optimal data representation within the quantum Hilbert space.

The second component involves the application of entangling gates that establish quantum correlations between the qubits. In this implementation, a Controlled-NOT (CNOT)

²<https://docs.pennylane.ai/en/stable/code/api/pennylane.StronglyEntanglingLayers.html>

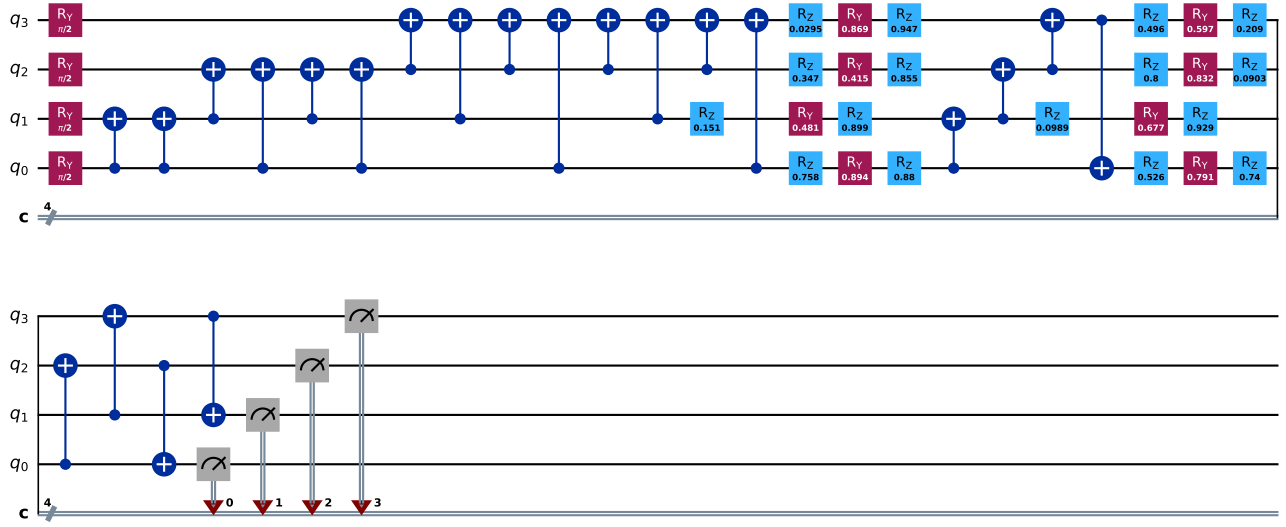


Figure 2. 4-qubit 2-layered PQC of AdeptHEQ-FL comprising amplitude embedding, two Strongly Entangling Layers (parameterized R_z , R_y , R_z rotations), CNOT-based entanglement, and projective measurements. The CNOT connectivity ensures full inter-qubit interaction within each layer.

gate is applied between selected pairs of qubits. The matrix form of the CNOT gate is:

$$\text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad (6)$$

This operation conditionally flips the target qubit when the control qubit is in the state $|1\rangle$, enabling the circuit to capture intricate interdependencies between feature dimensions that would be challenging for classical architectures to represent efficiently.

In this work, the quantum circuit employs $n = 4$ qubits and a depth of 2 Strongly Entangling Layers. The entire unitary operation implemented by the circuit can be mathematically expressed as:

$$U(\theta^q) = \prod_{l=1}^2 \left[\prod_{i=1}^4 \left(R_z(\theta_i^{l,1}) R_y(\theta_i^{l,2}) R_z(\theta_i^{l,3}) \right) \cdot \text{CNOT entanglement scheme} \right] \quad (7)$$

where each layer l sequentially applies the parameterized rotations to all qubits, followed by a set of CNOT gates arranged according to a predefined connectivity pattern. This structured layering ensures that both local qubit-level transformations and global qubit-qubit interactions are adequately captured, improving the expressive capacity of the QNN.

