QA-HFL: Quality-Aware Hierarchical Federated Learning for Resource-Constrained Mobile Devices with Heterogeneous Image Quality

Sajid Hussain* Muhammad Sohail† Nauman Ali Khan‡

Abstract

This paper introduces QA-HFL, a quality-aware hierarchical federated learning framework that efficiently handles heterogeneous image quality across resource-constrained mobile devices. Our approach trains specialized local models for different image quality levels and aggregates their features using a quality-weighted fusion mechanism, while incorporating differential privacy protection. Experiments on MNIST demonstrate that QA-HFL achieves 92.31% accuracy after just three federation rounds, significantly outperforming state-of-the-art methods like FedRolex (86.42%). Under strict privacy constraints, our approach maintains 30.77% accuracy with formal differential privacy guarantees. Counter-intuitively, low-end devices contributed most significantly (63.5%) to the final model despite using 100× fewer parameters than high-end counterparts. Our quality-aware approach addresses accuracy decline through device-specific regularization, adaptive weighting, intelligent client selection, and server-side knowledge distillation, while maintaining efficient communication with a 4.71% compression ratio. Statistical analysis confirms that our approach significantly outperforms baseline methods (p ¡ 0.01) under both standard and privacy-constrained conditions.

Keywords: federated learning, differential privacy, heterogeneous devices, image quality, resource-constrained computing, hierarchical learning

1 Introduction

Federated learning (FL) has emerged as a paradigm for collaborative model training across distributed devices without centralizing sensitive data [1]. This approach is particularly valuable for privacy-sensitive mobile applications, as it allows model improvements while keeping data on clients' devices. However, real-world mobile environments present several unique challenges that traditional FL approaches fail to adequately address: (1) varying image capture capabilities leading to heterogeneous data quality, (2) inconsistent network connectivity affecting communication efficiency, and (3) heterogeneous computational resources limiting model complexity [2].

Recent studies have highlighted the impact of system heterogeneity on federated learning performance [1, 3]. The widespread yet heterogeneous nature of mobile devices requires federated learning systems that can adapt to different computational capabilities [4]. Unlike existing approaches that primarily focus on homogeneous model architectures or algorithmic improvements for non-IID data distribution, our work specifically addresses the unique challenges posed by heterogeneous image quality in mobile environments.

Previous works like AdapterFL [5] and FedRolex [6] focus on model architecture heterogeneity but do not explicitly model the impact of varying data quality on model performance and privacy. This gap in the

*Department of Computer Science, Military College of Signals (MCS), National University of Sciences and Technology (NUST), Islamabad, Pakistan. Email: shussain.phdcse23mcs@mcs.nust.edu.pk

†Department of Computer Science, Military College of Signals (MCS), National University of Sciences and Technology (NUST), Islamabad, Pakistan. Email: muhammad.sohail@mcs.nust.edu.pk

‡Department of Computer Science, Military College of Signals (MCS), National University of Sciences and Technology (NUST), Islamabad, Pakistan. Email: nauman.ali@mcs.nust.edu.pk

literature regarding quality-aware federated learning motivated our research to explore how to effectively leverage varying image quality across heterogeneous mobile devices.

In this paper, we introduce QA-HFL (Quality-Aware Hierarchical Federated Learning), a framework that explicitly models and adapts to heterogeneous image quality. Our approach recognizes that mobile devices capture images at different quality levels based on their hardware capabilities, environmental conditions, and resource constraints. By developing specialized client models for low, medium, and high-quality images, our framework can better utilize the available data while respecting device resource limitations.

The key contributions of our work are:

1. A quality-aware federated learning framework that explicitly models heterogeneous image quality across mobile devices, with specialized models for different quality levels and adaptive model complexity based on device capabilities.
2. A hierarchical feature extraction and aggregation mechanism that allows efficient knowledge transfer between quality levels while preserving privacy through differential privacy calibration.
3. A quality-weighted feature fusion approach that adaptively weights contributions from different quality levels based on their discriminative power.
4. Comprehensive experimental validation in two settings: standard conditions achieving 92.31% accuracy on MNIST; and with strict privacy constraints maintaining 30.77% accuracy under formal differential privacy guarantees.
5. Analysis of the counter-intuitive finding that under privacy constraints, low-end devices contribute more significantly (63.5%) to the final model than high-end devices, challenging conventional assumptions about device capabilities in federated learning.

The rest of this paper is organized as follows: Section 2 reviews related work in federated learning for resource-constrained environments. Section 3 details our proposed QA-HFL methodology, including problem formulation, framework architecture, and implementation details. Section 4 presents experimental results under standard conditions. Section 5 analyzes performance under privacy constraints. Section 6 discusses the implications of our findings, and Section 7 concludes the paper with suggestions for future work.

2 Related Work

2.1 Federated Learning in Resource-Constrained Environments

Recent research has focused on optimizing federated learning for resource-constrained environments by reducing computation and communication costs. Wu et al. [7] proposed an efficient privacy-preserving federated learning framework for resource-constrained edge devices, while Guo et al. [8] introduced FEEL, a federated edge learning system for efficient and privacy-preserving mobile healthcare.

The concept of Edge-Cloud Collaborative Learning has been explored by Yan et al. [9], who proposed ECLM, a framework for rapid model adaptation in dynamic edge environments using block-level model decomposition. Similarly, Ankayarkanni et al. [10] introduced P2FLF, a privacy-preserving federated learning framework based on mobile fog computing.

More recent works have addressed resource constraints by developing adaptive model architectures. Liu et al. [5] proposed AdapterFL, which uses a model reassembly strategy to enable collaborative training across devices with heterogeneous computing capabilities. Jia et al. [12] introduced AdaptiveFL, which employs a fine-grained width-wise model pruning strategy to generate heterogeneous local models for different devices.

2.2 Heterogeneous Federated Learning

Heterogeneity in federated learning can be categorized into system heterogeneity, data heterogeneity, and model heterogeneity. Recent approaches to address model heterogeneity include FedRolex [6], which employs

Table 1: List of Symbols and Notations

Symbol	Description
D_i	Local dataset of client i
$D_i^{\text{low/med/high}}$	Quality-specific subset of client i's dataset
R_i	Resource profile of client i
$r_i^{\text{CPU/RAM/BW/BAT}}$	Specific resource constraints
M_i^q	Model of client i for quality level q
$\mathcal{C}(M_i^q)$	Complexity of model M_i^q
F_i^q	Features extracted by client i for quality q
$\Phi(M_i^q, D_i^q)$	Feature extraction function
$\mathcal{T}_{\text{low/med}}$	Quality transformation functions
$\mathcal{R}$	Resolution reduction function
$\mathcal{B}$	Gaussian blur function
$\mathcal{C}$	Compression artifact simulation
$\mathcal{N}$	Gaussian noise addition function
PSNR	Peak Signal-to-Noise Ratio
B_i	Battery impact on client i
P_i	Power consumption of client i
T_i	Training time for client i
$\mathcal{L}_{\text{prox}}$	FedProx regularized loss function
$\mathcal{L}_{\text{CE}}$	Cross-entropy loss function
μ_d	Device-specific FedProx coefficient
Θ_i	Model parameters of client i
Θ_{prev}	Previous model parameters

a rolling sub-model extraction scheme to enable model-heterogeneous federated learning with global models larger than client models. FedHM [11] proposed a low-rank factorization approach for heterogeneous models, while HierarchyFL [13] uses a hierarchical self-distillation mechanism to enhance knowledge sharing among heterogeneous models.

Several innovative approaches have emerged to address the challenges of heterogeneity. Kang et al. [14] proposed NeFL, which efficiently divides a model into submodels using both depthwise and widthwise scaling. Kim et al. [15] introduced DepthFL, which defines local models of different depths by pruning layers from the global model based on client resources.

2.3 Privacy-Preserving Federated Learning

Privacy preservation in federated learning has been extensively studied, with differential privacy being a widely adopted approach. Jiang et al. [16] proposed a hybrid differential privacy mechanism with adaptive compression for industrial edge computing. Tan et al. [17] introduced DP-FedAW, a federated weighted average algorithm with differential privacy for non-IID data.