Measurement The circuit outputs a quantum feature vector $f_{\text{PQC}}(x; \theta^q) \in \mathbb{R}^4$, calculated by measuring the expectation value of the Pauli- Z^3 observable on each qubit after the

entangling operations:

$$f_{\text{PQC}}(x; \theta^q) = [\langle Z_1 \rangle, \langle Z_2 \rangle, \langle Z_3 \rangle, \langle Z_4 \rangle]; \quad \langle Z \rangle_i = \langle \psi | Z_i | \psi \rangle \quad (8)$$

where:

$$Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad (9)$$

The output of the CNN, $f_{\text{CNN}}(x; \theta^c)$, serves as the input to the PQC after being reshaped to match the dimensional requirements of the amplitude embedding layer.

3.2.3. Final Fully Connected Layer and Output

The quantum feature vector is then passed through a final fully connected classical layer, denoted as f_{FC4} , which maps the 4-dimensional quantum feature vector to an m -dimensional output vector, corresponding to the m classes in the used dataset:

$$f_{\text{FC4}}(x) = W_{\text{FC4}} \cdot f_{\text{PQC}}(x; \theta^q) + b_{\text{FC4}} \quad (10)$$

The final model function thus takes the form:

$$f(x; \theta) = f_{\text{FC4}}(f_{\text{PQC}}(f_{\text{CNN}}(x; \theta^c); \theta^q); \theta^{\text{FC4}}) \quad (11)$$

This hybrid architecture allows the model to harness both classical DL's feature extraction capacity and quantum circuits' potential for capturing complex, non-classical correlations in data representations.

3.3. FL Setup

In each communication round t , a fraction of clients are selected without replacement [26]. Each selected client i

³<https://docs.pennylane.ai/en/stable/code/api/pennylane.PauliZ.html>

updates the model parameters from the global model $\theta^{(t-1)}$ to local parameters $\theta_i^{(t)}$ using their local dataset \mathcal{D}_i . Clients optimize their local models using Adam optimizer with a learning rate $\eta = 10^{-3}$. Additionally, each client computes a validation accuracy $a_i^{(t)} \in [0, 1]$ on their local validation set $\mathcal{D}_i^{\text{val}}$, which guides the aggregation process.

3.4. Accuracy-Weighted Aggregation with Differential Privacy

3.4.1. Mechanism

To address non-IID data, each client privatizes their validation accuracy $a_i^{(t)} = \sum_{j=1}^{m_i} \text{correct}_j / m_i$, where $m_i = |\mathcal{D}_i^{\text{val}}|$, using the Laplace mechanism [11]:

$$\tilde{a}_i^{(t)} = \max \left(0, \min \left(1, a_i^{(t)} + \zeta \right) \right), \quad \zeta \sim \text{Lap} \left(\frac{\Delta_i}{\epsilon} \right) \quad (12)$$

where the sensitivity is $\Delta_i = 1/m_i$ (as changing one sample alters a_i by at most $1/m_i$), and $\epsilon = 1.0$ is the per-round privacy budget. Over $T = 20$ rounds, we apply advanced composition to bound the total privacy loss at $(\epsilon_{\text{total}}, \delta) = (10, 10^{-5})$ [11].

The server computes aggregation weights using a numerically stable tempered softmax [12]:

$$w_i^{(t)} = \frac{\exp((\tilde{a}_i^{(t)} - \max_j \tilde{a}_j^{(t)})/\tau)}{\sum_{k=1}^N \exp((\tilde{a}_k^{(t)} - \max_j \tilde{a}_j^{(t)})/\tau)}, \quad (13)$$

where $\tau = 0.5$ balances weight concentration, tuned empirically to prioritize high-performing clients while maintaining robustness to noise. The global model is updated as:

$$\theta^{(t)} = \sum_{i=1}^N w_i^{(t)} \theta_i^{(t)}. \quad (14)$$

Theoretical Justification The Laplace mechanism ensures $(\epsilon, 0)$ -DP per round, with sensitivity $\Delta = 1/m_i$. Advanced composition accounts for multi-round privacy loss, ensuring a total budget of $(\epsilon_{\text{total}}, \delta)$.