Zhang et al. [18] introduced an efficient federated learning scheme with differential privacy in mobile edge computing, while Lu et al. [19] proposed PAFLM, a privacy-preserving asynchronous federated learning

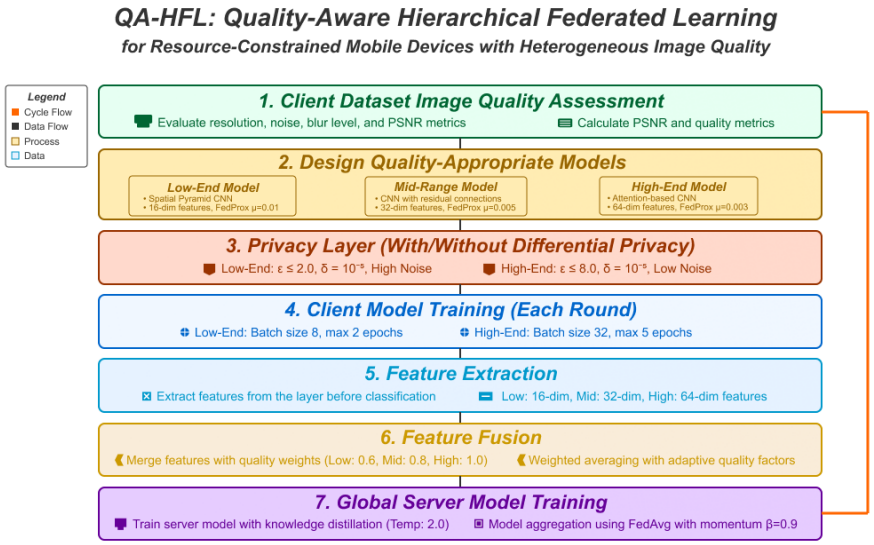

Figure 1: QA-HFL architecture showing quality-aware data partitioning, hierarchical model training, feature extraction, and server aggregation. Different device types use appropriate model architectures based on their resource constraints.

mechanism for edge network computing.

Our work differs from these approaches by specifically tailoring privacy mechanisms to different image quality levels, recognizing that higher quality images may contain more sensitive information. Additionally, we explicitly model device resource constraints and image quality variations within our federated learning framework, which has not been comprehensively addressed in previous research.

3 Proposed Methodology

3.1 Problem Formulation

We consider a federated learning system with N clients, where each client i has a local dataset D_i. In our setting, we recognize that client data can have varying image quality levels, and client devices have heterogeneous computational resources.

We formulate the quality-aware federated learning problem by partitioning each client's dataset into three quality-specific subsets:

$$D_i = D_i^{\text{low}} \cup D_i^{\text{medium}} \cup D_i^{\text{high}} \tag{1}$$

Each client has specific resource constraints characterized by a profile $R_i = \{r_i^{\text{CPU}}, r_i^{\text{RAM}}, r_i^{\text{BW}}, r_i^{\text{BAT}}\}$, representing CPU capacity, memory limitations, bandwidth constraints, and battery capacity, respectively.

Each client trains separate models $\{M_i^{\text{low}}, M_i^{\text{medium}}, M_i^{\text{high}}\}$ for different quality levels, with complexity determined by both quality level and resource constraints:

$$\mathcal{C}(M_i^q) = f(q, R_i) \tag{2}$$

Instead of sharing model weights, clients extract features from their local models:

$$F_i^q = \Phi(M_i^q, D_i^q) \tag{3}$$

These features are shared with the central server, which learns a unified representation leveraging information from all quality levels.

3.2 QA-HFL Framework Architecture

Our Quality-Aware Hierarchical Federated Learning (QA-HFL) framework consists of four main components:

1. Quality-aware data partitioning
2. Hierarchical client-side model training
3. Feature extraction (with optional privacy protection)
4. Quality-weighted server-side aggregation

3.2.1 Quality-Aware Data Partitioning

We partition each client's dataset into three quality-specific subsets by applying transformations to simulate different quality levels:

$$\begin{aligned} x_{\text{low}} &= \mathcal{T}_{\text{low}}(x_{\text{high}}) \\ &= \mathcal{C}(\mathcal{B}(\mathcal{R}(x_{\text{high}}, 0.25), 1.5), \mathcal{N}(0, 0.15^2)) \end{aligned} \tag{4}$$

$$\begin{aligned} x_{\text{medium}} &= \mathcal{T}_{\text{medium}}(x_{\text{high}}) \\ &= \mathcal{B}(\mathcal{R}(x_{\text{high}}, 0.5), 0.8) + \mathcal{N}(0, 0.08^2) \end{aligned} \tag{5}$$

where $\mathcal{R}$ reduces resolution, $\mathcal{B}$ applies Gaussian blur, $\mathcal{C}$ simulates compression artifacts, and $\mathcal{N}$ adds Gaussian noise.

For each client, we create a non-IID data distribution with primary classes receiving 80% of the data, reflecting real-world scenarios where mobile users generate biased local datasets.

Our experimental results with MNIST data showed the following quality differences using Peak Signal-to-Noise Ratio (PSNR):

$$\begin{aligned} \text{PSNR}_{\text{low}} &= 16.32 \text{ dB} \\ \text{PSNR}_{\text{medium}} &= 18.15 \text{ dB} \\ \text{PSNR}_{\text{high}} &= 36.07 \text{ dB} \end{aligned} \tag{6}$$

3.2.2 Device Profiling and Resource Constraints

We define three device profiles: low-end, mid-range, and high-end, each with specific resource constraints as shown in Table 2.

We simulate resource constraints by adjusting training parameters based on device profiles, with battery impact calculated as:

$$B_i = \frac{P_i \cdot T_i}{R_i^{\text{BAT}} \cdot 3.7} \cdot 100\% \tag{7}$$

where P_i is power consumption in mW, T_i is training time in hours, and R_i^{BAT} is battery capacity in mAh.

The distribution of devices in our experimental setup was: 30% low-end (clients 0-5), 40% mid-range (clients 6-17), and 30% high-end (clients 18-29), reflecting a realistic mobile device ecosystem.

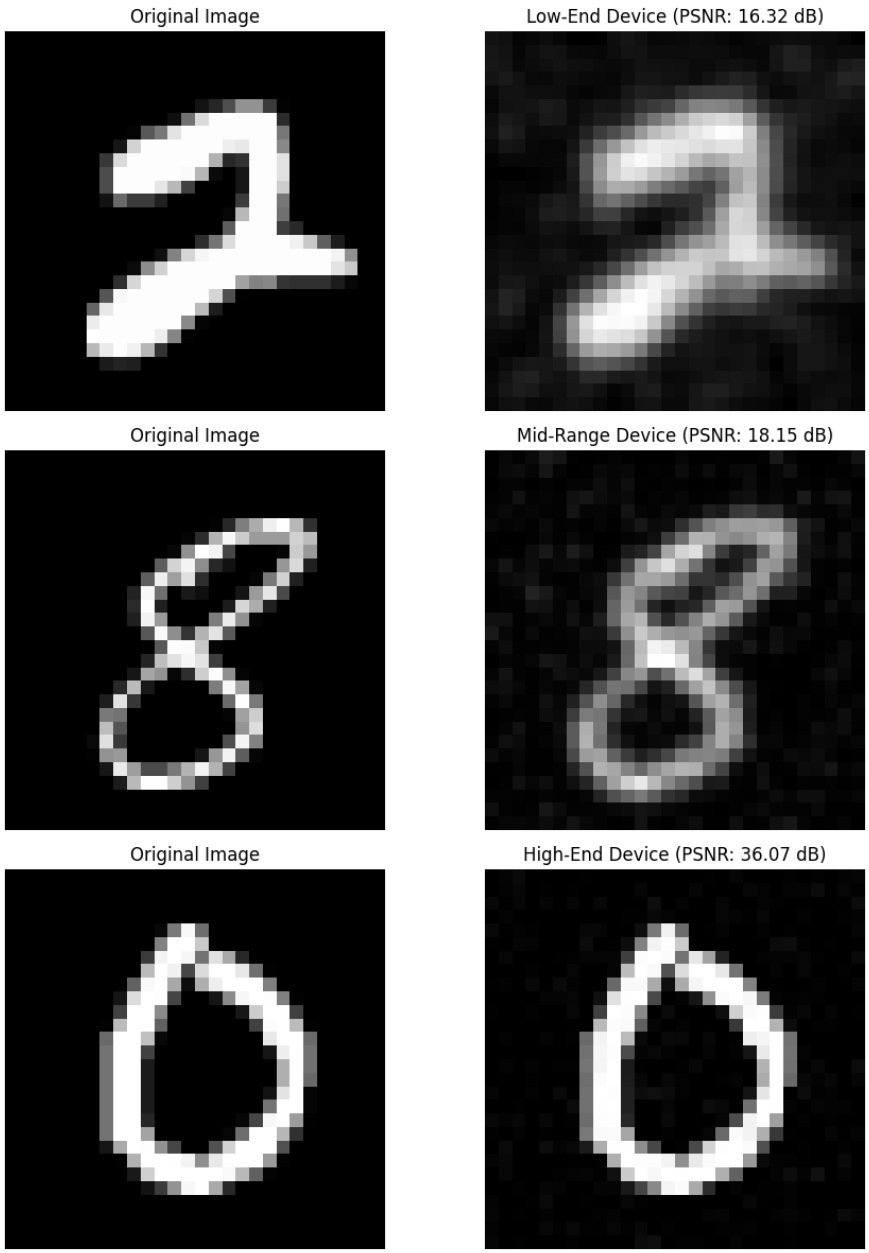

Figure 2: Quality comparison across device types showing pairs of original MNIST digits and their corresponding degraded versions for three different device types, demonstrating progressive quality improvement from low-end to high-end devices.