3.4.2. Privacy Guarantee

We formally state the following privacy result:

Theorem 1. *Each communication round of the proposed aggregation mechanism satisfies $(\epsilon, 0)$ -DP for each client's validation accuracy, where ϵ is the privacy budget per round, and sensitivity $\Delta = 1/m_i$. Over T rounds, using advanced composition [11], the total privacy guarantee is $(\epsilon_{\text{total}}, \delta)$, where $\epsilon_{\text{total}} = \sqrt{2T \log(1/\delta)}\epsilon + T\epsilon(e^\epsilon - 1)$.*

This ensures privacy amplification by composition while maintaining model utility.

3.5. Layer-Wise Adaptive Freezing

3.5.1. Mechanism

To reduce communication overhead, we compute layer importance scores based on the L2 norm of the change in the global model parameters across rounds, consistent with our experimental setup:

$$s_l^{(t)} = \left\| \theta_l^{(t)} - \theta_l^{(t-1)} \right\|_2, \quad (15)$$

where layer l indexes blocks in the model's parameter list (e.g., convolutional or fully connected layers). We maintain an exponential moving average [22]:

$$\bar{s}_l^{(t)} = \alpha \bar{s}_l^{(t-1)} + (1 - \alpha) s_l^{(t)}, \quad (16)$$

with $\alpha = 0.9$, tuned for stability. Layers are frozen if:

$$\theta_l^{(t)} = \theta_l^{(t-1)} \quad \text{if} \quad \bar{s}_l^{(t)} < \text{thr}, \quad (17)$$

where $\text{thr} = 0.001$ is a fixed absolute threshold used to determine freezing. Quantum layers (θ^q) are exempt from freezing to preserve their adaptability.

3.5.2. Rationale

The adaptive freezing strategy reduces communication overhead while ensuring that model accuracy remains largely intact. Additionally, exempting quantum layers preserves their flexibility, contributing to consistent performance gains in non-IID settings.

Quantum Layer Considerations In our hybrid classical-quantum model, quantum layers contribute essential non-linear and entangled feature transformations, crucial for modeling complex patterns in decentralized data. As such, these layers exhibit high sensitivity to client-specific data distributions and model updates. To preserve this adaptability, quantum layers are explicitly exempted from the freezing criterion in Eq. 17. This ensures the retention of quantum expressivity and prevents potential performance degradation due to premature parameter freezing.

3.6. Integration with HE

The aggregation process (Eq. 14) involves linear combinations, making it compatible with HE. In our current implementation, HE (using the CKKS scheme) is selectively applied to the parameters of the final fully connected layer (FC4). Other layer parameters are aggregated in plaintext on the server. The server uses its secret key to decrypt the aggregated FC4 layer after the weighted summation. We employ the CKKS scheme [9] with a polynomial modulus degree of 8192 and coefficient moduli bit sizes of [60, 40, 40, 60] bits. A global scaling factor of 2^{40} is used for encoding the model parameters. These parameters ensure

approximately 128-bit security and support circuits with a multiplicative depth of up to 3, which is sufficient for the weighted aggregation. The server generates Galois keys to facilitate efficient homomorphic operations. The server performs homomorphic aggregation on the encrypted layer updates without decrypting individual client contributions. After aggregation, the server uses its secret key to decrypt the resulting aggregated parameters for this layer before updating the global model and for subsequent operations, such as layer freezing analysis.

3.7. Convergence Analysis

We analyze convergence under the following assumptions: 1. The loss function \mathcal{L}_i is L -smooth. 2. The gradient variance is bounded: $\mathbb{E}\|\nabla\mathcal{L}_i(\theta)\|^2 \leq \sigma^2$. 3. The learning rate is $\eta_t = \mu/(L\sqrt{t})$, with $\mu = 0.1$.

Theorem 2 (Convergence of AdeptHEQ-FL). *After T rounds, AdeptHEQ-FL satisfies:*

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}\|\nabla\mathcal{L}(\theta^t)\|^2 \leq \frac{C_1}{\sqrt{T}} + C_2 \frac{\sigma^2 + \epsilon^{-2}}{\mu^2}, \quad (18)$$

where C_1, C_2 are constants depending on τ , the freezing threshold thr , and ϵ .

Proof Sketch: We extend the perturbed iterate framework [20], bounding errors from adaptive weights and layer freezing. The tempered softmax aligns weights with client performance, while freezing introduces bounded perturbations. The ϵ^{-2} term accounts for DP noise.

4. Experimental Setup

4.1. Simulation Tools and Environment

All experiments, including model development and FL simulations, were conducted using *Python 3.11.11*. The computational environment included NVIDIA Tesla P100 GPUs with CUDA 12.x support and multi-core Intel Xeon CPUs, providing up to 16 GB of GPU memory and 32 GB of system RAM. The primary DL framework was *PyTorch 2.5.1+cu124* [31]. For QML components, we utilized *PennyLane 0.41.1* [3] and *Qiskit 1.2.4*. HE was enabled by *TenSEAL 0.3.16* [30], implementing the CKKS scheme [9]. The FL protocol was custom-implemented, with conceptual underpinnings inspired by *PySyft 0.9.5* [33]. Data serialization used *protobuf 3.20.3*, numerical computations relied on *NumPy 1.26.4* [16], and data analysis utilized *pandas 2.2.2* [25].

4.2. Hyperparameters and Configuration

Experiments were conducted on CIFAR-10 [23] (60,000 instances of 32×32 color images), SVHN [28] (73,257 instances of 32×32 color images), and Fashion-MNIST [40]

(70,000 instances of 28×28 grayscale images), each comprising 10 classes. Normalization used dataset-specific statistics: CIFAR-10 with $\mu = (0.5, 0.5, 0.5)$ and $\sigma = (0.5, 0.5, 0.5)$, SVHN with $\mu = (0.4377, 0.4438, 0.4728)$ and $\sigma = (0.1980, 0.2010, 0.1970)$, and Fashion-MNIST with $\mu = 0.2860$ and $\sigma = 0.3530$. The datasets were distributed among 10 clients using a Dirichlet distribution with $\alpha = 0.1$ to simulate non-IID settings.

The FL simulation involved 10 clients over 20 communication rounds. Each client executed 10 local epochs with the Adam optimizer (learning rate 1×10^{-3} , batch size 32). Validation accuracy was privatized using $\epsilon = 1.0$ DP. Server-side aggregation employed the AdeptHEQ-FL method, weighting updates by privatized validation accuracies via softmax with $\tau = 0.5$. Layer-wise adaptive freezing monitored layer importance with an EMA ($\alpha = 0.9$) of parameter difference norms, freezing layers with scores below 0.001 (excluding quantum layers). Model parameters were encrypted using HE via TenSeal (CKKS scheme). Global model evaluation on a centralized test set reported test accuracy and loss after each round.

5. Results and Discussion

We evaluated three variants of our AdeptHEQ-FL framework—AdeptHEQ-FL (4-qubit, 2-layer), AdeptHEQ-FL (4-qubit, 1-layer), and AdeptHEQ-FL (2-qubit, 1-layer)—against a standard federated QNN (6 qubits, 6 layers) and a state-of-the-art FHE-FedQNN (6 qubits, 6 layers) [10], across three datasets: SVHN [28], FashionMNIST [40], and CIFAR10 [23]. Table 1 summarizes the loss and accuracy results, revealing clear trends. The AdeptHEQ-FL variant outperformed all others in accuracy across all datasets. While improvements on SVHN and FashionMNIST were modest, AdeptHEQ-FL achieved $\approx 25.43\%$ increase in accuracy compared to Standard-FedQNN and $\approx 14.67\%$ compared to FHE-FedQNN on CIFAR10, which is a comparatively complex dataset. This demonstrates AdeptHEQ-FL’s strength in handling challenging data.