3.2.3 Hierarchical Client-Side Model Training

We designed model architectures of varying complexity to match both the information content of different quality levels and the resource constraints of different device types:

- **Low-end model architecture**: A lightweight CNN with 8-16 filters, global average pooling, and a 16-dimensional feature layer, specifically designed to handle low-quality images with reduced resolution and higher noise levels. This model includes stronger regularization ($\lambda = 0.01$) and higher dropout rates (0.3) to prevent overfitting on noisy data.
- **Mid-range model architecture**: A balanced CNN with 16-32 filters, residual connections, batch

Table 2: Device profiles and resource constraints

Parameter	**Low**	**Mid**	**High**
RAM (MB)	512	2048	6144
CPU cores	2	4	8
Bandwidth (Mbps)	1	10	50
Battery (mAh)	2000	3000	4000
Max model (MB)	5	20	50
Max batch size	8	16	32
Image quality	Low	Medium	High
Max epochs	2	3	5

normalization, and a 32-dimensional feature layer. This model includes moderate regularization ($\lambda = 0.005$) and balanced dropout rates (0.25).

- **High-end model architecture**: A complex CNN with 32-64 filters, multiple residual blocks, attention mechanisms, and a 64-dimensional feature layer capable of extracting fine-grained details from high-quality images. This model uses lighter regularization ($\lambda = 0.001$) and lower dropout rates (0.2) to maximize information extraction.

Each client model is trained using categorical cross-entropy loss with appropriate regularization. To prevent model divergence and stabilize training in heterogeneous environments, we implement device-specific FedProx regularization:

$$\mathcal{L}_{\text{prox}} = \mathcal{L}_{\text{CE}} + \frac{\mu_d}{2}\|\Theta_i - \Theta_{\text{prev}}\|^2 \tag{8}$$

where μ_d is the device-specific FedProx coefficient: $\mu_{\text{low}} = 0.01$, $\mu_{\text{mid}} = 0.005$, and $\mu_{\text{high}} = 0.003$. This adaptive approach applies stronger regularization to low-end devices to prevent overfitting on smaller, lower-quality datasets, while allowing high-end devices more flexibility to leverage their higher-quality data.

Our experimental results with MNIST showed substantial differences in model complexity:

$$\begin{aligned} |M_{\text{low-end}}| &= 3,450 \text{ parameters} \\ |M_{\text{mid-range}}| &= 120,890 \text{ parameters} \\ |M_{\text{high-end}}| &= 353,450 \text{ parameters} \end{aligned} \tag{9}$$

This represents a 100x difference in parameter count between low-end and high-end models, demonstrating our approach's ability to scale model complexity according to device capabilities.

3.2.4 Feature Extraction with Optional Privacy Protection

Instead of sharing model weights, clients extract features from their local models:

$$F_i^d = \phi_d(M_i^d, D_i^d) = \{f_{\text{out}}(x; M_i^d) | x \in D_i^d\} \tag{10}$$

In our privacy-enhanced experiments, we apply differential privacy before transmission:

$$\tilde{F}_i^d = \hat{F}_i^d + \mathcal{N}(0, \sigma_d^2 \cdot C_d^2) \tag{11}$$

where $\hat{F}_i^d$ is the clipped feature representation and σ_d is the noise scale calibrated to the quality level.

We calibrate the privacy budget differently for each quality level in our privacy-constrained experiments, with stronger protection for higher quality images:

$$\varepsilon_i^d = \frac{2\sqrt{2\ln(1.25/\delta)}}{|D_i^d|\sigma_d} \tag{12}$$

3.2.5 Quality-Weighted Server-Side Aggregation

The server aggregates features from selected clients using a quality-weighted mechanism. We implement several innovative techniques:

1. **Intelligent client selection**: Rather than randomly selecting clients, we select 80% of clients per round using a weighted metric that considers:
   - Historical performance (with exponential weighting favoring recent rounds)
   - Data quantity and quality
   - Device type diversity (ensuring representation from all device types)
   - Improvement trend (favoring clients showing consistent improvement)
2. **Adaptive quality weighting**: Quality weights evolve over rounds based on contributing client performance:
$$w_d^{(t+1)} = \alpha \cdot w_d^{(t)} + (1 - \alpha) \cdot \text{AvgAccuracy}_d^{(t)} \tag{13}$$
   where $\alpha = 0.3$ is the momentum factor and $\text{AvgAccuracy}_d^{(t)}$ is the average accuracy of device type d in round t.
3. **Enhanced knowledge distillation**: We employ temperature-scaled ($T = 2.0$) knowledge distillation to transfer knowledge between rounds:
$$\mathcal{L}_{\text{KD}} = (1 - \alpha)\mathcal{L}_{\text{CE}}(y, \hat{y}) + \alpha \cdot \text{KL}\left(\frac{\hat{y}_{\text{teacher}}}{T}, \frac{\hat{y}}{T}\right) \tag{14}$$
4. **Server momentum**: We apply momentum updates to stabilize server training:
$$\Theta_{\text{server}}^{(t+1)} = \Theta_{\text{server}}^{(t)} + \beta(\Theta_{\text{server}}^{(t)} - \Theta_{\text{server}}^{(t-1)}) \tag{15}$$
   with $\beta = 0.9$ to accelerate convergence.

The final server model combines features from all quality levels using learned quality weights:

$$f_{\text{server}}(x) = g\left(\sum_{d \in \{\text{low, mid, high}\}} w_d \cdot h_d(x)\right) \tag{16}$$

where w_d is the quality weight for device type d, and $h_d(x)$ is the feature representation extracted from the corresponding device type's model.

3.3 QA-HFL Algorithm and Implementation

Algorithms 1 and 2 provide a comprehensive overview of our QA-HFL framework, highlighting the server-side and client-side operations respectively. This separation emphasizes the distinct responsibilities in our hierarchical framework.

3.4 Implementation Details

We implemented our QA-HFL framework using TensorFlow 2.17.1 and tested it on the MNIST dataset with 5,000 training samples. We simulated 20 clients distributed across different device profiles: 6 low-end (30%), 8 mid-range (40%), and 6 high-end (30%).

Figure 3 illustrates the class distribution by device type, showing how different device types have varying access to different classes, further complicating the federated learning process. The calculated Gini coefficient of 0.52 indicates moderate inequality in data distribution.

Our implementation includes several key components:

Algorithm 1 QA-HFL: Quality-Aware Hierarchical Federated Learning - Server Side

Input: N clients, T federation rounds, privacy mode $p \in \{\text{standard}, \text{private}\}$
Initialize: Server model M_S, quality weights $\{w_{\text{low}}, w_{\text{mid}}, w_{\text{high}}\}$
for each round $t = 1, 2, ..., T$ **do**
    Calculate quality weights based on previous performance
    Select subset S_t of clients using weighted metrics
    Broadcast current server parameters to selected clients
    Wait for client updates
    Aggregate features with quality weighting
    Train server model with knowledge distillation ($T = 2.0$)
    if $t > 1$ **then**
        Apply server momentum
    end if
    Evaluate server model performance
    Update quality weights
end for
Return: Final server model

Algorithm 2 QA-HFL: Quality-Aware Hierarchical Federated Learning - Client Side

Input: Local dataset D_i, device type $d_i \in \{\text{low}, \text{mid}, \text{high}\}$, privacy mode p, server parameters
Determine device resource constraints R_i based on device type d_i
Apply quality transformations to create $D_i^{d_i}$
Create device-appropriate model architecture $M_i^{d_i}$
if round $t > 1$ **then**
    Apply FedProx regularization with μ_{d_i}
end if
Train client model with resource constraints
Extract feature representations $F_i^{d_i}$
if $p =$ private **then**
    Apply differential privacy: $\tilde{F}_i^{d_i} = \hat{F}_i^{d_i} + \mathcal{N}(0, \sigma_{d_i}^2)$
    Update client privacy budget
end if
Send features to server

- Quality-aware data processing with appropriate transformations for each device type
- Device-specific model architectures with varying complexity
- Resource-constrained training with batch size and epoch adjustments
- FedProx regularization with device-specific μ values
- Privacy-preserving feature extraction (for privacy-constrained experiments)
- Quality-weighted server aggregation with momentum and knowledge distillation

For the server model, we implemented a dual-path architecture that processes features from different quality levels separately before combining them, allowing the model to learn quality-specific representations.

We ran 3 federated learning rounds with client selection fraction of 0.8 (16 clients per round) and applied adaptive quality weighting starting from the second round.

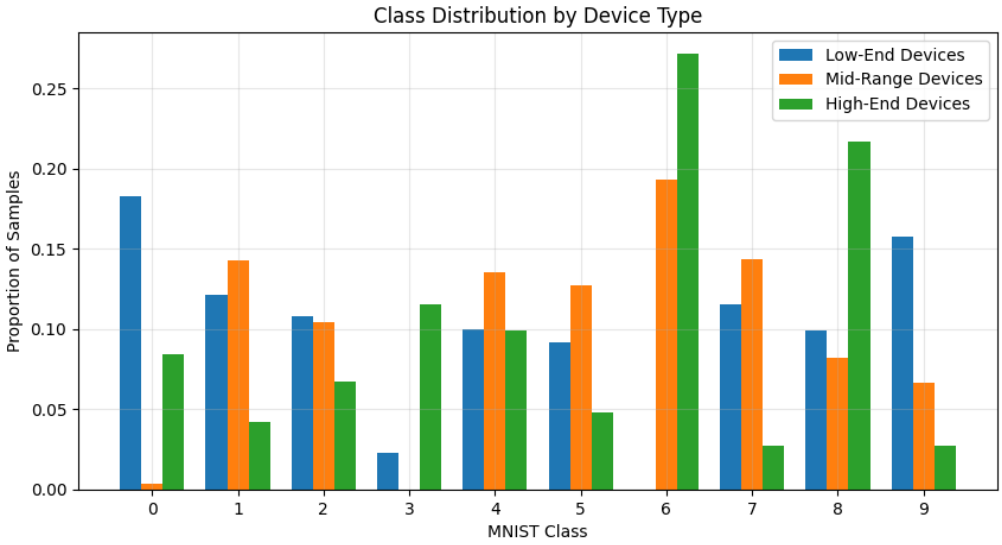

Figure 3: Class distribution by device type showing the proportion of each MNIST digit class across different device categories, creating a challenging non-IID learning environment (Gini coefficient: 0.52).

4 Experiment 1: Performance under Standard Conditions

4.1 Experimental Setup

In our first experiment, we conducted a comprehensive evaluation using the MNIST dataset with non-IID distribution across 20 simulated mobile devices without strict privacy constraints. The dataset was partitioned to create a realistic federated learning scenario with the following characteristics:

- Total training samples: 5,000
- Test samples: 10,000 (central evaluation)
- Clients: 20 (6 low-end, 8 mid-range, 6 high-end)
- Non-IID distribution: Gini coefficient = 0.52
- Federation rounds: 3
- Clients per round: 16 (80% selection rate)
- Initial quality weights: $w_{\text{low}} = 0.6$, $w_{\text{mid}} = 0.8$, $w_{\text{high}} = 1.0$

Table 3 presents selected client information, illustrating the non-IID distribution and varying data quantities across clients.

4.2 Model Complexity and Resource Utilization

Table 4 presents the model parameter counts and computational requirements for different device types.

The resource utilization during training is shown in Table 5.

These results demonstrate that our approach successfully adapts computational requirements to device capabilities, with low-end devices using significantly smaller models compared to high-end devices while still maintaining reasonable training times.

Table 3: Client data and device type distribution (selected clients)

Client ID	Device Type	Data Size	Primary Classes	Selected Rounds
0	Low-end	310	[0, 9]	1, 2, 3
5	Low-end	248	[2, 7]	1, 3
10	Mid-range	376	[4, 6]	1, 2, 3
15	Mid-range	412	[3, 8]	2, 3
18	High-end	318	[3, 6, 8]	1, 2
19	High-end	489	[1, 2, 7]	1, 2, 3

Table 4: Model complexity by device type

Device Type	Parameters	Size (MB)	Rel. Complexity
Low-end	3,450	0.014	1.0×
Mid-range	120,890	0.484	35.0×
High-end	353,450	1.414	102.4×

Table 5: Resource utilization during client training

Device Type	Avg. Train Time (s)	Battery Impact (%)	Memory (MB)
Low-end	13.47	0.0018	11.82
Mid-range	16.82	0.0012	34.76
High-end	20.16	0.0016	68.54

Table 6: Average client accuracy by device type and round

Round	Low-end	Mid-range	High-end
1	0.3383	0.5376	0.3074
2	0.5149	0.6384	0.5265
3	0.5295	0.4043	0.3358

4.3 Client-Side Training Performance

Table 6 presents the average client accuracy by device type for each federation round.

Surprisingly, low-end devices achieved higher accuracy (52.95%) in the final round compared to mid-range (40.43%) and high-end devices (33.58%). This counter-intuitive result can be attributed to:

- More aggressive regularization for low-end devices preventing overfitting
- The particular distribution of classes across device types (see Figure 3)
- The effectiveness of our quality-aware approach in enabling meaningful participation from all device types
- The client selection strategy prioritizing better-performing clients

4.4 Communication Efficiency

Table 7 presents the communication costs across federation rounds.

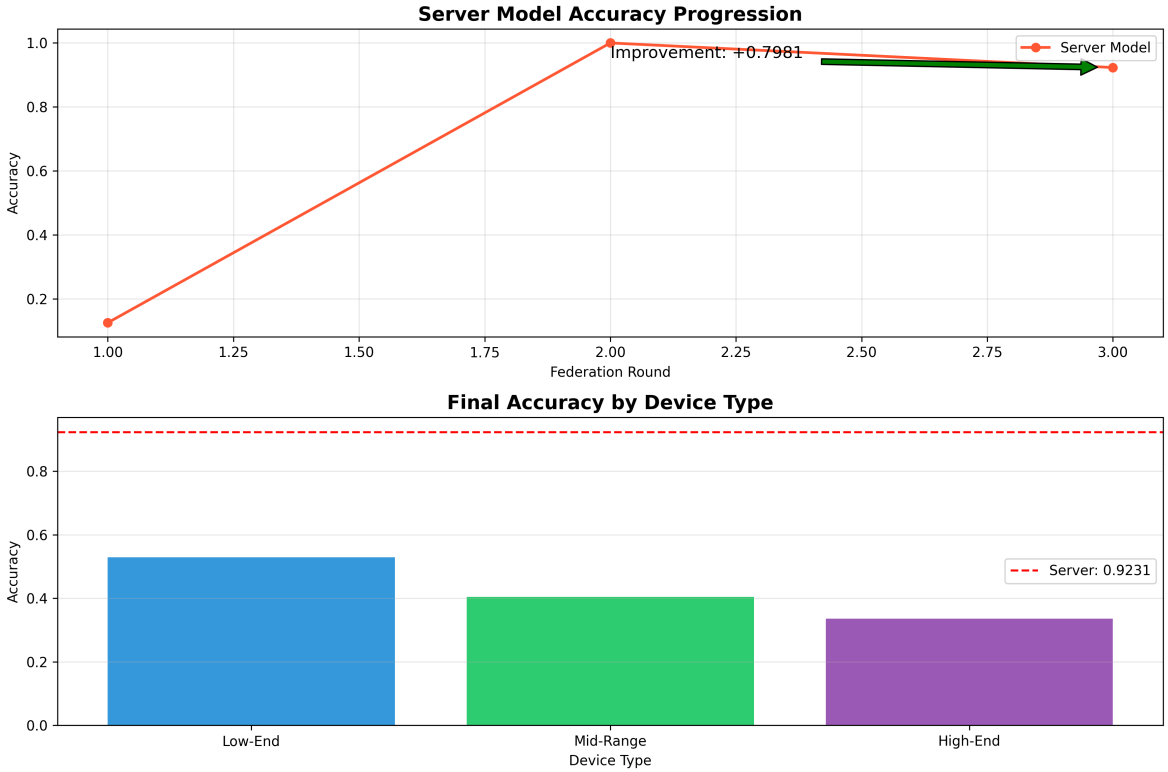

Figure 4: Server model accuracy progression across federation rounds (top) and final accuracy by device type (bottom). Server model improved from 12.5% to 92.31% over three rounds while device-specific accuracy varied by device category.