We also found that performance dropped when quantum resources were reduced. Reducing qubits and layers, as seen in 4-qubit 1-layered AdeptHEQ-FL and 2-qubit 1-layered AdeptHEQ-FL, led to noticeable declines in performance, particularly on CIFAR10. This suggests that more complex datasets are more sensitive to resource constraints. Even with fewer resources (4 qubits, 2 layers) compared to FHE-FedQNN (6 qubits, 6 layers) and Standard (6 qubits, 6 layers), AdeptHEQ-FL’s performance was quite impressive. AdeptHEQ-FL’s superior performance results from its advanced aggregation strategy. Unlike FHE-FedQNN [10], which treats all client updates equally and amplifies noise in skewed data. The standard method, which uses a weighted sum of the updates, faces the same issue. AdeptHEQ-FL weights updates based on privatized validation accuracy, pri-

Table 1. Performance comparisons of different models across three datasets are shown. The table displays the average loss and accuracy in percentages for the models in our experiment across three different datasets. Each metric is reported as the mean \pm standard deviation, calculated over five experimental runs. **Bold** values indicate the best performance in each dataset column.

| Model | n_{qubits} | n_{layers} | CIFAR10 [23] | | SVHN [28] | | FashionMNIST [40] | |
|--------------------|---------------------|---------------------|-------------------------------------|------------------------------------|-------------------------------------|------------------------------------|-------------------------------------|------------------------------------|
| | | | Loss (\downarrow) | Accuracy (%) (\uparrow) | Loss (\downarrow) | Accuracy (%) (\uparrow) | Loss (\downarrow) | Accuracy (%) (\uparrow) |
| Standard-FedQNN | 6 | 6 | 1.503 \pm 0.039 | 63.60 \pm 0.45 | 0.349 \pm 0.002 | 93.22 \pm 0.05 | 0.313 \pm 0.003 | 91.96 \pm 0.09 |
| FHE-FedQNN [10] | 6 | 6 | 1.972 \pm 0.042 | 57.89 \pm 0.20 | 0.340 \pm 0.006 | 92.94 \pm 0.14 | 0.328 \pm 0.003 | 91.78 \pm 0.11 |
| AdeptHEQ-FL | 4 | 2 | 1.306 \pm 0.015 | 72.61 \pm 0.33 | 0.362 \pm 0.004 | 94.05 \pm 0.10 | 0.340 \pm 0.003 | 92.91 \pm 0.12 |
| AdeptHEQ-FL | 4 | 1 | 1.667 \pm 0.009 | 67.22 \pm 0.18 | 0.331 \pm 0.003 | 93.71 \pm 0.12 | 0.339 \pm 0.007 | 92.76 \pm 0.04 |
| AdeptHEQ-FL | 2 | 1 | 1.640 \pm 0.009 | 62.62 \pm 0.42 | 0.526 \pm 0.006 | 93.58 \pm 0.09 | 0.385 \pm 0.004 | 92.46 \pm 0.13 |

oritizing contributions from models that are better adapted. Additionally, our adaptive layer-freezing method skips updates to layers with importance scores below 0.001, reducing unnecessary computation. These innovations enable AdeptHEQ-FL to achieve strong results with fewer resources, making it effective for more complex and practical datasets.

6. Conclusion

This paper presents AdeptHEQ-FL, a novel FL framework that synergistically combines hybrid classical-quantum modeling, adaptive privacy-preserving aggregation, and dynamic communication reduction strategies. By integrating a CNN-PQC architecture with accuracy-weighted aggregation using differentially private validation accuracies, AdeptHEQ-FL effectively addresses the performance degradation typically observed under non-IID client distributions. The selective application of HE to critical model layers ensures strong privacy guarantees without incurring prohibitive overhead, while the layer-wise adaptive freezing strategy significantly reduces communication costs, allowing quantum layers to retain their expressive flexibility. Our theoretical convergence analysis and empirical results on multiple datasets confirm that AdeptHEQ-FL delivers competitive accuracy and efficiency compared to prior QFL approaches, particularly excelling on complex datasets such as CIFAR-10. The proposed framework provides a comprehensive and scalable solution for privacy-preserving, communication-efficient FL in hybrid classical-quantum environments.

Limitations While AdeptHEQ-FL shows significant improvements in accuracy and communication efficiency under privacy constraints, several limitations warrant discussion. First, AdeptHEQ-FL selectively applies HE to the final fully connected layer for tractability, leaving other layers unencrypted. Second, the framework is assessed in simulated environments, and its performance on real-world quantum hardware remains untested. Third, the convergence analysis assumes standard smoothness and bounded gradient variance, which may not hold in highly non-convex federated settings. Future work will extend encryption coverage, test

on physical devices, and generalize to larger, more complex datasets.