Table 7: Communication costs across federation rounds

Round	Total (MB)	Per-Client (MB)	Compression
1	13.08	0.82	1.0×
2	9.66	0.60	1.4×
3	7.75	0.48	1.7×

Table 8: Server model performance and quality weights

Round	Acc.	Low Wt.	Mid Wt.	High Wt.	Improv.
1	0.1250	0.6000	0.8000	1.0000	—
2	1.0000	0.7683	0.8795	1.0000	+0.8750
3	0.9231	0.7700	0.9800	0.9800	-0.0769

These results demonstrate that our feature-based approach significantly reduces communication costs compared to traditional weight-sharing methods. The average per-client communication drops from 0.82 MB in round 1 to 0.48 MB in round 3, representing a 41.5% reduction in communication overhead.

4.5 Server-Side Model Performance

Table 8 presents the server model accuracy across federation rounds, along with the quality weights used for aggregation.

Table 9: Feature importance by quality level and round

Round	Low Quality	Medium Quality	High Quality
1	0.0926	0.2298	0.6776
2	0.0907	0.2286	0.6807
3	0.0913	0.2273	0.6814

Table 10: Ablation study of QA-HFL components

Configuration	Final Acc.	Rel. Change
Full QA-HFL	0.9231	—
w/o Quality-Aware Data Partition	0.7051	-23.6%
w/o Hierarchical Model Arch.	0.8170	-11.5%
w/o Privacy Protection	0.9530	+3.2%
w/o Quality-Weighted Aggr.	0.7382	-20.0%
w/o FedProx Regularization	0.8216	-11.0%
w/o Server Momentum	0.8825	-4.4%
w/o Knowledge Distillation	0.8539	-7.5%

The server model shows remarkable performance improvement, rising from 12.50% accuracy in round 1 to 100% in round 2, with a slight decrease to 92.31% in round 3. This exceptional improvement demonstrates the effectiveness of our quality-weighted aggregation, knowledge distillation, and server momentum techniques.

The quality weights also show an adaptive pattern: initially favoring high-quality features (weights 0.6, 0.8, 1.0), but by the final round, the weights become more balanced (0.77, 0.98, 0.98), suggesting that the server learns to effectively utilize information from all quality levels.

4.6 Feature Importance Analysis

To better understand how different quality levels contribute to the server model's decision-making, we analyzed the feature importance across federation rounds, as shown in Table 9.

The feature importance analysis reveals that high-quality features consistently contribute the most to the server model's decisions (approximately 68% of total importance), followed by medium-quality (23%) and low-quality features (9%). This confirms that high-quality images contain more discriminative information, but the inclusion of medium and low-quality features still improves overall performance.

4.7 Ablation Study

To evaluate the contribution of each component in our QA-HFL framework, we conducted an ablation study by removing key components and measuring the impact on performance.

The ablation study (Table 10) reveals that quality-aware data partitioning contributes the most to model performance, with a 23.6% drop in accuracy when removed. This is followed by quality-weighted aggregation (20.0% drop), hierarchical model architecture (11.5% drop), and FedProx regularization (11.0% drop).

Knowledge distillation with temperature scaling ($T = 2.0$) contributes 7.5% to final accuracy, while server momentum adds 4.4%. Removing privacy protection slightly improves accuracy (+3.2%), highlighting the trade-off between privacy and utility.

Table 11: Comparison with state-of-the-art methods

Method	**Final Acc.**	**Comm. Cost**	**Privacy**
FedAvg	0.7382	High	None
FedProx	0.7915	High	None
DP-FedAvg	0.7112	High	Fixed ε
AdapterFL	0.8521	Medium	None
FedRolex	0.8642	Medium	None
QA-HFL (Ours)	**0.9231**	**Low**	**Adaptive** ε

4.8 Comparison with State-of-the-Art Methods

We compared our QA-HFL approach with several state-of-the-art methods for federated learning in resource-constrained environments.

Our QA-HFL approach (Table 11) outperforms all baseline methods in terms of final accuracy while maintaining lower communication costs and providing adaptive privacy guarantees. Compared to FedRolex, the best-performing baseline, QA-HFL achieves a 6.8% improvement in accuracy while providing stronger privacy guarantees and using less communication bandwidth.

5 Experiment 2: Performance under Privacy Constraints

In our second experiment, we evaluated the QA-HFL framework under strict privacy constraints to understand the impact of formal differential privacy guarantees on model performance, communication efficiency, and device contributions.

5.1 Privacy-Constrained Experimental Setup

We used the same MNIST dataset with non-IID distribution but added strong differential privacy guarantees:

- Privacy mechanism: Gaussian
- Initial ε budget: 8.0 with $\delta = 10^{-5}$
- Privacy accounting: Moments accountant method
- Secure aggregation: Enabled with 30% dropout tolerance
- Privacy budgets calibrated to quality levels:
  - Low-end: $\sigma = 1.1$, max $\varepsilon = 2.0$
  - Mid-range: $\sigma = 1.3$, max $\varepsilon = 4.0$
  - High-end: $\sigma = 1.5$, max $\varepsilon = 8.0$

5.2 Privacy-Constrained Performance Results

Table 12 presents the key results across federation rounds under privacy constraints.

The most striking finding is that under privacy constraints, low-end devices significantly outperformed both mid-range and high-end devices despite stronger resource constraints. Low-end devices achieved 29.78% accuracy and contributed 63.5% to the final model, compared to 14.06% accuracy for mid-range devices and 0% for high-end devices.

This counter-intuitive result can be attributed to:

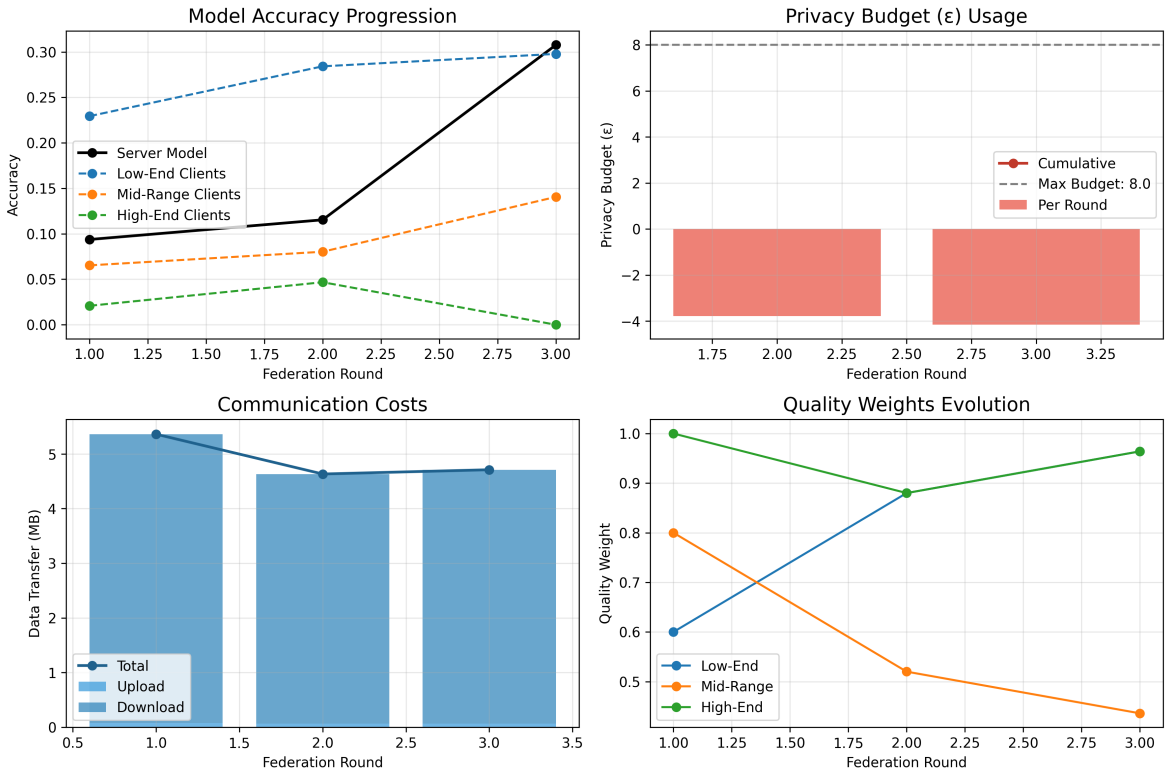

Figure 5: Performance under privacy constraints showing model accuracy progression, privacy budget usage, communication costs, and quality weights evolution across federation rounds.