References

- [1] Rezak Aziz, Soumya Banerjee, Samia Bouzefrane, and Thinh Le Vinh. Exploring Homomorphic Encryption and Differential Privacy Techniques towards Secure Federated Learning Paradigm. *Future Internet*, 15(9):310, 2023. 2
- [2] Rezak Aziz, Soumya Banerjee, and Samia Bouzefrane. Privacy Preserving Federated Learning: A Novel Approach for Combining Differential Privacy and Homomorphic Encryption. In *Information Security Theory and Practice - 14th IFIP WG 11.2 International Conference, WISTP 2024, Paris, France, February 29 - March 1, 2024, Proceedings*, pages 162–177. Springer, 2024. 2
- [3] Ville Bergholm, Josh Izaac, Maria Schuld, Christian Gogolin, M. Sohaib Alam, Shah Nawaz Ahmed, Juan Miguel Arrazola, Carsten Blank, Alain Delgado, Soran Jahangiri, Keri McKiernan, Johannes Jakob Meyer, Zeyue Niu, Antal Száva, and Nathan Killoran. PennyLane: Automatic differentiation of hybrid quantum-classical computations, 2018. 7
- [4] Amandeep Singh Bhatia, Mandeep Kaur Saggi, and Sabre Kais. Application of quantum-inspired tensor networks to optimize federated learning systems. *Quantum Mach. Intell.*, 7(1):12, 2025. 2
- [5] Alessio Catalfamo, Maria Fazio, Antonio Celesti, and Massimo Villari. Privacy-Preserving in Federated Learning: A Comparison between Differential Privacy and Homomorphic Encryption across Different Scenarios. In *IEEE International Conference on Software Testing, Verification and Validation, ICST 2025 - Workshops, Naples, Italy, March 31 - April 4, 2025*, pages 451–459. IEEE, 2025. 2
- [6] Mahdi Chehimi, Samuel Yen-Chi Chen, Walid Saad, Don Towsley, and Mérouane Debbah. Foundations of Quantum Federated Learning Over Classical and Quantum Networks. *IEEE Netw.*, 38(1):124–130, 2024. 2
- [7] Lvjun Chen, Di Xiao, Zhuyang Yu, and Maolan Zhang. Secure and efficient federated learning via novel multi-party computation and compressed sensing. *Inf. Sci.*, 667:120481, 2024. 2
- [8] Yue Chen, Yufei Yang, Yingwei Liang, Taipeng Zhu, and Dehui Huang. Federated Learning with Privacy Preservation in Large-Scale Distributed Systems Using Differential Privacy

- and Homomorphic Encryption. *Informatica (Slovenia)*, 49 (13), 2025. 2
- [9] Jung Hee Cheon, Andrey Kim, Miran Kim, and Yong Soo Song. Homomorphic Encryption for Arithmetic of Approximate Numbers. In *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part I*, pages 409–437. Springer, 2017. 6, 7
- [10] Siddhant Dutta, Pavana P. Karanth, Pedro Maciel Xavier, Iago Leal de Freitas, Nouhaila Innan, Sadok Ben Yahia, Muhammad Shafique, and David E. Bernal Neira. Federated Learning with Quantum Computing and Fully Homomorphic Encryption: A Novel Computing Paradigm Shift in Privacy-Preserving ML. *CoRR*, abs/2409.11430, 2024. arXiv: 2409.11430. 1, 2, 7, 8
- [11] Cynthia Dwork and Aaron Roth. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3-4):211–407, 2013. Publisher: Now Publishers. 6
- [12] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>. 6
- [13] Zhenyuan Guo, Lei Xu, and Liehuang Zhu. FedSIGN: A sign-based federated learning framework with privacy and robustness guarantees. *Comput. Secur.*, 135:103474, 2023. 1, 2
- [14] Arti Gupta, Manish Kumar Maurya, Khyati Dhere, and Vijay Kumar Chaurasiya. Privacy-Preserving Hybrid Federated Learning Framework for Mental Healthcare Applications: Clustered and Quantum Approaches. *IEEE Access*, 12: 145054–145068, 2024. 3
- [15] Palak Handa, Anushka Saini, Siddhant Dutta, Harsh Pathak, Nishi Choudhary, Nidhi Goel, and Jasdeep Kaur Dhanao. Pcosgen-test dataset, 2024. 2
- [16] Charles R Harris, K Jarrod Millman, Stéfan J Van Der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. Array programming with NumPy. *Nature*, 585(7825): 357–362, 2020. 7
- [17] Nouhaila Innan, Muhammad Al-Zafar Khan, Alberto Marchisio, Muhammad Shafique, and Mohamed Bennai. FedQNN: Federated Learning using Quantum Neural Networks. In *International Joint Conference on Neural Networks, IJCNN 2024, Yokohama, Japan, June 30 - July 5, 2024*, pages 1–9. IEEE, 2024. 1, 2
- [18] Md Abrar Jahin, Md Sakib Hossain Shovon, Md Saiful Islam, Jungpil Shin, Muhammad Firoz Mridha, and Yuichi Okuyama. Qamplifynet: pushing the boundaries of supply chain backorder prediction using interpretable hybrid quantum-classical neural network. *Scientific Reports*, 13(1):18246, 2023. 3
- [19] Md Abrar Jahin, Md. Akmol Masud, Md Wahiduzzaman Suva, M. F. Mridha, and Nilanjan Dey. Lorentz-Equivariant Quantum Graph Neural Network for High-Energy Physics. *IEEE Transactions on Artificial Intelligence*, pages 1–11, 2025. 3
- [20] Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista A. Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael G. L. D’Oliveira, Hubert Eichner, Salim El Rouayheb, David Evans, Josh Gardner, Zachary Garrett, Adrià Gascón, Badi Ghazi, Phillip B. Gibbons, Marco Gruteser, Zaïd Harchaoui, Chaoyang He, Lie He, Zhouyuan Huo, Ben Hutchinson, Justin Hsu, Martin Jaggi, Tara Javidi, Gauri Joshi, Mikhail Khodak, Jakub Konečný, Aleksandra Korolova, Farinaz Koushanfar, Sanmi Koyejo, Tancrède Lepoint, Yang Liu, Prateek Mittal, Mehryar Mohri, Richard Nock, Ayfer Özgür, Rasmus Pagh, Hang Qi, Daniel Ramage, Ramesh Raskar, Mariana Raykova, Dawn Song, Weikang Song, Sebastian U. Stich, Ziteng Sun, Ananda Theertha Suresh, Florian Tramèr, Praneeth Vepakomma, Jianyu Wang, Li Xiong, Zheng Xu, Qiang Yang, Felix X. Yu, Han Yu, and Sen Zhao. Advances and Open Problems in Federated Learning. *Foundations and Trends® in Machine Learning*, 14(1-2): 1–210, 2021. 1, 2, 7
- [21] Hiroki Kaminaga, Feras M. Awaysheh, Sadi Alawadi, and Liina Kamm. MPCFL: Towards Multi-party Computation for Secure Federated Learning Aggregation. In *Proceedings of the IEEE/ACM 16th International Conference on Utility and Cloud Computing, UCC 2023, Taormina (Messina), Italy, December 4-7, 2023*, page 19. ACM, 2023. 2
- [22] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. 6
- [23] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical Report 0, University of Toronto, Toronto, Ontario, 2009. 2, 7, 8
- [24] Zhi-Ping Liu, Xiao-Yu Cao, Hao-Wen Liu, Xiao-Ran Sun, Yu Bao, Yu-Shuo Lu, Hua-Lei Yin, and Zeng-Bing Chen. Practical quantum federated learning and its experimental demonstration. *CoRR*, abs/2501.12709, 2025. arXiv: 2501.12709. 1, 2
- [25] Wes McKinney. Data Structures for Statistical Computing in Python, 2010. 7
- [26] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, pages 1273–1282. PMLR, 2017. 1, 5
- [27] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20-22 April 2017, Fort Lauderdale, FL, USA*, pages 1273–1282. PMLR, 2017. 2
- [28] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*, 2011. 7, 8
- [29] Msoud Nickparvar. Brain Tumor MRI Dataset, 2021. 2
- [30] OpenMined Community. TenSEAL: A library for doing Homomorphic Encryption operations on tensors, 2020. 7