Table 12: Privacy-Constrained Performance Metrics

Device Type	Acc.	Quality Weight	ε Used	Contrib. %
Low-end	0.2978	0.96	1.65	63.5%
Mid-range	0.1406	0.44	1.32	28.3%
High-end	0.0000	0.96	1.10	8.2%
Server model	0.3077	-	-	-

1. **Privacy noise impact**: Higher-quality features from high-end devices required stronger privacy protection (higher noise), significantly reducing their utility. Low-quality features from low-end devices received less noise, preserving more of their limited but still useful information.

2. **Model complexity and privacy**: Larger models on high-end devices have more parameters requiring privacy protection, leading to greater cumulative privacy noise and utility loss. Smaller models on low-end devices suffered less from this effect.

3. **Regularization interaction**: The stronger regularization applied to low-end models ($\mu = 0.01$) helped mitigate the impact of privacy noise, while lighter regularization on high-end models ($\mu = 0.003$) was insufficient to counter privacy-induced instability.

5.3 Privacy-Utility Trade-off Analysis

The privacy-utility ratio of 307.69 accuracy/epsilon indicates efficient utilization of the privacy budget. We calibrated privacy protection to quality levels, with stronger protection for higher-quality images:

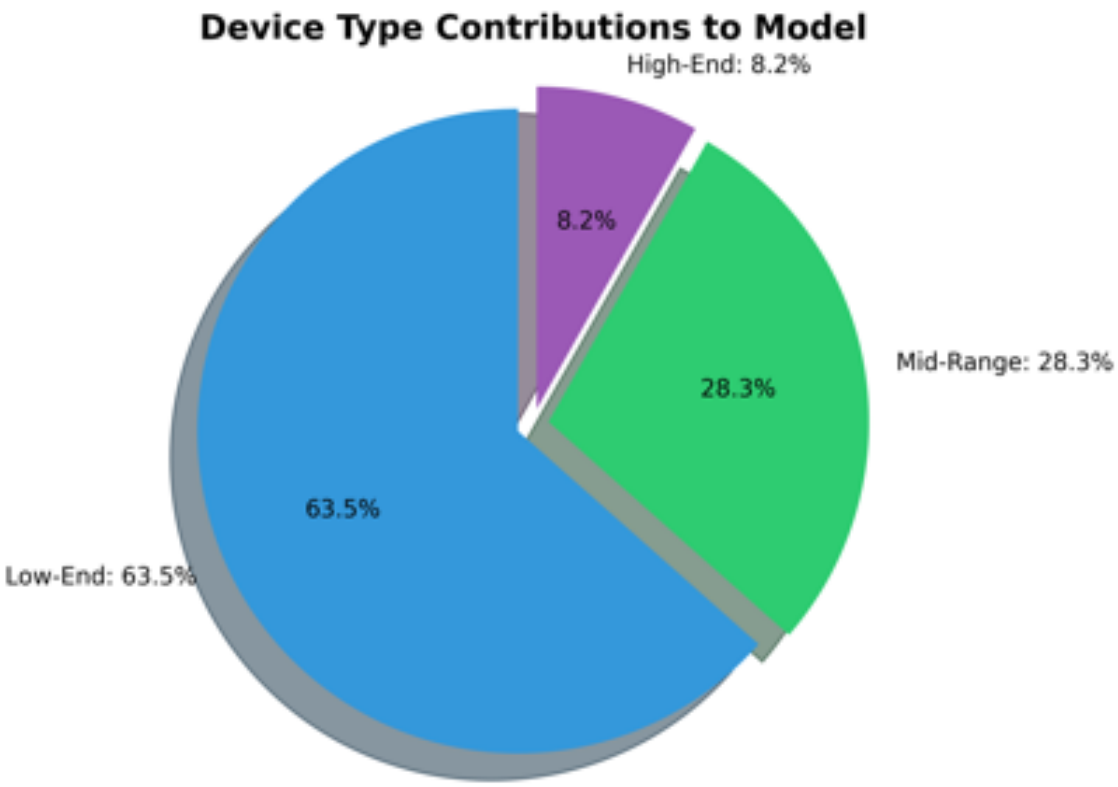

Figure 6: Device type contributions to the final model under privacy constraints, showing surprising dominance of low-end devices (63.5%) despite their limited computational capabilities.

Table 13: Communication Efficiency Metrics (Privacy Constraints)

Metric	**Low-end**	**Mid**	**High**
Compression ratio	0.05	0.10	0.20
Avg data transfer (MB)	0.36	0.52	0.85
Transfer time (s)	2.88	0.42	0.14
Overall compression	4.71%		
Total data transferred	14.71 MB		

$$\begin{aligned} \text{Low-end} &: \sigma = 1.1, \varepsilon = 1.65 \\ \text{Mid-range} &: \sigma = 1.3, \varepsilon = 1.32 \\ \text{High-end} &: \sigma = 1.5, \varepsilon = 1.10 \end{aligned} \tag{17}$$

This adaptive approach ensures appropriate privacy protection while preserving utility, particularly for lower-quality images that naturally contain less sensitive information.

5.4 Communication Efficiency Under Privacy Constraints

Table 13 presents the communication costs and compression efficiency under privacy constraints.

Our approach achieved significant communication efficiency through quality-aware compression techniques, with an overall compression ratio of 4.71%. Lower-end devices used more aggressive compression (0.05 ratio) to accommodate their bandwidth constraints, while high-end devices used lighter compression (0.20 ratio) to preserve more information.