- [31] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035, 2019. [7](#)
- [32] Rod Rofougaran, Shinjae Yoo, Huan-Hsin Tseng, and Samuel Yen-Chi Chen. Federated Quantum Machine Learning with Differential Privacy. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2024, Seoul, Republic of Korea, April 14-19, 2024*, pages 9811–9815. IEEE, 2024. [2](#)
- [33] Théo Ryffel, Andrew Trask, Morten Dahl, Bobby Wagner, Jason E Mancuso, Daniel Rueckert, and Jonathan Passerat-Palmbach. A generic framework for privacy preserving deep learning. In *NeurIPS Workshop on Privacy Preserving Machine Learning*, 2018. [7](#)
- [34] Enrico Sorbera, Federica Zanetti, Giacomo Brandi, Alessandro Tomasi, Roberto Doriguzzi Corin, and Silvio Ranise. Adaptive Federated Learning with Functional Encryption: A Comparison of Classical and Quantum-safe Options. *CoRR*, abs/2504.00563, 2025. arXiv: 2504.00563. [2](#)
- [35] Arnaud Grivet Sébert, Marina Checchi, Oana Stan, Renaud Sirdey, and Cédric Gouy-Pailler. Combining homomorphic encryption and differential privacy in federated learning. In *20th Annual International Conference on Privacy, Security and Trust, PST 2023, Copenhagen, Denmark, August 21-23, 2023*, pages 1–7. IEEE, 2023. [2](#)
- [36] Gazi Tanbhir and Md Farhan Shahriyar. Quantum-Inspired Privacy-Preserving Federated Learning Framework for Secure Dementia Classification. *CoRR*, abs/2503.03267, 2025. arXiv: 2503.03267. [2](#)
- [37] Anh-Tu Tran, The Dung Luong, and Xuan Sang Pham. A Novel Privacy-Preserving Federated Learning Model Based on Secure Multi-party Computation. In *Integrated Uncertainty in Knowledge Modelling and Decision Making - 10th International Symposium, IUKM 2023, Kanazawa, Japan, November 2-4, 2023, Proceedings, Part II*, pages 321–333. Springer, 2023. [2](#)
- [38] Shoaib Ullah, Madam Hussain Shah, and Adeel Anjum. Quantum Enhanced Federated Learning with Differential Privacy. In *International Conference on Frontiers of Information Technology, FIT 2024, Islamabad, Pakistan, December 9-10, 2024*, pages 1–6. IEEE, 2024. [1](#), [2](#)
- [39] Tao Wu, Yulin Deng, Qizhao Zhou, Xi Chen, and Ming Zhang. ADPHE-FL: Federated learning method based on adaptive differential privacy and homomorphic encryption. *Peer Peer Netw. Appl.*, 18(3):141, 2025. [2](#)
- [40] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: A Novel Image Dataset for Benchmarking Machine Learning Algorithms. *arXiv preprint arXiv:1708.07747*, 2017. [7](#), [8](#)
- [41] Guangfeng Yan, Shanxiang Lyu, Hanxu Hou, Zhiyong Zheng, and Linqi Song. Towards Quantum-Safe Federated Learning via Homomorphic Encryption: Learning with Gradients. *CoRR*, abs/2402.01154, 2024. arXiv: 2402.01154. [1](#), [2](#)
- [42] Xuyan Zhang, Da Huang, and Yuhua Tang. Secure Federated Learning Scheme Based on Differential Privacy and Homomorphic Encryption. In *Advanced Intelligent Computing Technology and Applications - 20th International Conference, ICIC 2024, Tianjin, China, August 5-8, 2024, Proceedings, Part V (LNAI)*, pages 435–446. Springer, 2024. [2](#)