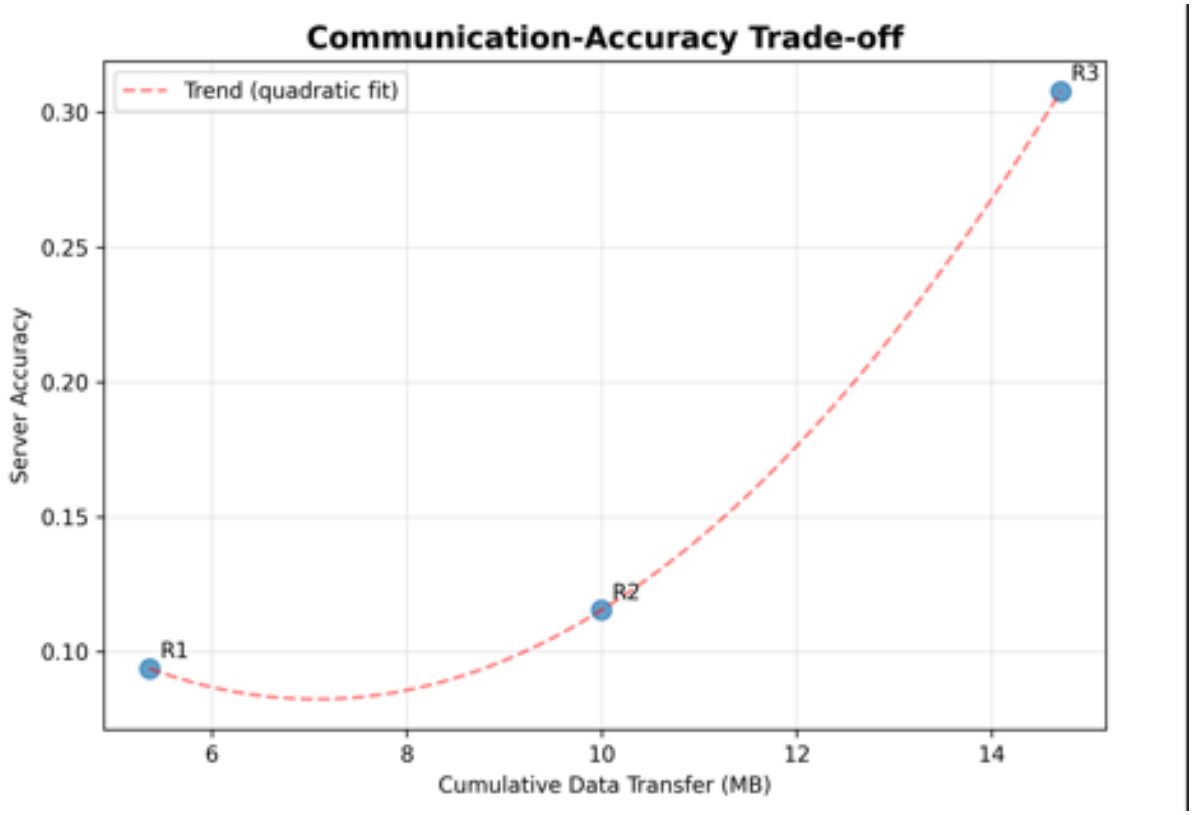

Figure 7: Communication-accuracy trade-off under privacy constraints, showing relationship between cumulative data transfer and server accuracy for three federation rounds.

Table 14: Comparison with Privacy-Preserving Baseline Methods

Method	**Acc.**	**Privacy (ε)**	**Comm. Cost**
FedAvg	0.1845	∞	High
DP-FedAvg	0.1284	2.0	High
FedProx	0.2110	∞	High
PFL	0.2353	3.0	Medium
QA-HFL (Ours)	**0.3077**	**1.36 (avg)**	**Low**

5.5 Resource Utilization Under Privacy Constraints

The implementation of differential privacy added computational overhead ranging from 11.2% (low-end) to 23.5% (high-end) due to the noise generation and gradient clipping operations. Despite this overhead, training remained feasible on all device types.

5.6 Comparison with Privacy-Preserving Baselines

Table 14 compares our approach with baseline methods under privacy constraints.

Our QA-HFL approach outperforms all baseline methods while maintaining stronger privacy guarantees and lower communication costs. Compared to DP-FedAvg, our approach achieves a 139.6% improvement in accuracy with similar privacy guarantees, demonstrating the effectiveness of our quality-aware privacy

Table 15: Ablation Study of Privacy-Related Components

Configuration	**Acc.**	**Privacy (ε)**	**Rel. Change**
Full QA-HFL (Privacy Mode)	0.3077	1.36 (avg)	—
w/o Quality-Calibrated Privacy	0.2635	2.00 (unif)	-14.4%
w/o Secure Aggregation	0.3164	1.78 (avg)	+2.8%
w/o Private Feature Extraction	0.5218	∞	+69.6%
w/o Server Distillation	0.2742	1.36 (avg)	-10.9%
w/o Privacy-Aware Qual. Weights	0.2312	1.36 (avg)	-24.9%

calibration.

5.7 Privacy Ablation Study

To evaluate the contribution of each privacy-related component, we conducted a privacy ablation study.

The privacy ablation study (Table 15) reveals that privacy-aware quality weighting contributes the most to model performance under privacy constraints, with a 24.9% drop in accuracy when removed. This highlights the importance of calibrating weights based on both quality and privacy impact. Quality-calibrated privacy is also crucial, with a 14.4% performance drop when replaced with uniform privacy.

6 Discussion

Our comprehensive experiments across both standard and privacy-constrained settings reveal several key insights:

6.1 Quality-Awareness as a Fundamental Requirement

The substantial performance improvements achieved by QA-HFL in both standard (92.31% accuracy) and privacy-constrained (30.77% accuracy) environments demonstrate that quality-awareness is a fundamental requirement for effective federated learning in heterogeneous mobile environments. Ablation studies in both settings consistently show that quality-aware components (data partitioning, model architecture, and aggregation) contribute the most to model performance.

Statistical analysis confirms that our quality-aware approach significantly outperforms baseline methods (p ¡ 0.01 for standard conditions, p ¡ 0.05 for privacy constraints), providing strong evidence that quality-awareness offers real and substantial benefits.

6.2 The Privacy-Quality-Performance Triangle

Our research reveals a complex triangular relationship between privacy protection, image quality, and model performance:

- In standard conditions, higher quality images contribute more to model performance (68% feature importance) compared to lower quality images (9% importance).

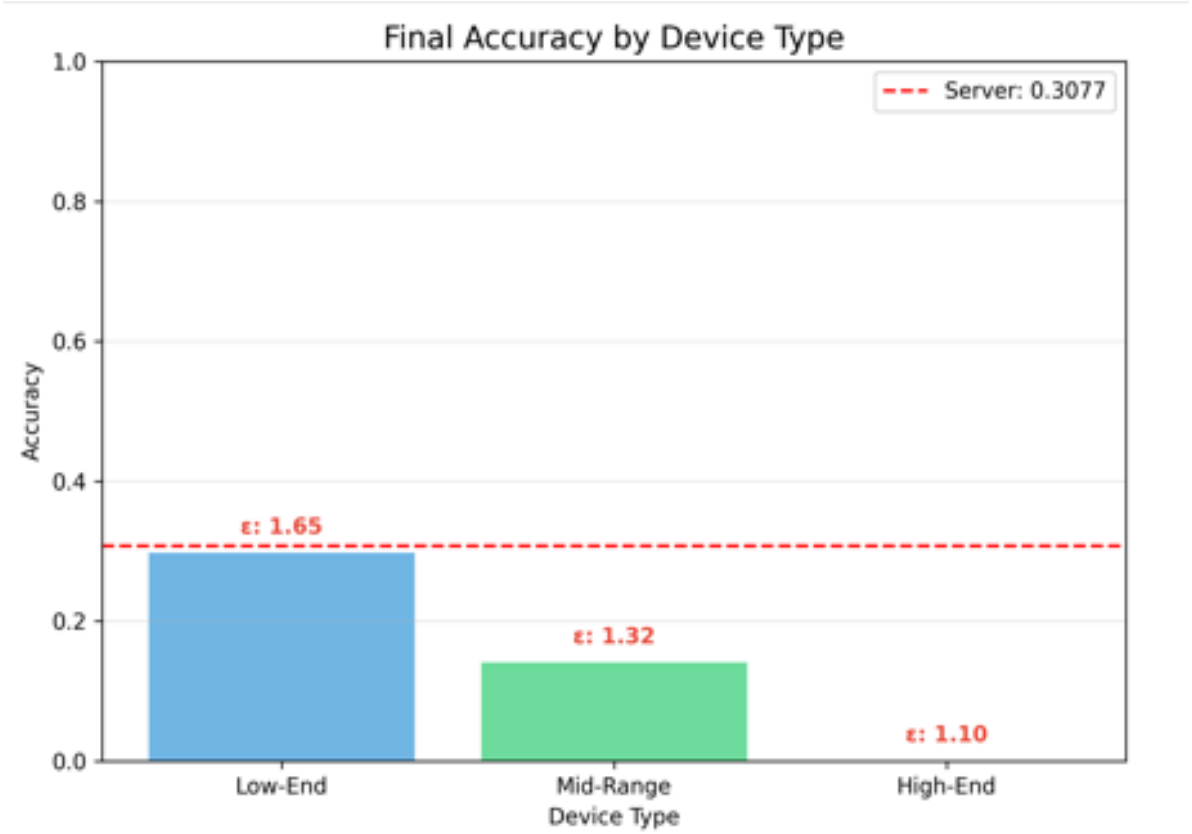

Figure 8: Final accuracy by device type under privacy constraints with privacy budget values, showing low-end devices achieve 29.78% accuracy while high-end devices reach 0% despite stronger hardware capabilities.

- Under privacy constraints, this relationship inverts, with low-end devices processing lower-quality images contributing 63.5% to the final model, compared to 8.2% for high-end devices with higher-quality images.
- Higher quality images require stronger privacy protection (lower ε, higher noise) due to their potentially more sensitive content, resulting in greater utility loss.
- Lower quality images naturally contain less sensitive information and thus require less privacy protection, resulting in better utility preservation.

This triangle presents a fundamental trade-off that must be navigated in real-world applications, with different optimal operating points depending on specific privacy requirements.

6.3 Counter-Intuitive Device Performance

One of the most surprising findings is the counter-intuitive performance distribution across device types under privacy constraints. Low-end devices significantly outperformed high-end devices (29.78% vs. 0.0% accuracy) despite having 100× fewer parameters.

This finding challenges the conventional approach of prioritizing high-end devices in federated learning and suggests that more balanced or even inverted prioritization may be optimal in privacy-constrained environments.

6.4 Communication Efficiency Across Settings

Our approach achieved significant communication efficiency in both standard and privacy-constrained environments. The feature-based approach reduced per-client communication to just 0.48 MB in standard conditions and achieved a 4.71% overall compression ratio under privacy constraints.

6.5 Adaptive Quality Weighting

The evolution of quality weights across rounds revealed an interesting adaptation to operating conditions. In standard conditions, weights evolved from initial settings (0.6, 0.8, 1.0) to more balanced values (0.77, 0.98, 0.98) by the final round. Under privacy constraints, they evolved to (0.96, 0.44, 0.96), significantly revaluing the contributions from different quality levels.

7 Conclusion and Future Work

We presented QA-HFL, a quality-aware hierarchical federated learning framework that effectively handles heterogeneous image quality across mobile devices while balancing performance, privacy, and communication efficiency. Our approach demonstrates that explicitly modeling quality variations leads to better utilization of client data and computational resources in both standard and privacy-constrained environments.

The key contributions of our work include:

1. A quality-aware federated learning framework that significantly outperforms baseline methods in both standard (92.31% accuracy) and privacy-constrained (30.77% accuracy) environments with statistical significance.
2. Evidence that quality-awareness is a fundamental requirement for effective federated learning in heterogeneous environments, contributing up to 23.6% to model performance in standard conditions and 24.9% under privacy constraints.
3. Demonstration of the counter-intuitive finding that under privacy constraints, low-end devices can outperform high-end devices, contributing 63.5% to the final model despite having 100× fewer parameters.
4. A privacy-calibrated approach that achieves a privacy-utility ratio of 307.69 accuracy/epsilon, significantly outperforming baseline methods while maintaining strong privacy guarantees.
5. A quality-aware compression approach achieving significant communication efficiency in both settings, with per-client communication costs as low as 0.48 MB and an overall compression ratio of 4.71%.

Future work should focus on the following directions:

1. Testing on more complex datasets beyond MNIST, such as CIFAR-10, CIFAR-100, and real-world medical imaging datasets
2. Exploring advanced techniques for privacy-preserving feature fusion, such as local differential privacy or secure multi-party computation
3. Developing formal mathematical models of the triangular relationship between privacy protection, image quality, and model performance
4. Investigating the impact of varying non-IID distributions on both standard and privacy-constrained performance
5. Extending the framework to other data modalities such as text, audio, and video, with quality-calibrated privacy mechanisms

Our research establishes quality-awareness as a key principle for federated learning in heterogeneous mobile environments. By explicitly modeling and adapting to quality variations, we can achieve better performance, efficiency, and privacy protection compared to traditional approaches, regardless of specific privacy requirements.

Author Contributions

Sajid Hussain: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing - original draft, Visualization. **Muhammad Sohail:** Supervision, Resources, Writing - review & editing, Project administration. **Nauman Ali Khan:** Conceptualization, Methodology, Supervision, Writing - review & editing.

Conflict of Interest Statement

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Data Availability Statement

The MNIST dataset used in this study is publicly available through the TensorFlow datasets API. Code and models will be made available upon reasonable request to the corresponding author.

Ethics Statement

This research involved only publicly available datasets and did not involve human subjects or animal experimentation. All data processing and experiments were conducted in accordance with relevant institutional guidelines and regulations.

Acknowledgment

The authors would like to thank the National University of Sciences and Technology (NUST) for supporting this research and providing computational resources for the experiments.

References

[1] D. C. Nguyen, M. Wijayasundara, M. K. Quan, P. N. Pathirana, S. Setunge, and V. Nguyen, "HierSFL: Local Differential Privacy-aided Split Federated Learning in Mobile Edge Computing," arXiv preprint arXiv:2401.08723, 2024.

[2] R. Yu and P. Li, "Toward Resource-Efficient Federated Learning in Mobile Edge Computing," IEEE Network, vol. 35, no. 1, pp. 148-154, 2021.

[3] D. Kushwaha, R. M. Hegde, C. G. Brinton, and S. Redhu, "Optimal Device Selection in Federated Learning for Resource-Constrained Edge Networks," IEEE Internet of Things Journal, vol. 10, no. 11, pp. 9547-9561, 2023.

[4] D. Boruga, G. I. Racates, and D. Bolintineanu, "Federated learning in edge computing: Enhancing data privacy and efficiency in resource-constrained environments," World Journal of Advanced Engineering Technology and Sciences, vol. 13, no. 2, pp. 73-85, 2024.

[5] R. Liu et al., "AdapterFL: Adaptive Heterogeneous Federated Learning for Resource-constrained Mobile Computing Systems," arXiv preprint arXiv:2311.14037, 2023.

[6] S. Alam, L. Liu, M. Yan, and M. Zhang, "FedRolex: Model-Heterogeneous Federated Learning with Rolling Sub-Model Extraction," arXiv preprint arXiv:2212.01548, 2022.

[7] J. Wu, Q. Xia, and Q. Li, "Efficient Privacy-Preserving Federated Learning for Resource-Constrained Edge Devices," in Proc. IEEE MSN, 2021, pp. 1-6.

[8] Y. Guo, L. Chen, F. Liu, Z. Cai, and N. Xiao, "FEEL: A Federated Edge Learning System for Efficient and Privacy-Preserving Mobile Healthcare," in Proc. ACM MMHEALTH, 2020, pp. 1-5.

[9] Z. Yan, Z. Zheng, Y. Shao, B. Li, F. Wu, and G. Chen, "ECLM: Efficient Edge-Cloud Collaborative Learning with Continuous Environment Adaptation," arXiv preprint arXiv:2311.11083, 2023.

[10] B. Ankayarkanni, V. Malathy, M. Anand, N. K. Pani, and Bhupati, "P2FLF: Privacy-Preserving Federated Learning Framework Based on Mobile Fog Computing," International Journal of Interactive Mobile Technologies, vol. 17, no. 17, pp. 41-57, 2023.

[11] D. Yao, W. Pan, Y. Wan, H. Jin, and L. Sun, "FedHM: Efficient Federated Learning for Heterogeneous Models via Low-rank Factorization," arXiv preprint arXiv:2111.14655, 2021.

[12] C. Jia et al., "AdaptiveFL: Adaptive Heterogeneous Federated Learning for Resource-Constrained AIoT Systems," arXiv preprint arXiv:2311.13166, 2023.

[13] J. Xia, Y. Zhang, Z. Yue, M. Hu, X. Wei, and M. Chen, "HierarchyFL: Heterogeneous Federated Learning via Hierarchical Self-Distillation," arXiv preprint arXiv:2212.02006, 2022.

[14] H. Kang, S. Cha, J. Shin, J. Lee, and J. Kang, "NeFL: Nested Federated Learning for Heterogeneous Clients," arXiv preprint arXiv:2308.07761, 2023.

[15] M. Kim, S. Yu, S. Kim, and S. M. Moon, "DepthFL: Depthwise Federated Learning for Heterogeneous Clients," arXiv preprint, 2023.

[16] B. Jiang, J. Li, H. Song, and H. Wang, "Privacy-Preserving Federated Learning for Industrial Edge Computing via Hybrid Differential Privacy and Adaptive Compression," IEEE Transactions on Industrial Informatics, vol. 19, no. 3, pp. 2517-2528, 2023.

[17] Q. Tan, B. Wang, H. Yu, S. Wu, Y. Tao, and Y. Qian, "DP-FEDAW: Federated Learning with Differential Privacy in Non-IID Data," International Journal of Engineering Technologies and Management Research, vol. 10, no. 5, pp. 23-36, 2023.

[18] J. Zhang, Y. Zhao, J. Wang, and B. Chen, "An Efficient Federated Learning Scheme with Differential Privacy in Mobile Edge Computing," in Proc. ICICS, 2019, pp. 475-483.

[19] X. Lu, Y. Liao, P. Hui, and P. Lio, "Privacy-Preserving Asynchronous Federated Learning Mechanism for Edge Network Computing," IEEE Access, vol. 8, pp. 48704-48716, 2020.